

INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

**ProQuest Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
800-521-0600**

UMI[®]

BOSTON UNIVERSITY
COLLEGE OF ENGINEERING

Dissertation

**PRICING AND RESOURCE ALLOCATION IN COMMUNICATION
NETWORKS AND SUPPLY CHAINS**

by

YONG LIU

B.E., Xi'an Jiaotong University, Xi'an, China, 1994
M.E., Xi'an Jiaotong University, Xi'an, China, 1997

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2002

UMI Number: 3040711

UMI[®]

UMI Microform 3040711

**Copyright 2002 by ProQuest Information and Learning Company.
All rights reserved. This microform edition is protected against
unauthorized copying under Title 17, United States Code.**

**ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346**

Approved by

First Reader



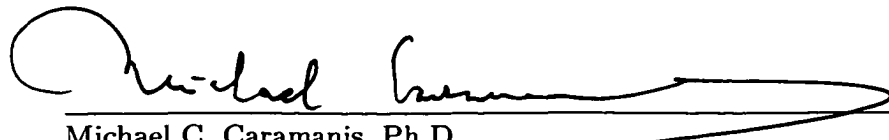
Ioannis Ch. Paschalidis, Ph.D.
Associate Professor of Manufacturing Engineering

Second Reader




Christos G. Cassandras, Ph.D.
Professor of Manufacturing / Electrical and Computer Engineering

Third Reader



Michael C. Caramanis, Ph.D.
Professor of Manufacturing Engineering

Fourth Reader



Jian-Qiang Hu, Ph.D.
Associate Professor of Manufacturing Engineering

Acknowledgements

First of all, I am so much grateful to my Ph.D. advisor, Professor Ioannis Paschalidis. Throughout all the time of my Ph.D. study and the writing of this thesis, he provided not only good teaching, sound advice, and lots of good ideas, but also caring, encouragement, and support. It is difficult to overstate my gratitude to him and this work would not have come into existence without him.

I am very grateful to Professor Christos Cassandras, Professor Michael Caramanis and Dr. Christos Panayiotou for extensive discussions and their suggestions on the relevant research. I also would like to thank Professors Jian-Qiang Hu, Professor Pirooz Vakili and Professor James Perkins for their help during my studies at Boston University.

I thank Hui Huang and Ping Zhang for their help and discussions on some topics in the thesis.

I am indebted to many student colleagues for providing a stimulating and fun environment in which to learn and grow. I am especially grateful to Chang Shu, Chang Su, Bei Zhou, Rui Yu, Xima Zhang, Zhen Wang, So-Yeon Baek, Keiko Matsunaga, Yi Zou, Liwu Hao, Jianyang Tai, Yanfeng Wang, Haidong Pan, Osman Anli, and Spyridon Vassilaras. I am also grateful to my friends in Xi'an Jiaotong University, Changgeng Xi, Chuhong Fei, Zhaohui Wang, and my brother Bin Liu for all the emotional support.

Lastly, and most importantly, I wish to thank my parents, Wanyu Liu and Lixin Chen. They raised me, taught me, supported me, and loved me. I dedicate this thesis to them.

class to serve as a substitute of another. For such networks, under certain conditions, we also show that static pricing is asymptotically optimal in the same regime of many small users.

Regarding problem (ii), we study QoS-capable supply chains consisting of a tandem of production facilities (stages). Unsatisfied external demand is backlogged. We quantify QoS by the stockout probabilities at various stages. We propose production policies in two separate cases: when each stage (a) has only local inventory information, and (b) has knowledge of the total downstream inventory. In case (a) the proposed policy guarantees service level requirements. In case (b) the proposed policy minimizes expected inventory costs subject to QoS constraints. In both cases policy parameters are obtained analytically, based on large deviations asymptotics, which leads to drastic computational savings compared to simulation. Our model can accommodate autocorrelated demand and production processes, both critical features of modern manufacturing systems. We demonstrate that detailed distributional information on demand and production processes, which is incorporated into large deviations asymptotics, is critical in inventory control decisions.

Contents

Acknowledgements	iii
Abstract	iv
Table of Contents	vi
List of Tables	x
List of Figures	xiii
1 Introduction	1
1.1 Motivation and Literature Review	1
1.1.1 Pricing in Multiservice Communication Networks	2
1.1.2 Inventory Control in Supply Chains	5
1.2 Contributions	7
1.2.1 Pricing in Multiservice Communication Networks	7
1.2.2 Inventory Control in Supply Chains	9
1.3 Organization of the Dissertation	11
1.4 Preliminaries	13
1.4.1 Dynamic Programming	13
1.4.2 Large Deviations	17
2 Pricing in Communication Systems: Single-Link Case	22
2.1 Problem Formulation	22
2.1.1 Revenue Maximization Problem	24

2.1.2	Welfare Maximization Problem	24
2.2	Optimal Dynamic Policy	25
2.2.1	Dynamic Programming Formulation	25
2.2.2	Some General Properties	27
2.2.3	Welfare Maximization Case	28
2.3	Static Pricing Policy	28
2.4	Price Aggregation Approach for Single-class Case	31
2.5	An Upper Bound on Optimal Revenue	35
2.6	Approximate DP Approach	37
3	Pricing in Communication Systems: The Network Case	41
3.1	The Model of Multi-Service Networks	41
3.1.1	Revenue Maximization Problem	43
3.1.2	Welfare Maximization Problem	43
3.2	Optimal Dynamic Policy	44
3.3	Static Pricing Policy	45
3.4	An Upper Bound on Optimal Performance	47
3.5	Asymptotic Optimality of Static Pricing	48
3.6	Structure of Asymptotically Optimal Static Pricing Policy	55
3.6.1	Revenue Maximization	55
3.6.2	Welfare Maximization	56
3.7	Large Scale Problems	57
3.7.1	A Simulation-based Method	58
3.7.2	Numerical Results	62
4	Pricing in Communication Networks with Demand Substitution Effects	67
4.1	The Model with Demand Substitution Effects	67
4.2	An Upper Bound on Optimal Performance	70
4.3	Asymptotic Optimality of Static Pricing	71

4.4	Structure of the Asymptotically Optimal Static Policy	78
4.5	Numerical Results	79
5	Inventory Control for Single-Stage Make-to-Stock System	80
5.1	The Model	80
5.2	Large Deviations Analysis	82
5.2.1	Stockout Probability	83
5.2.2	Inventory Cost	85
5.3	Numerical Results	86
6	Inventory Control in Supply Chain Management: The Local Inventory Case	88
6.1	The Model	88
6.2	The Decomposition Approach	91
7	Inventory Control in Supply Chain Management: The Multi-echelon Approach	97
7.1	The Multi-echelon Approach: A Global Information Case	98
7.2	Large Deviations Analysis of the Stockout Probability	99
7.2.1	Lower Bound	104
7.2.2	Upper Bound	106
7.2.3	Upper and Lower Bounds Match	110
7.3	Refining the Large Deviations Asymptotics	115
7.4	Approximating the expected inventory cost	120
7.5	Extensions: The Multiclass Case and Lost Sales	125
7.5.1	The Multiclass Case	125
7.5.2	A Model with Lost Sales	127
7.6	Numerical Results	134
7.6.1	A Two-Stage Supply Chain with the Decomposition Approach	134

7.6.2	A Two-Stage System with Multi-Echelon Approach	136
7.6.3	A Three-Stage System with Multi-Echelon Approach	140
7.6.4	Significance of Distributional Information	143
8	Summary and Future Research	145
8.1	Pricing in Multiservice Communication Networks	145
8.2	Inventory Control in Supply Chains	147
	Bibliography	149
	Curriculum Vitae	154

List of Tables

2.1	Numerical results of the optimal static policy for the two-class problem ($C = 155, r_1 = 4, r_2 = 1, \mu_1 = 1, \mu_2 = 2$).	30
2.2	Numerical results of price aggregation for single-class problem ($C = 30, r = 1$), with uniform price interval.	34
2.3	Numerical results of price aggregation for single-class problem ($C = 30, r = 1$), with varied price interval.	35
2.4	Numerical results of the upper bound for single-class problems ($C = 30, r = 1$).	37
2.5	Numerical results of the upper bound for two-class problems ($C = 155, r_1 = 4, r_2 = 1, \mu_1 = 1, \mu_2 = 2$).	37
2.6	Numerical results of the approximate DP method for some two-class problems. J^* is the optimal revenue derived by value iteration, \bar{J} is the upper bound of the problem obtained by the way in Section 2.5, J_{simu} is the simulation result with the policy obtained from approximate DP.	40
3.1	The services provided by the network of Figure 3.2.	63
3.2	Comparing the various policies for the network of Figure 3.2. We use $J(\mathbf{u}_{\text{ub}}^*)$ and J_{sim} to denote the performance of the policy obtained from the upper bound problem and the simulation-based optimization approach, respectively.	64
3.3	The prices for the network of Figure 3.2 under the policy obtained from the upper bound problem (\mathbf{u}_{ub}^*) and the policy obtained from the simulation-based optimization approach ($\mathbf{u}_{\text{sim}}^*$).	64
3.4	The services provided by the network of Figure 3.3.	65

3.5	Comparing the various policies for the network of Figure 3.3. We maintain the notation of Table 3.2.	65
4.1	The services with demand substitution effects provided by the network in Figure 3.2.	79
4.2	Comparing the various policies for the network of Section 4.5. $J_{u_{ub}^*}$ and J_{sim} denote the performance of the policy obtained from the upper bound problem and the simulation-based optimization approach, respectively.	79
4.3	The prices for the network of Section 4.5 under the policy obtained from the upper bound problem (u_{ub}^*) and the simulation-based optimization approach (u_{sim}^*).	79
5.1	Comparing the analytical the LD expression for the stockout probability with the simulated values. The first column lists the required stockout probabilities, ranging from 0.3 to 10^{-6} . The second column provides the hedging points obtained by the LD expression. The next two columns provide the hedging points used in the simulation and the stockout probability obtained.	87
7.1	Numerical results from the decomposition approach for Example 1. The simulated values for the expected shortfalls, that are used in computing the prefactors c_i , are $E[L^1] = 9.297$ and $E[L^2] = 1.098$. Expected inventory costs are reported in Table 7.4.	135
7.2	Numerical results from the decomposition approach for Example 2. The simulated values for the expected shortfalls, that are used in computing the prefactors c_i , are $E[L^1] = 2.420$ and $E[L^2] = 8.216$. Expected inventory costs are reported in Table 7.6.	136

7.3	<p>Numerical results for Example 1 operated under the multi-echelon policy. We denote by \mathbf{w}_A^* (3rd column) the hedging vector obtained by solving the optimization problem in (7.47). Similarly, \mathbf{w}_S^* (7th column) denotes the hedging vector obtained by brute-force simulation over integer points. The 4th column ($\mathbf{E}[C]$) lists the optimal value of the optimization problem in (7.47), that is, our analytical approximation of the total expected inventory cost of the policy in column 3. Column 5 and 6 list the stockout probability and expected inventory cost, respectively, obtained by simulating the policy of column 3. Column 8 and 9 list the stockout probability and expected inventory cost, respectively, obtained by simulating policy of column 7. . . .</p>	138
7.4	<p>The 3rd and 4th column report the hedging points obtained by the decomposition approach. We simulated the system to obtain $\mathbf{E}[(I^1)^+]$, $\mathbf{E}[I^2]$, $\mathbf{P}[I^1 \leq 0]$, and $\mathbf{E}[\text{Cost}]$. To make comparisons with the results in Table 7.3 note that holding costs h_1, h_2 for echelon 1 and 2, respectively, correspond to holding costs $h_1 + h_2$ and h_2 for stage 1 and 2, respectively, in the decomposition approach.</p>	139
7.5	<p>Numerical results for Example 2 operated under the multi-echelon policy. The notation and the structure of the table are the same as in Table 7.3. . .</p>	141
7.6	<p>Comparing the multi-echelon policy with the decomposition policy in Example 2.</p>	141
7.7	<p>Numerical results for Example 3 operated under the multi-echelon policy. The notation and the structure of the table are the same as in Table 7.3. . .</p>	143

List of Figures

2.1	The dynamic policy for a system with total bandwidth $C = 30$, which provides a single-class service ($r = 1$), and $\lambda_0 = 60$, $\lambda_1 = 5$, $\mu = 1$, (a) by price aggregation with uniform interval; (b) by price aggregation with varied interval; (c) by value iteration for DP (optimal policy).	36
2.2	The numerical results for Example 1. The control space \mathcal{U} is discretized into $L = 50$ intervals. The revenue obtained from the approximate policy is 161.24, while the optimal revenue is 167.68.	39
2.3	The numerical results of the approximate DP for Example 2. The control space \mathcal{U} is discretized into $L = 20$ intervals. The revenue obtained from this approximate policy is 920.529. The optimal revenue is 952.63.	40
3.1	$\lambda_i(u_i)$ and $\hat{\lambda}_i(u_i)$ of Equations (3.26) and (3.27), respectively, when $\lambda_{i,0} = 40$, $\lambda_{i,1} = 4$, and $\epsilon = 1$	61
3.2	A 5-node, 4-link network with 12 service classes.	63
3.3	A 9-node, 13-link network with 59 service classes. The link capacities are $C_1 = 840$, $C_2 = 420$, $C_3 = 420$, $C_4 = 420$, $C_5 = 420$, $C_6 = 300$, $C_7 = 420$, $C_8 = 420$, $C_9 = 420$, $C_{10} = 420$, $C_{11} = 420$, $C_{12} = 210$, and $C_{13} = 420$	66
5.1	The model of a make-to-stock system.	81
5.2	The largest root of $\Lambda_D^+(\theta) + \Lambda_B(-\theta) = 0$	84
5.3	The models of the demand and production processes.	86
6.1	The model of the supply chain.	88

6.2	The equivalent $G/G/1$ queue of stage i , $i = 1, \dots, M$, in a decoupled multi-stage supply chain.	93
7.1	The supply chain model under the echelon base-stock policy.	98
7.2	The paths for each stage at time slot 1 (one-level graph).	101
7.3	The paths of each stage at time slot 2 (two-level graph).	102
7.4	We depict the feasible set \mathcal{B} in a three-dimensional example. The set \mathcal{B} consists of three regions: region I, where $\theta_{G,1}^* = \theta_1^*$, region II, where $\theta_{G,1}^* = \beta_1 \theta_2^*$, and region III, where $\theta_{G,1}^* = \beta_2 \theta_3^*$. As sample points we select points $Q_1, Q_2, Q_3, Q_4, Q_5, Q_6, Q_7$, and potentially some additional points scattered in regions I–III.	117
7.5	We consider a system with two stages and assume that $\theta_1^* > \theta_2^*$. For a fixed w_1 , $f_1(w_1, \beta_1)$ is a function of β_1 . A, B, C, D, E, F are sample points. Suppose that we do not know $f_1(w_1, \frac{\theta_1^*}{\theta_2^*})$. We can approximate it by extrapolating segment BC and DE as shown in the graph. To maintain the continuity of $f_1(w_1, \beta_1)$ and convexity in region $[1, \theta_1^*/\theta_2^*]$ and $[\theta_1^*/\theta_2^*, \infty)$, we select Q as $f_1(w_1, \theta_1^*/\theta_2^*)$. For points to the right of F , we set them to be constant and equal to $f_1(w_1, \beta_F)$	118
7.6	The models of demand and production processes in Example 1, a two-stage system. We denote by \mathbf{r} the vector of demand or production amounts at each state of the corresponding Markov chain.	134
7.7	The optimal multi-echelon policies for Example 2 derived by simulation, where the service level requirement is $\epsilon_1 = 10^{-2}$. (a) When $h_1 = 1$ and $h_2 = 1$, $w_1^* = 22$, $w_2^* = 54$, $\mathbf{E}[\text{Cost}] = 64.04$; (b) when $h_1 = 5$ and $h_2 = 1$, $w_1^* = 18$, $w_2^* = 61$, $\mathbf{E}[\text{Cost}] = 129.22$. In each graph, the thick curve is the boundary of the set of the feasible policies (i.e., set of vectors (w_1, w_2) satisfying the service level constraints), the optimal policy obtained by our analytical approach is marked with a circle.	142

7.8	The models of demand and production processes in Example 3.	142
7.9	For the example of Section 7.6.4, we plot $\Lambda_D(\theta) + \Lambda_{B^1}(-\theta)$, $\Lambda_D(\theta) + \Lambda_{B^2}(-\theta)$ and identify the corresponding largest positive roots θ_1^* and θ_2^* , respectively.	144

Chapter 1

Introduction

1.1 Motivation and Literature Review

We consider pricing and resource allocation decisions in communication and supply networks that can provide *Quality of Service (QoS)* guarantees. In general, such systems exhibit strong stochasticity and are typically referred to as *stochastic networks*. Demand for the communication services or for the products of supply chains are random. Packet generation processes in the Internet are very bursty, while at various stages of a supply chain, the capacities of production and distribution facilities could also change over time. Recent developments have emphasized the importance of QoS. In communication networks, the emergence of real-time traffic (e.g., Internet telephony, video conferencing, video-on-demand) imposes rather stringent QoS requirements. In supply chains, increasing competition has led to the desire to devise production/distribution policies that guarantee a certain level of QoS. To adequately quantify QoS one needs to use performance metrics that are hard to analyze. In communication networks, such metrics include the probabilities of packet loss or delay, and the blocking probability of connection requests. In supply chains, such metrics include the stockout probability or the probability of exceeding a promised delivery time. The optimal allocation of available resources is an ageless problem that has found increased importance in today's competitive market environment. In this thesis, we will focus into two instances of resource allocation problems in stochastic networks: (1) pricing in multiservice communication networks, and (2) inventory control in supply chains.

In both cases the complexity of the problem will not allow us to pursue exact analysis.

Instead, we will employ powerful asymptotic techniques to analyze the problems in well-motivated limiting regimes. Based on our asymptotic results we will develop policies that are efficient even away from the limiting regimes.

1.1.1 Pricing in Multiservice Communication Networks

The Internet is a worldwide network of computer networks that use a common communication protocol, TCP/IP . In the late 1960s, the Advanced Research Projects Administration (ARPA), a division of the U.S. Defense Department, developed the ARPAnet to link together universities and high-tech defense contractors. In the mid-1980s, the National Science Foundation (NSF) created the NSFNET in order to provide connectivity to its supercomputer centers and to provide other general services. The NSFNET evolved directly out of ARPAnet and adopted the TCP/IP protocol to form a high-speed backbone for the developing Internet. Following that, the Internet grew rapidly, in terms of number of hosts, number of users, traffic and also technology, and eventually NSF decided that its operation could be more effectively administered by the private sector. In April 30, 1995 NSFNET ceased operation as its NSF funding ended. Now traffic in the United States is carried on several backbones operated by private, for-profit enterprises (see [MB98]). In this private setting, pricing is becoming increasingly important as Internet service providers need to develop efficient pricing schemes to recover costs and fund future expansions.

The most frequently used Internet applications are electronic mail (e-mail), file transfer and remote login. However, the converging digital technologies of publishing, telephony, television and computers are transforming the Internet into a unified integrated-service network. Digital video, audio, and interactive multimedia communication services are growing in popularity and have the potential to affect all sectors of society worldwide. The migration to an integrated-service network will have important implications for market structure and competition.

The current Internet offers a single service quality: “best effort” service. Packets are transported first-come, first-served with no guarantee of success. Some packets may ex-

perience severe delays, while others may be dropped and never arrive. However, different streams of data place different demands on the network. E-mail and file transfer require 100% data integrity, but can tolerate delay. Real-time voice broadcasts typically require higher bandwidth than file transfers, can only tolerate minor delays, but they can tolerate significant distortion. Real-time video broadcasts have low tolerance for delay and distortion. Because of different requirements, network efficiency can be increased if the different types of traffic are treated differently.

Until recently, nearly all users faced the same pricing structure for using the Internet. Local ISPs rent a fixed bandwidth connection from backbone ISPs and were charged an annual fee, which allowed for unlimited usage up to physical maximum flow rate. The end users were also charged a flat monthly fee for the Internet service provided by their local ISPs. Without an incentive to economize on usage, congestion can become quite serious. If everyone just stuck to ASCII e-mail congestion would not likely become a problem. However the demand for multimedia services is growing dramatically. Although administratively assigning different priorities to different types of traffic is appealing, it might be impractical as a long-run solution to reducing congestion because of the usual inefficiency of rationing, and because it is technically hard to implement.

An alternative approach for reducing congestion is overprovisioning. Overprovisioning means maintaining sufficient network capacity to support the peak demands without noticeable service degradation. This has been the most important mechanism used to date on the Internet. However, overprovisioning is costly. Given the explosive growth in demand and the long lead time needed to introduce new network infrastructure and protocols, the Internet may face serious problems very soon. Therefore it is time to seriously examine incentive-compatible allocation mechanisms, such as various forms of congestion pricing. We believe that pricing of network services is becoming increasingly important in this new environment. Apart from allowing providers to recover their operating expenses and fund future capacity expansions, it can lead to more efficient use of the network resources by providing sufficient incentives to users. Perhaps more importantly, it enables the creation of

a healthy market environment, where new network services can be (profitably) introduced and sustained.

Pricing in communication networks has received a lot of attention in the literature. MacKie-Mason and Varian [MMV94] proposed a “smart market” where individual packets bid for transport while the network only serves packets with bids above a certain cutoff amount, depending on the level of congestion. Kelly et al. [Kel97, KMT98] consider charges that increase with either realized flow rate or with the “share” of the network consumed by a traffic stream. Several researchers have looked at packet-based pricing schemes as an incentive for more efficient flow control (see e.g., Gibbens and Kelly [GK99], La and Anantharam [LA00], Kunniyur and Srikant [KS00]). Equilibrium properties of bandwidth and buffer allocation schemes are analyzed by Low [Low00]. Clark [Cla97] proposes an expected capacity-based pricing scheme where users are charged ahead of time on the basis of the expectation that they have of network usage and excess packets are dropped at times of congestion.

The emergence of real-time traffic substantially complicates the picture and requires QoS measures much harder to analyze (see [Kel96, BPT98a, BPT98b, Pas99]). The network model we will consider is more appropriate for real-time traffic that requires strict *Quality of Service (QoS)* guarantees. Such guarantees can often be translated into a preset resource amount that has to be allocated to a call at all links in its route through the network. If the resource is bandwidth this resource amount can be some sort of an *effective bandwidth* (see e.g., Kelly [Kel96] for a survey of effective bandwidth characterizations and Paschalidis [Pas99] for similar notions in a multiclass setting). In this setting, Kelly [Kel94] and Courcoubetis *et al.* [CKW97] propose the pricing of real-time traffic with QoS requirements, in terms of its effective bandwidth, and provide approximations that only involve time and volume charges.

In this dissertation, we consider the pricing problem in a multi-service communication network with fixed routing. We propose pricing strategies that aim at two distinct objectives: either maximizing the revenue of the network operator or maximizing the social

welfare of the users. For the case of revenue maximization, we have essentially a problem of *yield management*, or *revenue management*. Yield management can be described as the process of efficiently utilizing the limited available resources through pricing or other types of actions. It is widely practiced in capacity-constrained service industries, such as airlines, hotels and car rentals (see [SLD92, Ash97, BM95, Kim89a, Kim89b, CG95]).

Technically, the revenue maximization problem that we study is structurally similar to the work in the revenue management of airlines (see, [SSL97, GvR97]). The latter problem though is typically formulated as a finite horizon (e.g., [GvR97]), which is different than infinite horizon average cost setup. Our work is also related to problems of admission control in loss networks (see [Key90, OK92, Ros95, IS01]). This literature, however, assumes that the prices are fixed and is only concerned with admission decisions, while we wish to study optimal or near-optimal pricing schemes. Also, we use a decision-theoretic framework under an explicit model of users' reaction to prices (demand functions). Similar demand functions have been used in [LV93] under a somewhat different model.

1.1.2 Inventory Control in Supply Chains

Our focus thus far has been on the technology, costs and pricing of the network transport. However most of the value of the network is not in the transport, but in the value of the information being transported. The advent of the Internet and the trend of globalization radically transformed the manner in which business and supply chains are being managed today. Manufacturing has recently gone through significant restructuring. A recent survey [eco98] emphasized that "no factory is an island." Companies are becoming more global. They consist of factories, suppliers, distributors, and customer service centers scattered around the world. As a result, modern manufacturing enterprises have recognized that production can not be viewed separately from the physical distribution of goods. Instead, both activities should be perceived as indispensable parts of a *supply chain*. Fortunately for the U.S. and other industrialized countries, the higher-value parts of these new global supply chains tend to stay within their borders. In these countries, it is becoming more

important to manage supply chains, develop technology, and attend to customer needs; labor intensive operations migrate to places where wages are smaller. At the same time, the explosion of e-commerce has created a host of new companies that need to efficiently manage inventory in supply chains.

An additional reason that led to this integrated view of manufacturing enterprises, but also, a key defining characteristic of the new environment, is that manufacturing is becoming more *customer oriented*. In an era of increased competition, customers are more demanding and require products delivered in a timely manner wherever they happen to be located. In addition to product functionality, companies are recognizing the significance of *Quality of Service (QoS)* in acquiring and maintaining market share. E-commerce companies in particular, have been more adept at adopting new practices and emphasizing QoS.

Our research objective is to develop effective policies for inventory control in supply chains that address the difficulties present in the new manufacturing environment. The fundamental trade-off in inventory control is between *producing*, which accumulates inventory and incurs inventory costs, and *idling*, which leads to stockouts and unsatisfied demand. A *production policy* resolves this trade-off and determines at each point in time whether the production facilities at all stages of the supply chain should be producing or idling.

There is a large literature on production inventory systems (see Kapuscinski and Tayur [KT99] for a survey). The single-stage, single-class, version of the problem is significantly simpler. It has been shown in a variety of settings that a so called *base-stock* policy (produce when inventory falls below a certain level and idle otherwise) is optimal (see Evans [Eva67], Gavish and Graves [GG80], Sobel [Sob82], Federgruen and Zipkin [FZ86], Akella and Kumar [AK86], and Kapuchinski and Tayur [KT98]). In multiclass single-stage systems the optimal policy is not in general known. In these systems a production policy involves both idling and scheduling decisions (deciding on which classes to work on, if any). There have only been results for special cases (Zheng and Zipkin [ZZ90], Ha [Ha97], de Véricourt, Karaesmen and Dallery [dVKD00]) or approximations and heuristics for the general case (Wein [Wei92], Peña-Perez and Zipkin [PPZ97], Veatch and Wein [VW96], Glasserman

[Gla96], and Bertsimas and Paschalidis [BP01]). In a multiple-stage, single-class system, and *without capacity limits*, Clark and Scarf in their seminal paper [CS60] have shown the optimality of a production policy where each facility follows a base-stock policy based on the total inventory available in the downstream facilities (we will refer to this as *echelon* inventory). Their result has been generalized in several directions (Federgruen and Zipkin [FZ84], Chen and Song [CS01]). In the more general case where capacity limits exist and demand and service processes are autocorrelated, such a policy is not necessarily optimal. However the simplicity of its structure makes it attractive. Under a similar echelon policy Glasserman and Tayur [GT95] proposed a perturbation analysis approach to compute the hedging points in a capacitated single-class multi-stage system, and Glasserman [Gla97] has developed asymptotics to approximate stockout probabilities under renewal demand and constant production capacities.

1.2 Contributions

In this dissertation, we focus into two instances of resource allocation problems in stochastic networks: (i) revenue or welfare maximization in QoS-capable communication networks, and (ii) inventory control in supply chains subject to given QoS requirements. Because of the complexity of the problems, analyzing them exactly is intractable. Instead, we employ asymptotic techniques.

1.2.1 Pricing in Multiservice Communication Networks

We consider the pricing problem in a multi-service communication network with fixed routing. Different classes differ in bandwidth requirements, demand pattern, call duration, and routing. Links in the network have given finite capacities and the total resource requirement of all calls using a link cannot exceed the link's capacity. The network charges a fee per call which can depend on the current congestion level, and which affects user's demand for calls.

Our work is closer to Paschalidis and Tsitsiklis [PT00] that considered pricing of multiple services sharing a single resource. The single-link problem can be formulated as an infinite horizon, average reward dynamic programming (DP) problem. But for large-scale problems, i.e., when the system has many resources and provides multiple classes of services, solving the DP numerically becomes impractical. We develop a number of approximation approaches, such as price aggregation and approximate DP, by which we can obtain sub-optimal dynamic pricing policies. Furthermore, we generalize the main result of [PT00] in a network setting. In particular, we show that in a limiting regime of “many small users,” laws of large number take effect and a simple *static* pricing scheme is asymptotically optimal. That is, under stationarity assumptions, prices can remain fixed (distinct for various classes) and it is not necessary to employ a dynamic scheme according to which prices depend on the current congestion level. If demand is nonstationary and is characterized by time-of-day effects, which is widely agreed to be the case in communication networks, the proposed pricing scheme leads to *time-of-day pricing*. The “many small users” limiting regime we consider is quite appropriate for large networks such as the Internet where (backbone) capacities are large and individual sessions or calls occupy a small fraction of those capacities.

A static pricing scheme, such as the one we propose, has obvious implementation advantages: charges are predictable by users, evolve in a slower time-scale than congestion phenomena, and no elaborate real-time mechanism is needed to communicate prices to the users. Moreover, as we will see, prices can be computed in large-scale systems, which is not the case with the optimal dynamic pricing scheme. To that end, from our asymptotic optimality results we first identify an insightful, asymptotically optimal, structure of static prices under both revenue and welfare maximization objectives. According to this structure prices depend on a parsimonious number of parameters. We then employ a simulation-based optimization technique to tune those parameters. We report results from a number of numerical experiments, including, a large-scale one, indicating that this approach yields near-optimal policies.

We characterize the rate at which a static pricing policy converges to optimal in the regime of many small users. This allows us to obtain bounds on the suboptimality gap of static pricing away from the limiting regime. We provide examples where such bounds are useful in quickly assessing efficiency gains achieved by appropriately scaling the system.

We also investigate demand substitution effects. Namely, when the price for one class of service is too high, some of its customers may choose another classes as a non-perfect alternative. Hence, the price of a service will affect not only its own demand, but also demand for other services. Courcoubetis and Reiman [CR99] discuss calculating an optimal static pricing policy in the limiting regime of “many small users” in a single link system with substitution effects. We substantially extend their framework in a network setting, and show that in the limiting regime of many small users, an appropriate static pricing policy will be asymptotically optimal for both revenue and welfare maximization problems. Furthermore, the prices have a similar structure with the simpler case that does not take substitution effects into account. A simulation-based optimization approach is again applicable to tune those parameters and yield near-optimal policies.

The network model we propose is general enough to accommodate several situations of practical interest. It can be seen as modeling the pricing of bandwidth by a network provider who offers a menu of services to users. Users can in fact also be smaller “retail” providers, in which case calls can be seen as virtual circuits leased from the backbone provider. The model can also be seen as pricing the use of Web or other servers by an application service provider: a “call” is associated with a transaction that requires cooperation from a series of servers, thus, it ties up a fraction of their capacities until it is completed.

1.2.2 Inventory Control in Supply Chains

As for the inventory control in supply chain management, we propose and analyze two base-stock production policies. Our first policy uses only local inventory information at each stage of the supply chain. Our second policy has similar structure to the policy proposed by Clark and Scarf [CS60], that is, each stage makes decisions based on the total

downstream inventory (echelon inventory). In both cases, we introduce constraints that ensure that stockout probabilities stay bounded below given desirable levels. Such *service-level* constraints provide a more natural representation of customer satisfaction and are closely watched by manufacturing managers. This is in contrast to most of the work in the literature that considers policies minimizing expected *linear* inventory and backorder costs. Our analysis is general enough to accommodate dependencies in demand and production processes. In practice, demand for various products might have strong correlations with a variety of phenomena such as: sales events, weather patterns, state of the economy, etc. Moreover, manufacturing facilities are *stochastic* and *failure-prone*, which creates dependencies in the production process. Under such assumptions, analyzing stockout probabilities exactly is intractable. We instead rely upon *large deviations* techniques that lead to asymptotically tight approximations. As a result, we are able to analytically obtain the appropriate base-stock levels for both policies we consider. Related techniques have been recently used by Bertsimas and Paschalidis [BP01] to devise production policies in a multiclass, single-stage setting. Approximation techniques of this type, but in the simpler case of renewal demand and production processes, have been introduced by Glasserman [Gla96, Gla97].

On the technical front, we have been able to use the full power of *large deviations* techniques to accommodate temporal dependencies in the demand and production processes. Our results are “network” large deviations results (tandem queues in particular). Such results have only been obtained in limited network cases, as in Bertsimas, Paschalidis, and Tsitsiklis [BPT98b] which we use in the case of local inventory information. Our echelon inventory results take into account the strong coupling between different stages of the supply chain and, to the best of our knowledge, are the first of such “network” results to do so in the presence of stochastic and autocorrelated demand and production processes. It should be noted that stochastic production processes suffice to complicate the picture (see [BPT98b] on how the character of the departure process from a G/D/1 is significantly altered by introducing stochasticity in the service process). Our echelon inventory main result has an interesting interpretation: it identifies a bottleneck stage whose production capacity is

“responsible” for stockouts at stage 1. But this bottleneck stage *is not* necessarily the one with the smallest mean capacity; it is determined by more detailed *distributional* information on all stochastic processes involved. In this sense, such distributional information is critical in making inventory control decisions.

Our numerical results demonstrate that the large deviations asymptotics are *accurate* in a wide range of desired stockout probabilities, including relatively large ones. Key to this are some heuristics we propose to compute a prefactor in front of the large deviations exponential.

1.3 Organization of the Dissertation

The remainder of the dissertation is organized as follows:

In Section 1.4, we provide some background knowledge on dynamic programming and large deviations analysis. We will use these mathematical tools in the subsequent chapters.

In Chapter 2, we focus on pricing for a single resource and formulate the problems of maximizing revenue and social welfare as a dynamic programming problem. We derive properties of the optimal dynamic policy and provide an upper bound on the optimal performance. We develop an approximate dynamic programming approach for solving large problems efficiently. We also introduce the static pricing, a suboptimal policy.

In Chapter 3, we consider the pricing problem for fixed-routing multi-service communication networks and show that static pricing is asymptotically optimal in a regime of many small users. For both revenue and welfare maximization objectives we characterize the structure of the asymptotically optimal static prices. We employ a simulation-based approach to compute an effective policy away from the limiting regime. The approach can handle large realistic, instances of the problem. Illustrative numerical results are reported at the end of the chapter.

In Chapter 4 we extend the model we have considered in Chapter 3 to incorporate demand substitution effects. Our main results extend to this situation as well. Following

the development of Chapter 3, we develop an upper bound on the optimal performance, establish the asymptotic optimality of static pricing, and characterize the structure of the asymptotically optimal static policy. In the end of the chapter, we discuss the asymptotic optimality of static policy for welfare maximization problems with demand substitution effects.

We start the discussion on supply chain management with the single-stage inventory control problem in Chapter 5. We present our model and obtain the large deviations approximations on the stockout probability and the inventory cost, which are building blocks for analyzing multi-stage problems.

In Chapter 6, we consider multi-stage supply chains under a base-stock policy. Each stage has local inventory information only and we want to satisfy the service-level constraint on the finished goods of the supply chains, i.e., maintaining stockout probabilities at various stages below given thresholds. We propose a decomposition approach based on large deviations approximations and the results we obtained for single-stage system.

The policy obtained via the decomposition approach, although it maintains the service-level constraint at stage 1, might not necessarily be efficient in terms of expected inventory cost. Information of inventory availability in other stages might lead to lower such cost by giving the opportunity to trade-off inventory between different stages, i.e., lower the required safety stock in stages where inventory costs are high and compensate by increasing the safety stock in stages where costs are lower. In Chapter 7 we consider such a situation where each stage has knowledge of the total downstream inventory. We implement a so-called echelon base-stock policy. We analyze the supply chain under this policy using large deviations techniques and devise a production policy that minimizes expected inventory costs subject to given service-level constraints. We also discuss extensions to multiclass supply chains. Numerical results for both decomposition approach and the multi-echelon approach are presented at the end of this chapter.

Summary and directions for future research are in Chapter 8.

1.4 Preliminaries

In this section, we provide some background knowledge on the mathematical tools we will be using throughout the thesis, including dynamic programming ([Ber95]) and large deviations analysis ([DZ98, Buc90, SW95, DE97]). On dynamic programming, we discuss infinite horizon, average reward problems and the uniformization of continuous-time Markov chains. We also review some basic results on large deviations, which will also help in establishing some of our notation.

1.4.1 Dynamic Programming

Infinite Horizon Problems with Average Reward Per Stage

For infinite horizon dynamic programming problems, if the reward per stage is discounted or the system eventually enters a reward-free termination state, the optimal total expected reward will be finite. However in many situations, discounting is inappropriate and there is no natural reward-free termination state. In such situations it is often meaningful to optimize the average reward per stage.

On a notational remark, we will be denoting all vectors using boldface and assume that they are column vectors unless otherwise explicitly specified. For example, we will be writing $\mathbf{x} = (x_1, \dots, x_m)$ to identify the elements of a vector $\mathbf{x} \in \mathbb{R}^m$.

The average reward per stage starting from a state i under policy π in an infinite horizon problem is defined by

$$J_\pi(i) = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbf{E} \left[\sum_{k=0}^{N-1} g(\mathbf{x}_k, \mu_k(\mathbf{x}_k)) \mid \mathbf{x}_0 = i \right], \quad i = 1, \dots, n, \quad (1.1)$$

where \mathbf{x}_k and $\mu_k(\mathbf{x}_k)$ are the state and the decision of policy π at time k , and $g_k(\cdot)$ is the corresponding immediate reward at time k . We will discuss the maximization problem in what follows; the minimization problem can be handled essentially the same way. The results in this section can be proved under the following assumption.

Assumption A

One of the states, by convention state n , is such that for some integer $m > 0$, and for all initial states and all policies, n , is visited with positive probability at least once within the first m stages.

Assumption A can be shown to be equivalent to the assumption that the special state n is recurrent in the Markov chain corresponding to each stationary policy (see [Ber95]).

Proposition 1.4.1 *Under Assumption A, the following hold for the average reward per stage problem:*

1. *The optimal average reward J^* is the same for all initial states.*
2. *Bellman's equation takes the form*

$$J^* + h^*(i) = \max_{\mathbf{u} \in \mathcal{U}(i)} \left(g(i, \mathbf{u}) + \sum_{j=1}^n p_{ij}(\mathbf{u}) h^*(j) \right), \quad i = 1, \dots, n. \quad (1.2)$$

where J^* is the optimal average reward per stage, $\mathcal{U}(i)$ is the control space at state i , $h^*(i)$ is the relative or differential reward for state i , and for the special state n , $h^*(n) = 0$.

3. *If a scalar J and a vector $\mathbf{h} = (h(1), \dots, h(n))$ satisfy the Bellman equation in (1.2), then J is the average optimal reward per stage for each initial state i :*

$$J = \max_{\pi} J_{\pi}(i) = J^*(i), \quad i = 1, \dots, n.$$

In addition, out of all vectors \mathbf{h} satisfying this equation, there is a unique vector for which $h(n) = 0$. Furthermore, if $\mu^*(i)$ attains the maximum in Equation (1.2) for each i , the stationary policy μ^* is optimal, i.e., $J_{\mu^*}(i) = J$ for all i .

The computational methods for solving the average reward per stage problems include value iteration, policy iteration, linear programming, and suboptimal approaches such as adaptive aggregations. We provide an introduction to value iteration.

The most natural version of the value iteration method for the average reward problem is simply to select arbitrarily a terminal reward function, say J_0 , and to generate successively the corresponding optimal k -stage reward $J_k(i)$, $k = 1, 2, \dots$. This can be done by executing the DP algorithm starting with J_0 , i.e., by using the recursion

$$J_{k+1}(i) = \max_{\mathbf{u} \in \mathcal{U}(i)} \left(g(i, \mathbf{u}) + \sum_{j=1}^n p_{ij}(\mathbf{u}) J_k(j) \right), \quad i = 1, \dots, n.$$

The k -stage average reward $J_k(i)/k$ will converge to the optimal average reward per stage J^* as $k \rightarrow \infty$, that is,

$$\lim_{k \rightarrow \infty} \frac{J_k(i)}{k} = J^*.$$

The value iteration method above is simple and straightforward, but typically J_k diverges to ∞ or $-\infty$, which makes the calculation of $\lim_{k \rightarrow \infty} J_k(i)/k$ numerically cumbersome. Also this method will not provide us with a corresponding differential reward vector h^* . To bypass those difficulties, we can consider an alternative algorithm, known as *relative value iteration*:

$$h_{k+1}(i) = \max_{\mathbf{u} \in \mathcal{U}(i)} \left(g(i, \mathbf{u}) + \sum_{j=1}^n p_{ij}(\mathbf{u}) h_k(j) \right) - \max_{\mathbf{u} \in \mathcal{U}(s)} \left(g(s, \mathbf{u}) + \sum_{j=1}^n p_{sj}(\mathbf{u}) h_k(j) \right), \quad i = 1, \dots, n, \quad (1.3)$$

where s is some fixed state. Under Assumption A, it can be shown that the iterates $h_k(i)$ generated by (1.3) are bounded. If the iteration converges to some vector \mathbf{h} , then we have

$$J + h(i) = \max_{\mathbf{u} \in \mathcal{U}(i)} \left(g(i, \mathbf{u}) + \sum_{j=1}^n p_{ij}(\mathbf{u}) h(j) \right),$$

and

$$J = \max_{\mathbf{u} \in \mathcal{U}(s)} \left(g(s, \mathbf{u}) + \sum_{j=1}^n p_{sj}(\mathbf{u}) h(j) \right).$$

By Proposition 1.4.1(3), this implies that J is the optimal average reward per stage for all initial states, and \mathbf{h} is an associated relative reward vector.

Uniformization of a Continuous-time Markov Chain

For a continuous-time system with a finite number of states, we assume the time interval τ between the transition to state i and the transition to the next state is exponentially distributed with parameter $\nu_i(\mathbf{u})$, i.e.,

$$\mathbf{P}[\tau \leq t \mid i, \mathbf{u}] = 1 - e^{-\nu_i(\mathbf{u})t},$$

and τ is independent of earlier transition times, states, and controls. The transition probabilities are $p_{ij}(\mathbf{u})$, $i \neq j$. The parameter $\nu_i(\mathbf{u})$ is referred to as the *rate of transition* associated with state i and control \mathbf{u} , and is uniformly bounded in the sense that for some ν we have

$$\nu_i(\mathbf{u}) \leq \nu, \quad \text{for all } i, \mathbf{u} \in \mathcal{U}(i).$$

The average transition time associated with state i and control \mathbf{u} is $\mathbf{E}[\tau] = \frac{1}{\nu_i(\mathbf{u})}$. The state and control at any time t are denoted by $\mathbf{x}(t)$ and $\mathbf{u}(t)$, respectively, and stay constant between transitions. We use following notation:

t_k : The time of occurrence of the k th transition. By convention, we denote $t_0 = 0$.

$\tau_k = t_k - t_{k-1}$: The k th transition time interval.

$\mathbf{x}_k = \mathbf{x}(t_k)$: We have $\mathbf{x}(t) = \mathbf{x}_k$ for $t_k \leq t < t_{k+1}$.

$\mathbf{u}_k = \mathbf{u}(t_k)$: We have $\mathbf{u}(t) = \mathbf{u}_k$ for $t_k \leq t < t_{k+1}$.

We consider an average reward function of the form

$$J = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbf{E} \left[\int_0^T g(\mathbf{x}(t), \mathbf{u}(t)) dt \right] = \lim_{N \rightarrow \infty} \frac{1}{t_N} \mathbf{E} \left[\int_0^{t_N} g(\mathbf{x}(t), \mathbf{u}(t)) dt \right],$$

where g is a given reward function. Similar to discrete-time problems, a policy is a sequence $\pi = \{\mu_0, \mu_1, \dots\}$, where each μ_k is a function mapping states to controls with $\mu_i \in \mathcal{U}(i)$ for

all states i . Under π , the control applied in the interval $[t_k, t_{k+1})$ is $\mu_k(\mathbf{x}_k)$. The average reward for policy π is given by

$$J_\pi(\mathbf{x}_0) = \lim_{N \rightarrow \infty} \frac{1}{t_N} \mathbf{E} \left[\sum_{k=0}^{N-1} \int_{t_k}^{t_{k+1}} g(\mathbf{x}_k, \mu_k(\mathbf{x}_k)) dt \mid \mathbf{x}_0 = i \right], \quad i = 1, \dots, n. \quad (1.4)$$

In general, the transition rate $\nu_i(\mathbf{u})$ depends on the state and the control, but we can convert them to a uniform transition rate by allowing fictitious transitions from a state to itself. Let ν be a new uniform transition rate with $\nu_i(\mathbf{u}) \leq \nu$ for all i and \mathbf{u} , and define new transition probabilities by

$$\bar{p}_{ij}(\mathbf{u}) = \begin{cases} \frac{\nu_i(\mathbf{u})}{\nu} p_{ij}(\mathbf{u}), & \text{if } i \neq j, \\ 1 - \frac{\nu_i(\mathbf{u})}{\nu}, & \text{if } i = j. \end{cases}$$

We can convert a continuous-time Markov chain problem with transition rate $\nu_i(\mathbf{u})$, transition probabilities $p_{ij}(\mathbf{u})$, and the average reward in (1.4) into a discrete-time Markov chain problem with a uniform transition rate ν , transition probabilities $\bar{p}_{ij}(\mathbf{u})$, the reward per stage

$$\bar{g}(i, \mathbf{u}) = \frac{g(i, \mathbf{u})}{\nu},$$

and the average reward per stage as in (1.1).

1.4.2 Large Deviations

Consider a sequence of i.i.d. random variables X_i , $i \geq 1$, with mean $\mathbf{E}[X_1] = \bar{X}$. The strong law of large numbers asserts that $\frac{\sum_{i=1}^n X_i}{n}$ converges to \bar{X} , as $n \rightarrow \infty$, with probability 1 (w.p.1). Thus, for large n the event $\sum_{i=1}^n X_i > na$, where $a > \bar{X}$, (or $\sum_{i=1}^n X_i < na$, for $a < \bar{X}$) is a rare event. More specifically, its probability behaves as $e^{-nr(a)}$, as $n \rightarrow \infty$, where the function $\tau(\cdot)$ determines the rate at which the probability of this event is diminishing. Cramér's theorem [Cra38] determines $\tau(\cdot)$, and is considered the first Large Deviations

statement. In particular,

$$r(a) = \sup_{\theta} \left(\theta a - \log \mathbf{E}[e^{\theta X_1}] \right).$$

For random variables, Gärtner-Ellis Theorem (see [Buc90] and [DZ98]) establishes a *Large Deviations Principle (LDP)*, which is a generalization of Cramér's theorem. Consider a sequence $\{S_1, S_2, \dots\}$ of random variables, with values in \mathbb{R} and define

$$\Lambda_n(\theta) \triangleq \frac{1}{n} \log \mathbf{E}[e^{\theta S_n}]. \quad (1.5)$$

For the applications we have in mind, S_n is a partial sum process. Namely, $S_n = \sum_{i=1}^n X_i$, where X_i , $i \geq 1$, are identically distributed, possibly dependent, random variables. Let $\{S_n\}$ satisfy the following assumption.

Assumption B

1. *The limit*

$$\Lambda(\theta) \triangleq \lim_{n \rightarrow \infty} \Lambda_n(\theta) = \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{E}[e^{\theta S_n}] \quad (1.6)$$

exists for all θ , where $\pm\infty$ are allowed both as elements of the sequence $\Lambda_n(\theta)$ and as limit points.

2. *The origin is in the interior of domain $D_\Lambda \triangleq \{\theta \mid \Lambda(\theta) < \infty\}$ of $\Lambda(\theta)$.*

3. *$\Lambda(\theta)$ is differentiable in the interior of D_Λ and the derivative tends to infinity as θ approaches the boundary of D_Λ .*

4. *$\Lambda(\theta)$ is lower semicontinuous, i.e., $\liminf_{\theta_n \rightarrow \theta} \Lambda(\theta_n) \geq \Lambda(\theta)$, for all θ .*

We will refer to $\Lambda(\cdot)$ as the *limiting log-moment generating function*. Let us also define

$$\Lambda^*(a) \triangleq \sup_{\theta} (\theta a - \Lambda(\theta)), \quad (1.7)$$

which will be referred to as the *large deviation rate function*.

Theorem 1.4.2 (Gärtner-Ellis) *Under Assumption B, the following inequalities hold*

Upper Bound: For every closed set F

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{P} \left[\frac{S_n}{n} \in F \right] \leq - \inf_{a \in F} \Lambda^*(a).$$

Lower Bound: For every open set G

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{P} \left[\frac{S_n}{n} \in G \right] \geq - \inf_{a \in G} \Lambda^*(a).$$

We say that $\{S_n\}$ satisfies a LDP with *good rate function* $\Lambda^*(\cdot)$. The term “good” refers to the fact that the level sets $\{a \mid \Lambda^*(a) \leq k\}$ are compact for all $k \leq \infty$, which is a consequence of Assumption B (see [DZ98] for a proof). LDP is satisfied by many general processes, such as renewal processes, Markov-modulated processes, and stationary processes with mild mixing conditions (see [DZ95] and [Cha95]), that are widely used in modeling traffic in communications networks, manufacturing systems and other areas.

The Gärtner-Ellis Theorem (Theorem 1.4.2) intuitively asserts that for large enough n and for small $\epsilon > 0$

$$\mathbf{P}[S_n \in (na - n\epsilon, na + n\epsilon)] \sim e^{-n\Lambda^*(a)}. \quad (1.8)$$

This can be viewed as an extension of Cramér’s theorem to autocorrelated stochastic processes. The notation “ \sim ” should be interpreted as “asymptotically behaves”: more rigorously, the logarithm of the probability divided by n converges to $-\Lambda^*(a)$, as $n \rightarrow \infty$.

It is important to note that $\Lambda(\cdot)$ and $\Lambda^*(\cdot)$ are convex duals (Legendre transforms of each other)[DZ98]. Namely, along with (1.7), it also holds

$$\Lambda(\theta) = \sup_a (\theta a - \Lambda^*(a)). \quad (1.9)$$

In the sequel, we are also estimating the tail probabilities of the form $\mathbf{P}[S_n \leq na]$ or $\mathbf{P}[S_n \geq na]$. We therefore define large deviations rate functions associated with such tail probabilities. Consider the case where $S_n = \sum_{i=1}^n X_i$, the random variables X_i , $i \geq 1$, being identically distributed, and let $m = \mathbf{E}[X_1]$. It is easily shown (see Dembo and

Zeitouni [DZ98]) that $\Lambda^*(m) = 0$. Let us now define

$$\Lambda^{*+}(a) \triangleq \begin{cases} \Lambda^*(a) & \text{if } a > m \\ 0 & \text{if } a \leq m \end{cases} \quad (1.10)$$

and

$$\Lambda^{*-}(a) \triangleq \begin{cases} \Lambda^*(a) & \text{if } a < m \\ 0 & \text{if } a \geq m. \end{cases} \quad (1.11)$$

Notice that $\Lambda^{*+}(a)$ is non-decreasing and $\Lambda^{*-}(a)$ is non-increasing function of a , respectively. The convex duals of these functions are

$$\Lambda^+(\theta) \triangleq \begin{cases} \Lambda(\theta) & \text{if } \theta \geq 0 \\ +\infty & \text{if } \theta < 0 \end{cases} \quad (1.12)$$

and

$$\Lambda^-(\theta) \triangleq \begin{cases} \Lambda(\theta) & \text{if } \theta \leq 0 \\ +\infty & \text{if } \theta > 0 \end{cases} \quad (1.13)$$

respectively. In particular, $\Lambda^{*-}(a) = \sup_{\theta}(\theta a - \Lambda^-(\theta))$ and $\Lambda^{*+}(a) = \sup_{\theta}(\theta a - \Lambda^+(\theta))$.

Using the Gärtner-Ellis Theorem it can be shown that for all $\epsilon_1, \epsilon_2 > 0$ there exists n_0 such that for all $n \geq n_0$

$$e^{-n(\Lambda^{*-}(a)+\epsilon_2)} \leq \mathbf{P}[S_n \leq na] \leq e^{-n(\Lambda^{*-}(a)-\epsilon_1)}, \quad (1.14)$$

and

$$e^{-n(\Lambda^{*+}(a)+\epsilon_2)} \leq \mathbf{P}[S_n \geq na] \leq e^{-n(\Lambda^{*+}(a)-\epsilon_1)}. \quad (1.15)$$

On a notational remark, in the sequel we will be denoting by

$$S_{i,j}^X \triangleq \sum_{k=i}^j X_k, \quad i \leq j, \quad (1.16)$$

the partial sums of the random sequence $\{X_i; i \in \mathbb{Z}\}$. We will be also denoting by $\Lambda_X(\cdot)$ and $\Lambda_X^*(\cdot)$ the limiting log-moment generating function and the large deviations rate function, respectively, of the process X .

Chapter 2

Pricing in Communication Systems: Single-Link Case

In this chapter, we consider the pricing problem for multi-service communication systems. We focus on the single resource case and formulate the problem as a dynamic programming problem. We derive properties of the optimal dynamic policy and provide an upper bound on the optimal performance. We propose an approximate dynamic programming approach for solving large problems efficiently. We also introduce static pricing, a suboptimal policy. Our analysis can also be applied to welfare maximization problems.

2.1 Problem Formulation

In this section, we introduce a detailed model for the operation of the single resource. We adopt a model that has been introduced in [PT00]. We assume that the link provides M classes of service to its customers and has a total capacity C . Each call of class i pays a fee of u_i upon arrival, and the arrival processes of all classes are Poisson. We assume the arrival rate of class i calls is a known function of price u_i , which will be referred to as the *demand function* and denoted by $\lambda_i(u_i)$. We will write $\mathbf{u} = (u_1, \dots, u_M)$ and $\boldsymbol{\lambda} = (\lambda_1(u_1), \dots, \lambda_M(u_M))$. We will be making the following assumption for demand functions.

Assumption C

For every $i = 1, \dots, M$, $\lambda_i(u_i) \geq 0$, and there exists a price $u_{i,\max}$ beyond which $\lambda_i(u_i)$ becomes zero. Furthermore, the function $\lambda_i(u_i)$ is continuous and strictly decreasing in the

range $u_i \in [0, u_{i,\max}]$.

Due to this assumption, the demand is at its peak when prices are zero. We will use $\lambda_0 = (\lambda_{1,0}, \dots, \lambda_{M,0}) \triangleq \lambda(\mathbf{0})$ to denote the peak demand vector, where $\mathbf{0}$ is the vector of all zeroes.

A customer of class i requires r_i units of bandwidth and stays in the system for an exponentially distributed time interval with parameter μ_i . We will write $\mathbf{r} = (r_1, \dots, r_M)$ and $\boldsymbol{\mu} = (\mu_1, \dots, \mu_M)$.

Let $n_i(t)$ be the number of customers of class i being connected at time t . We will make the convention that $n_i(t)$ is a right-continuous function of time. The state of system at time t can be identified by the vector of $n_i(t)$, $i = 1, \dots, M$, which is denoted by $\mathbf{n}(t) = (n_1(t), \dots, n_M(t))$. A request of class i can be accepted only if there is sufficient bandwidth to accommodate it, that is,

$$\mathbf{r}'(\mathbf{n}(t) + \mathbf{e}_i) \leq C,$$

where prime denotes the transpose, and \mathbf{e}_i is the i th unit vector, namely, a vector with all its components zero except the i th component which is equal to one. Otherwise, the customer will be informed that the service is unavailable for now and the request will be lost. The state space for the problem is

$$\mathcal{S} \triangleq \{\mathbf{n} \mid \mathbf{r}'\mathbf{n} \leq C\}.$$

A *pricing policy* is a rule that determines the pricing vector $\mathbf{u}(t) = (u_1(t), \dots, u_M(t))$ at any time t as a function of the state $\mathbf{n}(t)$. Under the assumptions put in place, for any given pricing policy the system evolves as a continuous-time Markov chain with state $\mathbf{n}(t) \in \mathcal{S}$. We are interested in pricing policies for two distinct objectives: revenue maximization and social welfare maximization.

2.1.1 Revenue Maximization Problem

Let us fix a pricing policy $\mathbf{u}(t)$. Assuming that there is enough bandwidth to accept a class i call, the instantaneous expected revenue rate from those calls is $\lambda_i(u_i(t))u_i(t)$, since class i arrivals are Poisson with rate $\lambda_i(u_i(t))$. If there is not sufficient bandwidth to accept class i calls we can, without loss of generality, set $u_i(t) = u_{i,\max}$ and bring the instantaneous expected revenue rate to zero. Thus, the total expected long-term average revenue is given by

$$J = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{i=1}^M \mathbf{E} \left[\int_0^T \lambda_i(u_i(t))u_i(t)dt \right] = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbf{E} \left[\int_0^T \boldsymbol{\lambda}(\mathbf{u}(t))' \mathbf{u}(t)dt \right]. \quad (2.1)$$

The above limit is easily seen to exist for any pricing policy, because the state $\mathbf{n} = \mathbf{0}$, corresponding to an empty system, is recurrent.

2.1.2 Welfare Maximization Problem

To formulate the welfare maximization problem, we will interpret the demand model as follows. Potential calls of class i are generated according to a Poisson process with constant rate $\lambda_{i,0}$, which is the peak arrival rate of class i introduced earlier. A potential call of class i , if it goes through, results in a user utility of U_i , where U_i is a non-negative random variable taking values in the range $[0, u_{i,\max}]$. Let $f_i(u_i)$ be the continuous probability density function of U_i . We assume that a potential class i call decides to join the system if and only if the utility it will extract, U_i , exceeds the prevailing price u_i . This implies that class i calls are realized according to a randomly modulated Poisson process with rate $\lambda_i(u_i(t)) = \lambda_{i,0} \mathbf{P}[U_i \geq u_i(t)]$. Furthermore, the expected utility conditioned on the fact that a call has been established, under a current price of u_i is equal to $\mathbf{E}[U_i | U_i \geq u_i]$. Hence, the expected long-term average rate at which utility is generated is given by

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{i=1}^M \mathbf{E} \left[\int_0^T \lambda_i(u_i(t)) \mathbf{E}[U_i | U_i \geq u_i(t)] dt \right]. \quad (2.2)$$

This is an objective of exactly the same form as in the case of revenue maximization, except that the instantaneous revenue rate $\lambda_i(u_i)u_i$ of class i is replaced by $\lambda_i(u_i)\mathbf{E}[U_i | U_i \geq u_i(t)]$. Thus, the two problems can be treated using the same set of tools. According to the utility assumptions put in place we have:

$$\lambda_i(u_i) = \lambda_{i,0} \int_{u_i}^{u_{i,\max}} f_i(v) dv,$$

and

$$\lambda_i(u_i)\mathbf{E}[U_i | U_i \geq u_i] = \lambda_{i,0} \int_{u_i}^{u_{i,\max}} v f_i(v) dv.$$

Example: Suppose that the utility derived from a call of class i is uniformly distributed in the range $[0, u_{i,\max}]$. Then the demand function of class i is

$$\lambda_i(u_i) = \lambda_{i,0} \left(1 - \frac{u_i}{u_{i,\max}}\right),$$

and

$$\mathbf{E}[U_i | U_i \geq u_i] = \frac{u_i + u_{i,\max}}{2}.$$

2.2 Optimal Dynamic Policy

In this section, we show how to obtain an optimal pricing policy using dynamic programming, under both objectives of revenue and welfare maximization. We derive some properties of the optimal policies. We first consider the revenue maximization problem, then comment on the corresponding results for the welfare maximization.

2.2.1 Dynamic Programming Formulation

Let $\mathcal{C}(\mathbf{n})$ be the set of classes whose calls are lost at state \mathbf{n} , thus

$$\mathcal{C}(\mathbf{n}) = \{i \mid \mathbf{r}'(\mathbf{n} + \mathbf{e}_i) > C, \quad i = 1, \dots, M\}.$$

Our objective is to find an optimal pricing strategy that maximizes the expected revenue of the firm over an infinite horizon, that is maximize the expression in Equation (2.1). From Assumption C, the control space for the problem is

$$\mathcal{U} \triangleq \{\mathbf{u} \mid u_i \in [0, u_{i,\max}], i = 1, \dots, M\}.$$

Using the uniformization technique for continuous-time Markov chain, we obtain the Bellman equation for the revenue maximization problem

$$J^* + h(\mathbf{n}) = \max_{\mathbf{u} \in \mathcal{U}} \left[\sum_{\substack{i=1 \\ i \notin \mathcal{C}(\mathbf{n})}}^M \lambda_i u_i + \sum_{\substack{i=1 \\ i \in \mathcal{C}(\mathbf{n})}}^M \frac{\lambda_i}{\nu} h(\mathbf{n} + \mathbf{e}_i) + \sum_{i=1}^M \frac{n_i \mu_i}{\nu} h(\mathbf{n} - \mathbf{e}_i) \right. \\ \left. + \left(1 - \sum_{\substack{i=1 \\ i \in \mathcal{C}(\mathbf{n})}}^M \frac{\lambda_i}{\nu} - \sum_{i=1}^M \frac{n_i \mu_i}{\nu} \right) h(\mathbf{n}) \right], \quad \mathbf{n} \in \mathcal{S}, \quad (2.3)$$

where $\nu = \sum_{i=1}^M \left(\lambda_{i,0} + \mu_i \left\lceil \frac{C}{r_i} \right\rceil \right)$.

For simplicity, we use $H(\mathbf{n}, \mathbf{u})$ to denote the expected revenue rate at state \mathbf{n} under price \mathbf{u} . And we define the dynamic programming operator T as follows: for any function f , $(Tf)(\mathbf{n})$ is defined to be equal to the right-hand side of Equation (2.3), with h replaced by f . In particular, Equation (2.3) can be written as

$$J^* + h(\mathbf{n}) = \max_{\mathbf{u} \in \mathcal{U}} H(\mathbf{n}, \mathbf{u}) = (Th)(\mathbf{n}), \quad \mathbf{n} \in \mathcal{S}. \quad (2.4)$$

The optimal solution to the problem is the solution of the system of nonlinear equations in (2.4). However, the computational complexity increases with the size of the state space, which is exponential in the number of classes M . For multi-class problem with large capacity C , it is impractical to find the optimal price of each class by either solving the system of nonlinear equations directly or using standard techniques such as value iteration and policy iteration (see [Ber95]). To overcome the difficulty, we will examine alternative approaches in the following sections.

2.2.2 Some General Properties

We present some properties of the optimal dynamic policy. The first result establishes the monotonicity of the relative rewards. It corresponds to the intuitive fact that it is always more desirable to have more free resources, as they lead to additional revenue generating opportunities in the future (see [PT00] for the proof).

Theorem 2.2.1 (Monotonicity of $h(\mathbf{n})$, [PT00]) *For all i and all \mathbf{n} such that $\mathbf{r}'(\mathbf{n} + \mathbf{e}_i) \leq C$. we have $h(\mathbf{n}) \geq h(\mathbf{n} + \mathbf{e}_i)$, where \mathbf{e}_i denotes the i th unit vector.*

Theorem 2.2.2 (The infinite bandwidth case, [PT00]) *If there are no capacity constraints ($C = \infty$), the optimal revenue is given by*

$$J_\infty = \max_{\mathbf{u} \in \mathcal{U}} \sum_{i=1}^M \lambda_i(u_i) u_i.$$

and the optimal price vector is some constant \mathbf{u}_∞ that does not depend on the state \mathbf{n} . Furthermore, we have $J^ \leq J_\infty$.*

We now show that resource limitations always result in higher prices in comparison to the unconstrained case.

Theorem 2.2.3 ([PT00]) *There exists an optimal policy \mathbf{u}^* such that for every state \mathbf{n} , we have $\mathbf{u}^*(\mathbf{n}) \geq \mathbf{u}_\infty$.*

Proof: Fix some state \mathbf{n} . From the Bellman equation, we see that for all $i \notin \mathcal{C}(\mathbf{n})$, an optimal price $u_i^*(\mathbf{n})$ can be chosen by maximizing

$$\lambda_i(u_i) u_i + \frac{\lambda_i(u_i)}{\nu} (h(\mathbf{n} + \mathbf{e}_i) - h(\mathbf{n})).$$

Consider a value of u_i which is less than the i th component $u_{i,\infty}$ of \mathbf{u}_∞ . Then, $\lambda_i(u_i) u_i \leq \lambda_i(u_{i,\infty}) u_{i,\infty}$, by the definition of $u_{i,\infty}$. By Theorem 2.2.1, we have $h(\mathbf{n} + \mathbf{e}_i) - h(\mathbf{n}) \leq 0$. Also, by monotonicity of the demand function, we have $\lambda_i(u_i) \geq \lambda_i(u_{i,\infty})$. Using all of the

above inequalities, we obtain

$$\lambda_i(u_i)u_i + \frac{\lambda_i(u_i)}{\nu}(h(\mathbf{n} + \mathbf{e}_i) - h(\mathbf{n})) \leq \lambda_i(u_{i,\infty})u_{i,\infty} + \frac{\lambda_i(u_{i,\infty})}{\nu}(h(\mathbf{n} + \mathbf{e}_i) - h(\mathbf{n})).$$

This implies that u_i cannot be strictly better than $u_{i,\infty}$, and proves the result. ■

2.2.3 Welfare Maximization Case

The case of welfare maximization, can be treated similarly. Bellman equation remains the same, except that the reward rate $\lambda(\mathbf{u})'\mathbf{u}$ is replaced by $\sum_i \lambda_i(u_i)\mathbf{E}[U_i | U_i \geq u_i]$. As in Theorem 2.2.1, the relative rewards $h(\mathbf{n})$ are again monotonically non-increasing in \mathbf{n} . If the bandwidth is infinite, welfare is maximized by admitting every user, and the optimal price \mathbf{u}_∞ is equal to zero. For a finite capacity network, the optimal prices are non-negative, which provides a trivial extension of Theorem 2.2.3.

2.3 Static Pricing Policy

The optimal dynamic policy has a number of drawbacks: it (i) is computationally hard to obtain, (ii) leads to prices that fluctuate in a very short time scale, which is inconvenient for customers since they prefer to predict their cost in advance, and (iii) requires implementing an elaborate feedback mechanism to communicate prices to users. In this section, we introduce a suboptimal policy, to be called *static pricing* strategy and develop methods to obtain an optimal static price.

We define *static pricing* to be the policy that keeps the price constant and independent of the level of congestion, i.e.,

$$\mathbf{u}(t) = (u_1, \dots, u_M) = \mathbf{u}, \quad \forall \mathbf{n} \in \mathcal{S}.$$

Therefore the demand function is also constant,

$$\lambda(\mathbf{u}(t)) = (\lambda_1(u_1), \dots, \lambda_M(u_M)).$$

Under the constant demand, for each class of service, there exist some states where the system does not have enough resources to accept a new request of this class; then the request is blocked. Let $\pi(\mathbf{n})$ be the steady-state probability of state \mathbf{n} , $\sum_{\mathbf{n} \in \mathcal{S}} \pi(\mathbf{n}) = 1$. The blocking probability of service i is given by

$$\mathbf{P}_{\text{loss}}^i(\mathbf{u}) = \sum_{\substack{\mathbf{n} \in \mathcal{S} \\ i \in \mathcal{C}(\mathbf{n})}} \pi(\mathbf{n}).$$

We are interested in obtaining the best static pricing policy. The average revenue under price \mathbf{u} is

$$J(\mathbf{u}) = \sum_{i=1}^M \lambda(u_i) u_i (1 - \mathbf{P}_{\text{loss}}^i(\mathbf{u})).$$

By solving the following nonlinear programming problem, we obtain an optimal static price and corresponding revenue rate,

$$J_s = \max_{\mathbf{u} \in \mathcal{U}} \sum_{i=1}^M \lambda(u_i) u_i (1 - \mathbf{P}_{\text{loss}}^i(\mathbf{u})).$$

We need to calculate the blocking probabilities in each iteration of the nonlinear optimization algorithm, which is computationally expensive. The following algorithm from [Ros95] determines the blocking probability for each class of service.

Algorithm 2.3.1 *Recursive algorithm to calculate the blocking probabilities $\mathbf{P}_{\text{loss}}^i(\mathbf{u})$:*

1. Set $g(0) = 1$ and $g(c) = 0$, for $c < 0$;
2. For $c = 1, \dots, C$, set

$$g(c) = \frac{1}{c} \sum_{i=1}^M r_i \frac{\lambda(u_i)}{\mu_i} g(c - r_i);$$

3. Set

$$G = \sum_{c=0}^C g(c);$$

4. For $c = 0, \dots, C$, set

$$q(c) = \frac{g(c)}{G};$$

5. For $i = 1, \dots, M$, set

$$\mathbf{P}_{\text{loss}}^i = \sum_{c=C-r_i+1}^C q(c).$$

The computational complexity of Algorithm 2.3.1 is $O(CM)$, that is, linear in the size parameters. But $g(c)$ increases geometrically with respect to λ_i 's, so for a large system where the arrival rates are large, the algorithm would encounter numerical problems such as overflow.

Table 2.1 lists the numerical results for a two-class system with linear demand function: $\lambda_i(u_i) = \lambda_{i,0} - \lambda_{i,1}u_i$, $i = 1, 2$. J^* is the optimal revenue rate, J_s is the revenue rate of optimal static policy. Note that ρ_1 and ρ_2 in Table 2.1 can be interpreted as the potential load of the two classes on the link.

$\lambda_{1,0}$	$\lambda_{1,1}$	$\rho_1 \triangleq \frac{\lambda_{1,0}}{\mu_1 C / r_1}$	$\lambda_{2,0}$	$\lambda_{2,1}$	$\rho_2 \triangleq \frac{\lambda_{2,0}}{\mu_2 C / r_2}$	J^*	J_s
40.0	4.0	1.032	350.0	35.0	1.129	952.63	945.79
40.0	4.0	1.032	500.0	50.0	1.613	1281.65	1270.4
80.0	8.0	2.064	350.0	35.0	1.129	977.28	965.33
80.0	8.0	2.064	500.0	50.0	1.613	1288.97	1273.9
160.0	16.0	4.128	1280.0	128.0	4.129	2235.13	2206.1
320.0	32.0	8.256	2560.0	256.0	8.258	2613.36	2588.9
640.0	64.0	16.512	5120.0	512.0	16.516	2820.47	2804.1

Table 2.1: Numerical results of the optimal static policy for the two-class problem ($C = 155$, $r_1 = 4$, $r_2 = 1$, $\mu_1 = 1$, $\mu_2 = 2$).

2.4 Price Aggregation Approach for Single-class Case

An alternative approach is solving the DP by linear programming (see [Ber95]). The optimal solution, J^* , of the DP solves the following linear programming (LP) problem,

$$\begin{aligned} \min_{J, h(\mathbf{n})} \quad & J \\ \text{s.t.} \quad & J + h(\mathbf{n}) \geq H(\mathbf{n}, \mathbf{u}) \quad \text{for all } \mathbf{n} \in \mathcal{S} \text{ and } \mathbf{u} \in \mathcal{U}, \end{aligned} \tag{2.5}$$

where H , defined in (2.4), denotes the expected revenue rate at a given state \mathbf{n} and price \mathbf{u} . Unfortunately, for a system with large capacity, the dimension and the number of constraints of this program can be very large and its solution can be computationally impractical. We propose a price aggregation approach, which is suboptimal, for the single-class case.

Consider a single-class problem where the state space is $\mathcal{S} = \{0, \dots, C\}$ and the control space is $\mathcal{U} = [0, u_{\max}]$. We can formulate the following LP problem,

$$\begin{aligned} \min_{J, h(n)} \quad & J \\ \text{s.t.} \quad & J \geq \lambda(u)u + \lambda(u)h(1) - \lambda(u)h(0), \quad \forall u \in \mathcal{U} \\ & J \geq \lambda(u)u + \lambda(u)h(n+1) + n\mu h(n-1) - (\lambda(u) + n\mu)h(n), \quad 0 < n < C, \forall u \in \mathcal{U} \\ & J \geq C\mu h(C-1) - C\mu h(C), \quad \forall u \in \mathcal{U}. \end{aligned} \tag{2.6}$$

The dual problem is

$$\begin{aligned} \max_{q(n, u)} \quad & \sum_{n=0}^{C-1} \sum_{u \in \mathcal{U}} q(n, u) \lambda(u) u \\ \text{s.t.} \quad & \sum_{u \in \mathcal{U}} (q(0, u) \lambda(u) - q(1, u) \mu) = 0 \\ & \sum_{u \in \mathcal{U}} (q(n-1, u) \lambda(u) - q(n, u) (\lambda(u) + n\mu) + q(n+1, u) (n+1)\mu) = 0, \quad 0 < n < C \\ & \sum_{u \in \mathcal{U}} (q(C-1, u) \lambda(u) - q(C, u) C\mu) = 0 \end{aligned} \tag{2.7}$$

$$\sum_{n=0}^C \sum_{u \in \mathcal{U}} q(n, u) = 1$$

$$q(n, u) \geq 0, \quad n = 0, \dots, C, \quad u \in \mathcal{U},$$

where $q(n, u)$ can be interpreted the steady-state probability of state n under price u .

Theorem 2.4.1 *The optimal solution to the dual problem in (2.7) is*

$$q(n, u) = \begin{cases} \pi_n & \text{if } u = u^*(n) \\ 0 & \text{otherwise,} \end{cases} \quad (2.8)$$

where $u^*(\cdot)$ is the optimal dynamic policy and π_n is the steady-state probability of state n under this policy.

Proof: It can be seen that $u^*(C) = u_{\max}$, so that $\lambda(u^*(C)) = 0$. Let π_n be the steady-state probability of state n under the optimal policy, where $\pi_n \geq 0$ and

$$\sum_{n=0}^C \pi_n = 1.$$

The optimal revenue rate is

$$J^* = \sum_{n=0}^C \pi_n \lambda(u^*(n)) u^*(n) = \sum_{n=0}^{C-1} \pi_n \lambda(u^*(n)) u^*(n),$$

which is also the optimal value of the LP problem in (2.6).

Note that $q(n, u)$ given by Equation (2.8) is a feasible solution to the dual problem (2.7), since the constraints become the detailed balance equations of the Markov chain that governs the evolution of the system. The corresponding cost is given by

$$\sum_{n=0}^{C-1} \sum_{u \in \mathcal{U}} q(n, u) \lambda(u) u = \sum_{n=0}^{C-1} \pi_n \lambda(u^*(n)) u^*(n) = J^*.$$

From the strong duality theorem, we conclude that Equation (2.8) is the optimal solution

to the dual problem (2.7). ■

To address the computational problems associated with large scale problems, we discretize u into L intervals $(\alpha_0, \alpha_1), (\alpha_1, \alpha_2), \dots, (\alpha_{L-1}, \alpha_L)$, where $\alpha_0 = 0, \alpha_0 < \dots < \alpha_L$, and $\alpha_L = u_{\max}$. Note that for any state n and price interval $U_j = (\alpha_{j-1}, \alpha_j), j = 1, \dots, L$, we define $\hat{q}(n, j) \triangleq \sum_{u \in U_j} q(n, u)$ as the total steady-state probability of state n under price in the interval U_j . By the monotonicity of $\lambda(\cdot)$, we obtain

$$\hat{q}(n, j)\lambda(\alpha_j) \leq \sum_{u \in U_j} q(n, u)\lambda(u) \leq \hat{q}(n, j)\lambda(\alpha_{j-1}), \quad (2.9)$$

Using the above inequality, we can aggregate the constraints of problem (2.7) over u , and obtain the following LP problem,

$$\begin{aligned} \max_{\hat{q}(n, j)} \quad & \sum_{n=0}^{C-1} \sum_{j=1}^L \hat{q}(n, j)\lambda(\alpha_{j-1})\alpha_j & (2.10) \\ \text{s.t.} \quad & \sum_{j=1}^L (\hat{q}(0, j)\lambda(\alpha_j) - \hat{q}(1, j)\mu) \leq 0 \\ & \sum_{j=1}^L (\hat{q}(0, j)\lambda(\alpha_{j-1}) - \hat{q}(1, j)\mu) \geq 0 \\ & \sum_{j=1}^L (\hat{q}(n-1, j)\lambda(\alpha_j) + \hat{q}(n+1, j)(n+1)\mu - \hat{q}(n, j)(\lambda(\alpha_{j-1}) + n\mu)) \leq 0, \quad 0 < n < C \\ & \sum_{j=1}^L (\hat{q}(n-1, j)\lambda(\alpha_{j-1}) + \hat{q}(n+1, j)(n+1)\mu - \hat{q}(n, j)(\lambda(\alpha_j) + n\mu)) \geq 0, \quad 0 < n < C \\ & \sum_{j=1}^L (\hat{q}(C-1, j)\lambda(\alpha_j) - \hat{q}(C, j)C\mu) \leq 0 \\ & \sum_{j=1}^L (\hat{q}(C-1, j)\lambda(\alpha_{j-1}) - \hat{q}(C, j)C\mu) \geq 0 \\ & \sum_{n=0}^C \sum_{j=1}^L \hat{q}(n, j) = 1 \\ & \hat{q}(n, j) \geq 0, \quad \forall n, j. \end{aligned}$$

The aggregated LP (2.10) does not have a solution of the form (2.8). But motivated by (2.8), for a state n , we can construct a policy $\hat{u}(n)$ as follows:

$$\hat{u}(n) = \frac{\alpha_{j^*-1} + \alpha_{j^*}}{2},$$

where

$$j^* = \arg \max_j \hat{q}(n, j),$$

i.e., $\hat{u}(n)$ is the middle of most likely price interval.

For the single-class system with a linear demand function $\lambda(u) = \lambda_0 - \lambda_1 u$, Table 2.2 reports numerical results of the system with different parameters (λ_0 , λ_1 and μ). The total bandwidth is $C = 30$ and the bandwidth required for one connection is $r = 1$. The optimal revenue (dynamic policy) is J^* and the revenue of the optimal static policy is J_s . We partition the control space (feasible price) into L uniform intervals, \bar{J}_L is the optimal value of the aggregated LP (2.10), J_L^* is revenue rate obtained by the dynamic policy $\hat{u}(n)$. The policy we get for the first row of Table 2.2 is shown in Figure 2.1(a). The optimal dynamic policy is depicted in Figure 2.1(c).

λ_0	λ_1	μ	$\rho \triangleq \frac{\lambda_0}{\mu C / r}$	J^*	J_s	L	\bar{J}_L	J_L^*
60.0	5.0	1.0	2.0	167.68	165.92	80	175.83	167.62
50.0	5.0	1.0	1.67	121.48	120.70	60	127.39	121.43
50.0	5.0	0.5	3.33	92.49	91.12	140	96.98	92.40
45.0	5.0	1.0	1.5	99.82	99.43	50	104.80	99.79
45.0	5.0	0.5	3.0	79.47	78.30	130	83.25	79.38
45.0	5.0	0.2	7.5	42.03	41.57	260	44.05	41.81
45.0	0.5	0.1	15	232.18	230.60	410	243.47	229.55
45.0	0.5	0.05	30	12.28	12.23	790	12.87	12.20

Table 2.2: Numerical results of price aggregation for single-class problem ($C = 30$, $r = 1$), with uniform price interval.

It has been proved in [PT00] that for the finite capacity single-class system, the optimal price is a nondecreasing function of the state (resources being occupied), and from Theorem 2.2.3, the optimal prices are always larger than the optimal price (constant) for the infinite

capacity system. We observed that when the system is nearly empty, the optimal price varies slowly with the state \mathbf{n} . Based on these observations, we can form price intervals with varied sizes, more specifically, we use large intervals for low prices and small intervals for high prices, to decrease the total number of intervals and guarantee sufficient accuracy at the same time. Table 2.3 shows the results of this approach. Comparing with the approach of uniform intervals, we can see that we can solve a smaller LP problem to obtain an equally good result. Figure 2.1(b) depicts the policy obtained from this method.

λ_0	λ_1	μ	$\rho \triangleq \frac{\lambda_0}{\mu C/r}$	J^*	J_s	L	\bar{J}_L	J_L^*
60.0	5.0	1.0	2.0	167.68	165.92	57	175.83	167.62
50.0	5.0	1.0	1.67	121.48	120.70	48	127.39	121.43
50.0	5.0	0.5	3.33	92.49	91.12	78	96.97	92.31
45.0	5.0	1.0	1.5	99.82	99.43	40	104.80	99.79
45.0	5.0	0.5	3.0	79.47	78.30	74	83.375	79.31
45.0	5.0	0.2	7.5	42.03	41.57	101	44.05	41.91
45.0	0.5	0.1	15	232.18	230.60	157	243.68	231.38
45.0	0.5	0.05	30	12.28	12.23	289	12.89	10.74

Table 2.3: Numerical results of price aggregation for single-class problem ($C = 30$, $r = 1$), with varied price interval.

From these examples, we can see that using the price aggregation approach, we can get a satisfactory dynamic policy and with this policy prices would change in a longer time scale than the state \mathbf{n} . Such a policy is more attractive to implement and more acceptable by users. Moreover, this approach provides an upper bound on the optimal revenue which is useful in assessing the degree of suboptimality of suboptimal policies. The approach can be extended to multi-class case, but the notation will be more cumbersome.

2.5 An Upper Bound on Optimal Revenue

In most cases, it would be hard to find the optimal policy and/or optimal revenue rate, we have to therefore resort suboptimal policies. The critical question is how do we determine that a suboptimal policy is “good enough”. To this end, in this section we will obtain an upper bound on the optimal revenue rate.




Figure 2.1: The dynamic policy for a system with total bandwidth $C = 30$, which provides a single-class service ($r = 1$), and $\lambda_0 = 60$, $\lambda_1 = 5$, $\mu = 1$, (a) by price aggregation with uniform interval; (b) by price aggregation with varied interval; (c) by value iteration for DP (optimal policy).

The idea is that while the system stays at the steady-state, the mean arrival and departure rate of the customers are equal. As in Section 2.4, we divide the price into L intervals, (α_0, α_1) , (α_1, α_2) , \dots , (α_{L-1}, α_L) , where $\alpha_0 = 0$, $\alpha_0 < \dots < \alpha_L$, and $\alpha_L = u_{\max}$. Let $\pi(k)$ be the steady-state probability of being in price interval k and n the average number of the customers at steady-state. The problem is formulated using the flow balance equations. Consider a system with M classes of service, and let $\pi(i, k)$ be the steady-state probability of the price of the i th type of service being in price interval k . The following LP problem provides an upper bound on the optimal revenue rate.

$$\begin{aligned}
& \max_{\pi(i,k)} && \sum_i \sum_k \pi(i, k) \lambda_i (\alpha_{k-1}) \alpha_k && (2.11) \\
& \text{s.t.} && \sum_k \pi(i, k) = 1, \quad \forall i \\
& && n_i \mu_i \geq \sum_k \lambda_i (\alpha_k) \pi(i, k), \quad \forall i \\
& && n_i \mu_i \leq \sum_k \lambda_i (\alpha_{k-1}) \pi(i, k), \quad \forall i \\
& && \sum_i n_i r_i \leq C \\
& && \pi(i, k) \geq 0, \quad \forall i, k \\
& && n_i \geq 0, \quad \forall i.
\end{aligned}$$

The computational examples are given in Table 2.4 and Table 2.5, where L is the number of price intervals and \bar{J}_L denotes the optimal values of the LP in (2.11). By solving a nonlinear programming (NLP) problem, [PT00] provides another tighter upper bound.

λ_0	λ_1	μ	$\rho \triangleq \frac{\lambda_0}{\mu C/r}$	J^*	L	\bar{J}_L
60.0	5.0	1.0	2.0	167.68	10000	180.04
50.0	5.0	1.0	1.67	121.48	10000	125.03
50.0	5.0	0.5	3.33	92.49	10000	105.04
45.0	5.0	1.0	1.5	99.82	10000	101.27
45.0	5.0	0.5	3.0	79.47	10000	90.03
45.0	5.0	0.2	7.5	42.03	10000	46.84
45.0	0.5	0.1	15	232.18	10000	252.38
45.0	0.5	0.05	30	12.28	10000	13.09

Table 2.4: Numerical results of the upper bound for single-class problems ($C = 30$, $r = 1$).

$\lambda_{1,0}$	$\lambda_{1,1}$	$\rho_1 \triangleq \frac{\lambda_{1,0}}{\mu_1 C/r_1}$	$\lambda_{2,0}$	$\lambda_{2,1}$	$\rho_2 \triangleq \frac{\lambda_{2,0}}{\mu_2 C/r_2}$	J^*	L	\bar{J}_L
40.0	4.0	1.032	350.0	35.0	1.129	952.63	20000	972.95
40.0	4.0	1.032	500.0	50.0	1.613	1281.65	20000	1317.47
80.0	8.0	2.064	350.0	35.0	1.129	977.28	20000	1012.56
80.0	8.0	2.064	500.0	50.0	1.613	1288.97	20000	1329.89
160.0	16.0	4.128	1280.0	128.0	4.129	2235.13	20000	2349.78
320.0	32.0	8.256	2560.0	256.0	8.258	2613.36	20000	2725.89
640.0	64.0	16.512	5120.0	512.0	16.516	2820.47	20000	2915.03

Table 2.5: Numerical results of the upper bound for two-class problems ($C = 155$, $r_1 = 4$, $r_2 = 1$, $\mu_1 = 1$, $\mu_2 = 2$).

2.6 Approximate DP Approach

For large, multi-class problems, it is nearly impossible to get the optimal dynamic price through value iteration, and even the price aggregation method in Section 2.4 becomes impractical. In this section, we explore an approximate dynamic programming approach through which we can obtain a suboptimal dynamic policy.

The idea is that we still use the LP formulation in Equation (2.5) to obtain the policy,

$$\begin{aligned} \min_{J, h(\mathbf{n})} \quad & J \\ \text{s.t.} \quad & J + h(\mathbf{n}) \geq H(\mathbf{n}, \mathbf{u}), \quad \forall \mathbf{u} \in \mathcal{U}, \mathbf{n} \in \mathcal{S}, \end{aligned} \tag{2.12}$$

where the control space \mathcal{U} is discretized. When the system is very large, we will have many feasible states \mathbf{n} and large number of (\mathbf{n}, \mathbf{u}) pairs. This means the number of variables ($h(\mathbf{n})$ and J) and constraints is very large, and the LP problem (2.12) is hard to solve. To address this problem, we approximate the reward function $h(\mathbf{n})$ with the *linear* form

$$h(\mathbf{n}) = \sum_{k=1}^K \psi_k w_k(\mathbf{n}), \tag{2.13}$$

where $\psi = (\psi_1, \dots, \psi_K)$ is a vector of parameters, and $w_k(\mathbf{n})$ are given functions of the state \mathbf{n} . This amounts to approximating the reward function $h(\mathbf{n})$ by a linear combination of K given function $w_k(\mathbf{n})$, $k = 1, \dots, K$. It is then possible to determine $\psi = (\psi_1, \dots, \psi_K)$ by plugging (2.13) in $H(\mathbf{n}, \mathbf{u})$ in the preceding linear programming problem. This new problem has a small number of variables but still many constraints. We will use a cutting-plane method guided by simulation. In particular, we initially solve an LP similar to (2.12). But its constraint is only a subset of the constraint in (2.12) corresponding to just a few states, and for any of those states, \mathbf{n} , we select the price that maximizes the instantaneous revenue rate H as its price,

$$\hat{u}(\mathbf{n}) = \max_{\mathbf{u} \in \mathcal{U}} H(\mathbf{n}, \mathbf{u}).$$

Then we use simulation to evaluate the performance of this policy. Starting from one of these states, as the simulation goes on, the system will get to a new state $\bar{\mathbf{n}}$ (not included in the original LP). We check the constraints $J + h(\bar{\mathbf{n}}) \geq H(\bar{\mathbf{n}}, \mathbf{u})$ for that new state, and if for all possible prices, the constraints are satisfied, we find the price that maximizes the instantaneous revenue rate and use that as our policy for this state and continue the simulation; if there are some constraints that are violated, we add the violated constraints


into the original LP problem and solve the extended problem, then update the policy and continue the simulation. Eventually, we will find a “good” dynamic policy for the problem.

We present numerical results for two examples to show the efficiency of the approximate DP approach. In those examples, we selected the following approximation for the reward function $h(\mathbf{n})$:

$$h(\mathbf{n}) = \mathbf{n}'\mathbf{Q}\mathbf{n} + \mathbf{c}'\mathbf{n},$$

where \mathbf{Q} is an $M \times M$ symmetric matrix and \mathbf{c} is an M -dimensional vector, M is the number of classes. Here, we set the relative reward of the empty state to be zero.

Example 1 is a single-class system with a total bandwidth of $C = 30$. The bandwidth required for each connection is $r = 1$. The demand function is linear, which has the form $\lambda(u) = \lambda_0 - \lambda_1 u$. The parameters are $\lambda_0 = 60$, $\lambda_1 = 5$. The departure rate is $\mu = 1$.



(a) The relative reward

(b) The dynamic policy

Figure 2.2: The numerical results for Example 1. The control space \mathcal{U} is discretized into $L = 50$ intervals. The revenue obtained from the approximate policy is 161.24, while the optimal revenue is 167.68.

Figure 2.2 depicts the relative reward and dynamic policy for the system in Example 1 obtained by solving the LP with the quadratic approximation of the relative reward function (compared with the optimal values). We can see that the quadratic approximation captures the basic feature of the relative reward function.

In Example 2, the system with total bandwidth $C = 155$ provides two classes of service ($M = 2$). The first class requires 4 units of bandwidth, the second class requires 1 unit of bandwidth ($r_1 = 4, r_2 = 1$). Both classes have linear demand functions, the parameters

are $\lambda_{1,0} = 40$, $\lambda_{1,1} = 4$, $\lambda_{2,0} = 350$, $\lambda_{2,1} = 35$. The service rate is $\mu_1 = 1$ and $\mu_2 = 2$, respectively. For this two-class problem, Figure 2.3 depicts the dynamic policy obtained by the approximate DP, and through the simulation, the revenue of this policy is 920.529. The optimal revenue is 952.63.

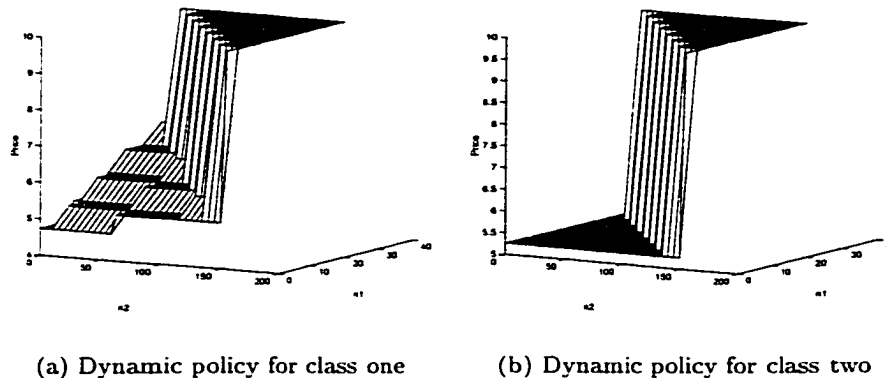


Figure 2.3: The numerical results of the approximate DP for Example 2. The control space \mathcal{U} is discretized into $L = 20$ intervals. The revenue obtained from this approximate policy is 920.529. The optimal revenue is 952.63.

The most important advantage of this approximate DP method is that it can solve very large problems. Table 2.6 lists the numerical results of this method for some large problems. For these problems, it is very hard to get the optimal revenue. So we cannot guarantee the result of approximated DP is close to the optimal value, but we believe that it is possible to get a good policy if we can find an accurate and simple approximation of the reward function.

C	$\lambda_{1,0}$	$\lambda_{1,1}$	τ_1	μ_1	$\lambda_{2,0}$	$\lambda_{2,1}$	τ_2	μ_2	J^*	J_{simu}	J
10	40	4	4	1	350	35	1	2	164.63	159.54	188.76
155	40	4	4	1	350	35	1	2	952.63	920.53	972.95
155	70	4	4	1	550	35	1	2	–	2074.44	2260.99
1550	400	40	4	1	3500	350	1	2	–	8956.29	9729.53
8500	400	40	4	1	35000	3500	1	2	–	85430.68	87781.79

Table 2.6: Numerical results of the approximate DP method for some two-class problems. J^* is the optimal revenue derived by value iteration, J is the upper bound of the problem obtained by the way in Section 2.5, J_{simu} is the simulation result with the policy obtained from approximate DP.

Chapter 3

Pricing in Communication Systems: The Network Case

In this chapter, we consider revenue and welfare maximization problems for fixed-routing multi-service communication networks and show that static pricing is asymptotically optimal in a regime of many small users. In particular, the performance of an optimal (dynamic) pricing strategy is closely matched by a suitably chosen static policy. For both revenue and welfare maximization objectives we characterize the structure of the asymptotically optimal static prices. We employ a simulation-based approach to compute an effective policy away from the limiting regime. Numerical examples are reported. The approach can handle large realistic instances of the problem.

3.1 The Model of Multi-Service Networks

In this section we will introduce the model of the multi-service network we wish to study. We consider a network with L links. The capacity of each link j is C_j of bandwidth for $j = 1, \dots, L$. We will write $\mathbf{C} = (C_1, \dots, C_L)$. The network provides M classes of service. Each service class is distinguished by its demand pattern, bandwidth requirement, call duration, and routing through the network. Classes have a fixed route through the network. In particular, class i requires r_{ji} units of bandwidth from link j , for $i = 1, \dots, M$ and $j = 1, \dots, L$. The routing matrix will be denoted by $\mathbf{R} = \{r_{ji}\}$, i.e., an $L \times M$ matrix with the (j, i) element being equal to r_{ji} . The route of class i is characterized by the

sequence of links $j_{i_1}, j_{i_2}, \dots, j_{i_l}$ it traverses; we will denote it by

$$\mathcal{R}_i = \{(j_{i_1}, j_{i_2}, \dots, j_{i_l}) \mid 1 \leq j_{i_1}, \dots, j_{i_l} \leq L, r_{j_{i_k}} > 0, k = 1, \dots, l\}, \quad i = 1, \dots, M.$$

We will write $j \in \mathcal{R}_i$ if link j is any link in the sequence $(j_{i_1}, j_{i_2}, \dots, j_{i_l})$. To exclude trivial cases, we will be assuming that $\mathcal{R}_i \neq \emptyset$ for each class i . For all other links j that are not in route \mathcal{R}_i it is understood that $r_{j_i} = 0$.

As the single-link case, we assume that calls of class $i = 1, \dots, M$ arrive according to a Poisson process and stay in the system for a time interval which is exponentially distributed with rate μ_i . Let $\boldsymbol{\mu} = (\mu_1, \dots, \mu_M)$. The network charges a fee u_i per call of class i , which can depend on the current congestion level and which affects user's demand for calls. We will assume that the demand functions are known and denoted by $\lambda_i(u_i)$ for class i , $i = 1, \dots, M$. We will write $\mathbf{u} = (u_1, \dots, u_M)$ and $\boldsymbol{\lambda}(\mathbf{u}) = (\lambda_1(u_1), \dots, \lambda_M(u_M))$. We will be making the same assumption as Assumption C in Section 2.1 about demand functions. We will use $\boldsymbol{\lambda}_0 = (\lambda_{1,0}, \dots, \lambda_{M,0}) \triangleq \boldsymbol{\lambda}(\mathbf{0})$ to denote the peak demand vector.

Let $n_i(t)$ be the number of class i calls that are in progress at time t . We will denote by $\mathbf{n}(t) = (n_1(t), \dots, n_M(t))$ the state of the system at time t . An incoming class i call is accepted if all the links along its route have enough available bandwidth, that is, if

$$\mathbf{R}(\mathbf{n}(t) + \mathbf{e}_i) \leq \mathbf{C},$$

where \mathbf{e}_i is the i th unit vector. If this latter condition is violated, an incoming call is rejected and lost for the system. Let $\mathcal{S} = \{\mathbf{n} \mid \mathbf{R}\mathbf{n} \leq \mathbf{C}\}$ denote the state space for the system, i.e., the set of states at which capacity constraints are satisfied.

A pricing policy is a rule that determines the pricing vector $\mathbf{u}(t) = (u_1(t), \dots, u_M(t))$ at any time t as a function of the state $\mathbf{n}(t)$. As in Chapter 2, we are interested in pricing policies for revenue maximization and social welfare maximization.

3.1.1 Revenue Maximization Problem

The formulation of revenue maximization problem is the same as in single-link case. For a pricing policy $\mathbf{u}(t)$, at time t , if there is enough bandwidth to accept a class i call, the instantaneous expected revenue rate from those calls is $\lambda_i(u_i(t))u_i(t)$, since class i arrivals are Poisson with rate $\lambda_i(u_i(t))$; if there is not sufficient bandwidth to accept class i calls we can, without loss of generality, set $u_i(t) = u_{i,\max}$ such that $\lambda_i(u_i(t)) = 0$. Thus, the total expected long-term average revenue is given by

$$J = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{i=1}^M \mathbf{E} \left[\int_0^T \lambda_i(u_i(t))u_i(t)dt \right] = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbf{E} \left[\int_0^T \boldsymbol{\lambda}(\mathbf{u}(t))' \mathbf{u}(t)dt \right]. \quad (3.1)$$

3.1.2 Welfare Maximization Problem

As the single-link case, we associate a utility U_i with a potential call of class i . U_i is a random variable in $[0, u_{i,\max}]$ with density function $f_i(u_i)$; potential call arrival of class i is a Poisson process with rate $\lambda_{i,0}$. Class i calls are realized according to a randomly modulated Poisson process with rate $\lambda_i(u_i(t)) = \lambda_{i,0} \mathbf{P}[U_i \geq u_i(t)]$. The expected utility conditioned on the fact that a call has been established, under current price of u_i is equal to $\mathbf{E}[U_i | U_i \geq u_i]$. The expected long-term average rate at which utility is generated is given by

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{i=1}^M \mathbf{E} \left[\int_0^T \lambda_i(u_i(t)) \mathbf{E}[U_i | U_i \geq u_i(t)] dt \right]. \quad (3.2)$$

According to the $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{i=1}^M \mathbf{E} \left[\int_0^T \lambda_i(u_i(t)) \mathbf{E}[U_i | U_i \geq u_i(t)] dt$

$$\lambda_i(u_i) = \lambda_{i,0} \int_{u_i}^{u_{i,\max}} f_i(v) dv, \quad (3.3)$$

and

$$\lambda_i(u_i) \mathbf{E}[U_i | U_i \geq u_i] = \lambda_{i,0} \int_{u_i}^{u_{i,\max}} v f_i(v) dv. \quad (3.4)$$

3.2 Optimal Dynamic Policy

We will start the analysis by considering optimal (dynamic) pricing policies. Under both objectives of revenue and welfare maximization the problem can be formulated using stochastic dynamic programming (DP). We first consider revenue maximization.

The state of the system $\mathbf{n}(t)$ evolves as a continuous-time Markov chain and its total transition rate out of any state is bounded by

$$\nu = \sum_{i=1}^M \left(\lambda_{0,i} + \mu_i \max_{j \in \mathcal{R}_i} \left[\frac{C_j}{r_{ji}} \right] \right).$$

The Markov chain can be uniformized, leading to a Bellman equation of the form

$$J^* + h(\mathbf{n}) = \max_{\mathbf{u} \in \mathcal{U}} \left[\sum_{i \notin \mathcal{C}(\mathbf{n})} \lambda_i(u_i) u_i + \sum_{i \in \mathcal{C}(\mathbf{n})} \frac{\lambda_i(u_i)}{\nu} h(\mathbf{n} + \mathbf{e}_i) + \sum_{i=1}^M \frac{n_i \mu_i}{\nu} h(\mathbf{n} - \mathbf{e}_i) + \left(1 - \sum_{i \in \mathcal{C}(\mathbf{n})} \frac{\lambda_i(u_i)}{\nu} - \sum_{i=1}^M \frac{n_i \mu_i}{\nu} \right) h(\mathbf{n}) \right]. \quad (3.5)$$

where $\mathcal{U} = \{\mathbf{u} \mid 0 \leq u_i \leq u_{i,\max}, \forall i\}$ is the set of possible price vector and $\mathcal{C}(\mathbf{n}) = \{i \mid \mathbf{R}(\mathbf{n} + \mathbf{e}_i) \not\subseteq \mathbf{C}\}$ is the set of classes whose calls cannot be admitted in state \mathbf{n} . Here J^* and $h(\mathbf{n})$ denote the optimal expected revenue rate and the relative reward in state \mathbf{n} . This DP formulation is in fact almost identical to the one in Section 2.2.1, the only difference being the definition of $\mathcal{C}(\mathbf{n})$ which has been extended to the network setting. It has been argued there that the standard infinite-horizon average-cost dynamic programming theory applies (see [Ber95]), thus, there exists a stationary policy which is optimal. We will use $\mathbf{u}^*(\mathbf{n})$ to denote an optimal policy to explicitly indicate its dependence on the state of the system. Such a policy can be found by solving Bellman's equation using standard DP algorithms. However, Bellman's "curse of dimensionality" prohibits us from solving realistic instances of the problem. Consequently, we are interested in exploring simpler, yet not too far from the optimal, alternatives. Before we proceed with this agenda we state some properties of the optimal policy. These properties are simple extensions of the results in Section 2.2.2 for

the single-link system, thus, we omit the proofs.

- Theorem 3.2.1**
1. **(Monotonicity of $h(\mathbf{n})$)** For all i and all \mathbf{n} such that $\mathbf{R}(\mathbf{n} + \mathbf{e}_i) \leq \mathbf{C}$, we have $h(\mathbf{n}) \geq h(\mathbf{n} + \mathbf{e}_i)$, where \mathbf{e}_i denotes the i th unit vector.
 2. **(The infinite bandwidth case)** If there are no capacity constraints on all links in the network (i.e., $C_j = \infty, \forall j$), the optimal revenue is given by

$$J_\infty = \max_{\mathbf{u} \in \mathcal{U}} \sum_{i=1}^M \lambda_i(u_i) u_i,$$

and the optimal price vector is some constant \mathbf{u}_∞ that does not depend on the state \mathbf{n} . Furthermore, we have $J^* \leq J_\infty$.

3. There exists an optimal policy \mathbf{u}^* such that for every state \mathbf{n} , we have $\mathbf{u}^*(\mathbf{n}) \geq \mathbf{u}_\infty$.

The case of welfare maximization, can be treated similarly.

3.3 Static Pricing Policy

Possibly the simplest pricing policy is a *static* policy defined as the policy under which prices are fixed to some vector \mathbf{u} independent of the state of the system. According to this policy the system evolves as a continuous-time Markov chain which has a unique stationary distribution. In particular, the steady-state distribution has a product form and under a static pricing policy \mathbf{u} is given by (see [Kel91] and [Ros95])

$$\pi_{\mathbf{n}}(\mathbf{u}) = \mathbf{P}[\mathbf{n}(t) = \mathbf{n} \mid \mathbf{u}(t) = \mathbf{u}] = \frac{1}{G(\mathbf{u})} \prod_{i=1}^M \frac{(\rho_i(u_i))^{n_i}}{n_i!}, \quad \mathbf{n} \in \mathcal{S}, \quad (3.6)$$

where $G(\mathbf{u})$ is a normalizing constant given by

$$G(\mathbf{u}) = \sum_{\mathbf{n} \in \mathcal{S}} \prod_{i=1}^M \frac{(\rho_i(u_i))^{n_i}}{n_i!},$$

and $\rho_i(u_i) = \lambda_i(u_i)/\mu_i$ is the load offered by class i .

According to the static pricing policy the prices stay fixed which results in a constant arrival rate $\lambda(\mathbf{u})$ independent of the state of the system. As a result we can not eliminate demand by raising prices when available resources are not sufficient to accept a call. Thus, deviating from our earlier convention, we will be blocking calls that arrive to find no sufficient resources. Consequently, the blocking probability has to be taken into account when calculating revenue. The blocking probability of class i calls under static policy \mathbf{u} is given by

$$\mathbf{P}_{\text{loss}}^i(\mathbf{u}) = \sum_{\{\mathbf{n} \mid \mathbf{R}(\mathbf{n}+\mathbf{e}_i) \not\leq \mathbf{C}\}} \pi_{\mathbf{n}}(\mathbf{u}). \quad (3.7)$$

The optimal revenue by a static policy is given by

$$J_s = \max_{\mathbf{u} \in \mathcal{U}} J(\mathbf{u}) = \max_{\mathbf{u} \in \mathcal{U}} \sum_{i=1}^M \lambda_i(u_i) u_i (1 - \mathbf{P}_{\text{loss}}^i(\mathbf{u})), \quad (3.8)$$

and it can be no better than the optimal (dynamic) revenue, i.e., $J_s \leq J^*$.

The calculation of the optimal static revenue J_s and the corresponding optimal static policy \mathbf{u}_s suffers from a similar “curse of dimensionality” as in the case of dynamic policies. In particular, to calculate the blocking probability one needs to compute the steady-state probabilities $\pi_{\mathbf{n}}(\mathbf{u})$ which depend on the normalizing constant G . Computing this constant for networks with arbitrary topologies is an NP-complete problem (see [Lou90]). Efficient schemes exist for special topologies and the so call *reduced load approximation* can be used to approximate the blocking probabilities in arbitrary networks (see [Ros95]). Numerical difficulties, though, exist for the reduced load approximation for large systems. To overcome high dimensionality problems we are interested in scalable and efficient ways of computing “good” static policies.

For the case of welfare maximization, the same discussion applies, with $\lambda_i(u_i)u_i$ replaced by $\lambda_i(u_i)\mathbf{E}[U_i \mid U_i \geq u_i]$.

3.4 An Upper Bound on Optimal Performance

We will next develop an upper bound on the optimal revenue J^* . Such a bound is useful because it can help us bound the suboptimality gap of suboptimal policies we consider in this thesis. It will also be instrumental in establishing our asymptotic optimality results.

Let us denote by $u_i(\lambda_i)$ the inverse of the demand function $\lambda_i(u_i)$, which exists due to Assumption C. Let us also denote define $F_i(\lambda_i) \triangleq \lambda_i u_i(\lambda_i)$ and $F_i(u_i) \triangleq \lambda_i \mathbf{E}[U_i \mid U_i \geq u_i(\lambda_i)]$, $i = 1, \dots, M$, for the case of revenue and welfare maximization, respectively. We assume that the functions F_i are concave. This is true, for example, when the demand function $\lambda_i(u_i)$ is linear. The following theorem provides an upper bound on J^* .

Theorem 3.4.1 *Consider the following nonlinear optimization problem*

$$\begin{aligned} \max_{\lambda_i, n_i} \quad & \sum_{i=1}^M F_i(\lambda_i) \\ \text{s.t.} \quad & \lambda_i = n_i \mu_i, \quad i = 1, \dots, M, \\ & \sum_i n_i r_{ji} \leq C_j, \quad j = 1, \dots, L, \end{aligned} \tag{3.9}$$

and let J_{ub} denote the optimal objective value. If $F_i(\lambda_i)$ is a concave function for all $i = 1, \dots, M$, then $J^* \leq J_{\text{ub}}$.

Proof: Consider an optimal dynamic pricing policy \mathbf{u}^* . Without loss of generality, we assume that the price u_i^* becomes large enough (e.g., $u_{i,\text{max}}$) and the arrival rate $\lambda_i(u_i^*)$ is equal to zero, whenever the state \mathbf{n} is such that a class i call cannot be admitted (i.e., $\mathbf{R}(\mathbf{n} + \mathbf{e}_i) \not\leq \mathbf{C}$). In the system operating under the optimal policy, we can view the arrival rate, λ_i , and the number of class i customers in the system, n_i , as random variables. Let $\mathbf{E}[\cdot]$ denote the expectation with respect to the steady-state distribution under this particular policy \mathbf{u}^* . At any time, we have $\sum_i n_i r_{ji} \leq C_j, \forall j$, which implies that $\sum_i \mathbf{E}[n_i] r_{ji} \leq C_j, \forall j$. Furthermore, Little's law implies $\mathbf{E}[\lambda_i] = \mu_i \mathbf{E}[n_i]$. Thus, $\mathbf{E}[n_i], \mathbf{E}[\lambda_i], i = 1, \dots, M$, form a feasible solution of the problem (3.9). Using the concavity of F_i and Jensen's inequality, we

have

$$J_{\text{ub}} \geq \sum_{i=1}^M F_i(\mathbf{E}[\lambda_i]) \geq \sum_{i=1}^M \mathbf{E}[F_i(\lambda_i)] = J^*,$$

where the last equality used the optimality of the policy under consideration. ■

3.5 Asymptotic Optimality of Static Pricing

We will now proceed with establishing our main results for the model considered in Section 3.1, namely, the asymptotic optimality of static pricing and the derivation of guarantees on the suboptimality gap away from the limiting regime.

The limiting regime we will consider is one of “*many small users*”, in the sense that link capacities become large compared to the bandwidth of a typical call. More specifically, we start with a base system with finite demand function $\lambda(\mathbf{u})$ and finite capacity \mathbf{C} and then scale by increasing both demand and capacity by a scaling factor $c \geq 1$. We will use a superscript c to denote various quantities in the scaled system. In particular, in the scaled system the capacity is $\mathbf{C}^c = c(C_1, \dots, C_L)$ and the demand function is given by $\lambda^c(\mathbf{u}) = c(\lambda_1(u_1), \dots, \lambda_M(u_M))$. Note that in the revenue maximization problem we simply scale the given demand function. In the welfare maximization problem it suffices to scale the peak demand rate as $\lambda_0^c = c\lambda_0$ and keep unaltered the behaviour of the users summarized in the utility density function $f_i(u_i)$ (see Section 3.1). This results in a demand function $\lambda_i^c(u_i) = c\lambda_{0,i}\mathbf{P}[U_i \geq u_i]$. The remaining system parameters $\boldsymbol{\mu}$ and \mathbf{R} are held fixed. The base system corresponds to the case $c = 1$.

In the scaled system the upper bound, J_{ub}^c , is obtained by maximizing $\sum_i c\lambda_i(u_i)u_i$ in the revenue maximization case and $\sum_i c\lambda_i(u_i)\mathbf{E}[U_i | U_i \geq u_i]$ in the welfare maximization case. The constraints in the upper bound calculation become

$$\sum_i \frac{c\lambda_i(u_i)r_{ji}}{\mu_i} \leq cC_j, \quad \forall j, \tag{3.10}$$

which are identical to the constraints for the base system (cf. (3.9)). Hence, there exists

an optimal solution $\mathbf{u}_{\text{ub}}^* = (u_{\text{ub},1}^*, \dots, u_{\text{ub},M}^*)$, which is independent of c , and it holds that $J_{\text{ub}}^c = cJ_{\text{ub}}^1$. In proving our asymptotic optimality result we will first consider the blocking probabilities in the scaled system. We will use the convention that for any static policy \mathbf{u} for which $\lambda_i(u_i) = 0$, $\mathbf{P}_{\text{loss}}^i(\mathbf{u}) = 0$. We will denote by \mathcal{O} the set of classes with nonzero demand at \mathbf{u}_{ub}^* , i.e., $\mathcal{O} = \{i \in \{1, \dots, M\} \mid \lambda_i(u_{\text{ub},i}^*) > 0\}$. We will also denote by \mathcal{O}_j , $j = 1, \dots, L$, the set of classes $i \in \mathcal{O}$ that use link j , i.e., $\mathcal{O}_j = \{i \in \mathcal{O} \mid r_{ji} > 0\}$. We will assume that $\mathcal{O} \neq \emptyset$; otherwise $J_{\text{ub}} = 0$ which can only happen in the trivial cases that $\mathbf{C} = \mathbf{0}$ or $\lambda(u_i) = 0$ for all u_i and i . Recall also that we have assumed $\mathcal{R}_k \neq \emptyset$; otherwise class k can be eliminated from the system. All classes $i \notin \mathcal{O}$ are shut out of the system under \mathbf{u}_{ub}^* , do not contribute to the revenue or the social welfare, and according to our convention have zero blocking probability.

Proposition 3.5.1 *Consider either the revenue maximization problem or the welfare maximization problem and let \mathbf{u}_{ub}^* be the optimal solution to the upper bound problem in the scaled system with parameter c . For any $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_M) > \mathbf{0}$, consider the static policy \mathbf{u}^ε given by $u_i^\varepsilon = u_{\text{ub},i}^* + \varepsilon_i$, $i = 1, \dots, M$. Let $\mathbf{P}_{\text{loss}}^{k,c}(\mathbf{u}^\varepsilon)$ be the blocking probability of class k calls in the scaled system, under policy \mathbf{u}^ε . For every class $k \in \mathcal{O}$ and all c , we have*

$$\mathbf{P}_{\text{loss}}^{k,c}(\mathbf{u}^\varepsilon) \leq \sum_{j \in \mathcal{R}_k} \exp \left\{ \inf_{\theta \geq 0} \xi_{jk}^\varepsilon(c, \theta) \right\}, \quad (3.11)$$

where

$$\xi_{jk}^\varepsilon(c, \theta) \triangleq c \sum_{i \in \mathcal{O}_j} \Lambda_{ji}^\varepsilon(\theta) + \theta r_{jk}, \quad (3.12)$$

and

$$\Lambda_{ji}^\varepsilon(\theta) \triangleq \frac{\lambda_i(u_i^\varepsilon)(e^{\theta r_{ji}} - 1) - \theta r_{ji} \lambda_i(u_{\text{ub},i}^*)}{\mu_i}. \quad (3.13)$$

Furthermore, for all $k \in \mathcal{O}$ and $j \in \mathcal{R}_k$, $\inf_{\theta \geq 0} \xi_{jk}^\varepsilon(c, \theta) \rightarrow -\infty$ as $c \rightarrow \infty$ and

$$\lim_{c \rightarrow \infty} \mathbf{P}_{\text{loss}}^{k,c}(\mathbf{u}^\varepsilon) = 0. \quad (3.14)$$

Proof: Since $\varepsilon > 0$ and due to Assumption C, $\lambda_i^c(u_i^\varepsilon) = 0$ for all $i \notin \mathcal{O}$ and all c . Thus, no customer exists in the system from those classes $i \notin \mathcal{O}$ and according to our convention the corresponding blocking probabilities are zero. We will next concentrate on classes $i \in \mathcal{O}$.

Let n_i^c (respectively $n_{i,\infty}^c$) be the random variable which is equal to the number of active class i calls, in steady-state, in the scaled system, under prices \mathbf{u}^ε and with capacity $c\mathbf{C}$ (respectively, with infinite capacity). By defining the arrival processes in these two systems on a common probability space we can see that for all sample paths n_i^c is smaller than $n_{i,\infty}^c$. Using this fact, for any class $k \in \mathcal{O}$ we have

$$\begin{aligned} \mathbf{P}_{\text{loss}}^{k,c}(\mathbf{u}^\varepsilon) &= \mathbf{P} \left[\bigcup_{j \in \mathcal{R}_k} \sum_{i \in \mathcal{O}_j} r_{ji} n_i^c > cC_j - r_{jk} \right] \\ &\leq \mathbf{P} \left[\bigcup_{j \in \mathcal{R}_k} \sum_{i \in \mathcal{O}_j} r_{ji} n_{i,\infty}^c > cC_j - r_{jk} \right]. \end{aligned} \quad (3.15)$$

In the above, note that since $k \in \mathcal{O}$ and $\mathcal{R}_k \neq \emptyset$, there exists at least one $j \in \mathcal{R}_k$ and \mathcal{O}_j contains at least class k . Using the fact that \mathbf{u}_{ub}^* satisfies the constraint (3.10), we obtain

$$\begin{aligned} \mathbf{P} \left[\bigcup_{j \in \mathcal{R}_k} \sum_{i \in \mathcal{O}_j} r_{ji} n_{i,\infty}^c > cC_j - r_{jk} \right] &\leq \mathbf{P} \left[\bigcup_{j \in \mathcal{R}_k} \sum_{i \in \mathcal{O}_j} r_{ji} n_{i,\infty}^c > \sum_{i \in \mathcal{O}_j} \frac{c\lambda_i(u_{\text{ub},i}^*) r_{ji}}{\mu_i} - r_{jk} \right] \\ &\leq \sum_{j \in \mathcal{R}_k} \mathbf{P} \left[\sum_{i \in \mathcal{O}_j} r_{ji} n_{i,\infty}^c > \sum_{i \in \mathcal{O}_j} \frac{c\lambda_i(u_{\text{ub},i}^*) r_{ji}}{\mu_i} - r_{jk} \right], \end{aligned} \quad (3.16)$$

where the last inequality is due to the fact $\mathbf{P}[\bigcup_j X_j] \leq \sum_j \mathbf{P}[X_j]$.

Note next that the random variable $n_{i,\infty}^c$ is equal to the number of customers in an $M/M/\infty$ queue with arrival rate $c\lambda_i(u_i^\varepsilon)$ and service rate μ_i for each server. Its moment-generating function is

$$\mathbf{E}[e^{\theta n_{i,\infty}^c}] = e^{\frac{c\lambda_i(u_i^\varepsilon)}{\mu_i} (\theta - 1)}.$$

and by independence we obtain

$$\mathbf{E} \left[e^{\sum_{i \in \mathcal{O}_j} \theta r_{ji} n_{i,\infty}^c} \right] = \exp \left\{ c \sum_{i \in \mathcal{O}_j} \frac{\lambda_i(u_i^\epsilon)}{\mu_i} (e^{\theta r_{ji}} - 1) \right\}.$$

Using the Markov inequality and the above, for any $j \in \mathcal{R}_k$ and $\theta \geq 0$, we obtain

$$\begin{aligned} \mathbf{P} \left[\sum_{i \in \mathcal{O}_j} r_{ji} n_{i,\infty}^c > c \sum_{i \in \mathcal{O}_j} \frac{\lambda_i(u_{\text{ub},i}^*) r_{ji}}{\mu_i} - r_{jk} \right] \\ \leq \mathbf{E} \left[e^{\sum_{i \in \mathcal{O}_j} \theta r_{ji} n_{i,\infty}^c} \right] \exp \left\{ -\theta c \sum_{i \in \mathcal{O}_j} \frac{\lambda_i(u_{\text{ub},i}^*) r_{ji}}{\mu_i} + \theta r_{jk} \right\} \\ = \exp \left\{ c \sum_{i \in \mathcal{O}_j} \left(\frac{\lambda_i(u_i^\epsilon) (e^{\theta r_{ji}} - 1) - \theta r_{ji} \lambda_i(u_{\text{ub},i}^*)}{\mu_i} \right) + \theta r_{jk} \right\} \\ = \exp \left\{ c \sum_{i \in \mathcal{O}_j} \Lambda_{ji}^\epsilon(\theta) + \theta r_{jk} \right\} \\ = \exp \{ \xi_{jk}^\epsilon(c, \theta) \}. \end{aligned} \quad (3.17)$$

where $\Lambda_{ji}^\epsilon(\theta)$ and $\xi_{jk}^\epsilon(c, \theta)$ were defined in (3.13) and (3.12), respectively. Optimizing the right hand side of (3.17) over all $\theta \geq 0$ to obtain the tightest bound yields

$$\mathbf{P} \left[\sum_{i \in \mathcal{O}_j} r_{ji} n_{i,\infty}^c > c \sum_{i \in \mathcal{O}_j} \frac{\lambda_i(u_{\text{ub},i}^*) r_{ji}}{\mu_i} - r_{jk} \right] \leq \exp \left\{ \inf_{\theta \geq 0} \xi_{jk}^\epsilon(c, \theta) \right\}. \quad (3.18)$$

Combining (3.18) with (3.15) and (3.16) yields

$$\mathbf{P}_{\text{loss}}^{k,c}(\mathbf{u}^\epsilon) \leq \sum_{j \in \mathcal{R}_k} \exp \left\{ \inf_{\theta \geq 0} \xi_{jk}^\epsilon(c, \theta) \right\},$$

which establishes the exponential bound in (3.11).

Let us now consider what happens as $c \rightarrow \infty$. For large c , $\xi_{jk}^\epsilon(c, \theta)$ will be dominated

by $c \sum_{i \in \mathcal{O}}, \Lambda_{ji}^\varepsilon(\theta)$. At $\theta = 0$,

$$\begin{aligned} \sum_{i \in \mathcal{O}_j} \Lambda_{ji}^\varepsilon(0) &= 0, \\ \sum_{i \in \mathcal{O}_j} \left. \frac{\partial \Lambda_{ji}^\varepsilon(\theta)}{\partial \theta} \right|_{\theta=0} &= \sum_{i \in \mathcal{O}_j} \frac{r_{ji}(\lambda_i(u_i^\varepsilon) - \lambda_i(u_{\text{ub},i}^*))}{\mu_i}. \end{aligned}$$

From Assumption C, for every $i \in \mathcal{O}$ and any $\varepsilon > 0$ we have $\lambda_i(u_i^\varepsilon) < \lambda_i(u_{\text{ub},i}^*)$. Furthermore, for every $i \in \mathcal{O}_j$, $r_{ji} > 0$. Therefore, $\sum_{i \in \mathcal{O}_j} \Lambda_{ji}^\varepsilon(\theta)$ achieves its minimum over $\theta \geq 0$ at some $\theta_j^*(\varepsilon) > 0$ at which it holds $\sum_{i \in \mathcal{O}_j} \Lambda_{ji}^\varepsilon(\theta_j^*(\varepsilon)) < 0$. Note also that for all $j \in \mathcal{R}_k$

$$\inf_{\theta \geq 0} \xi_{jk}^\varepsilon(c, \theta) \leq \left[c \sum_{i \in \mathcal{O}_j} \Lambda_{ji}^\varepsilon(\theta_j^*(\varepsilon)) + \theta_j^*(\varepsilon) r_{jk} \right],$$

and for large enough c the right hand side of the above is $O(c \sum_{i \in \mathcal{O}_j} \Lambda_{ji}^\varepsilon(\theta_j^*(\varepsilon)))$ which converges to $-\infty$ as $c \rightarrow \infty$. This establishes (3.14). \blacksquare

Remarks:

1. It should be noted that for small values of c the bound in (3.11) could be trivial, meaning that the right hand side might be larger than one.
2. As $c \rightarrow \infty$, however, the bound in (3.11) converges to zero exponentially fast like $\exp\{c \sum_{i \in \mathcal{O}_j} \Lambda_{ji}^\varepsilon(\theta_j^*(\varepsilon))\}$, where $\theta_j^*(\varepsilon) = \arg \inf_{\theta \geq 0} \sum_{i \in \mathcal{O}_j} \Lambda_{ji}^\varepsilon(\theta)$ and $\sum_{i \in \mathcal{O}_j} \Lambda_{ji}^\varepsilon(\theta_j^*(\varepsilon)) < 0$.

We are now ready to state our asymptotic optimality result. We have seen that J_{ub}^c is linear in c . The optimal performance $J^{*,c}$ and the optimal performance J_s^c achieved by a static policy are also roughly linear in c . Thus, we will divide such quantities by c to make comparisons. The following theorem summarizes the result.

Theorem 3.5.2 *Consider either the revenue or the welfare maximization problem and as-*

sume that the functions $F_i(\lambda_i)$ are concave. Then,

$$\lim_{c \rightarrow \infty} \frac{1}{c} J_s^c = \lim_{c \rightarrow \infty} \frac{1}{c} J^{*,c} = \lim_{c \rightarrow \infty} \frac{1}{c} J_{\text{ub}}^c.$$

Proof: To simplify the exposition we will provide the proof for the revenue maximization case: welfare maximization can be treated similarly. For some $\varepsilon > 0$, let $\boldsymbol{\varepsilon} = \varepsilon \mathbf{e}$, where \mathbf{e} is the vector of all ones, and consider the static pricing policy \mathbf{u}^ε given by $u_i^\varepsilon = u_{\text{ub},i}^* + \varepsilon$, $i = 1, \dots, M$. Let $J^c(\mathbf{u}^\varepsilon)$ be the resulting average revenue, which is no more than the optimal static revenue J_s^c . Thus,

$$\lim_{c \rightarrow \infty} \frac{1}{c} J_s^c \geq \lim_{c \rightarrow \infty} \frac{1}{c} J^c(\mathbf{u}^\varepsilon) = \lim_{c \rightarrow \infty} \sum_{i=1}^M \lambda_i(u_i^\varepsilon) u_i^\varepsilon (1 - \mathbf{P}_{\text{loss}}^{i,c}(\mathbf{u}^\varepsilon)) = \sum_{i=1}^M \lambda_i(u_i^\varepsilon) u_i^\varepsilon.$$

In the last equality above we used the fact that for all $\varepsilon > 0$ demand is zero at \mathbf{u}^ε for all classes $i \notin \mathcal{O}$, and Proposition 3.5.1 for all classes $i \in \mathcal{O}$. Since the above inequality holds for any $\varepsilon > 0$, we take $\varepsilon \rightarrow 0$, which implies $\mathbf{u}^\varepsilon \rightarrow \mathbf{u}_{\text{ub}}^*$ and, by the continuity of the demand functions,

$$\lim_{c \rightarrow \infty} \frac{1}{c} J_s^c \geq \sum_{i=1}^M \lambda_i(u_{\text{ub},i}^*) u_{\text{ub},i}^* = J_{\text{ub}}^1.$$

On the other hand, due to the suboptimality of the static policy and Theorem 3.4.1, $J_s^c \leq J^{*,c} \leq J_{\text{ub}}^c = c J_{\text{ub}}^1$, and the result follows. ■

Theorem 3.5.2 establishes that in the limit $c \rightarrow \infty$ the upper bound of Theorem 3.4.1 is tight and the optimal solution of the upper bound problem, which is a static policy, is asymptotically optimal. Furthermore, Proposition 3.5.1 can be seen as characterizing the rate of convergence. This characterization allows us to determine how we should scale any given system to provide guarantees on the suboptimality gap of appropriately chosen static pricing policies. The following proposition describes the result. We state the result for the revenue maximization problem. It can be easily generalized to the welfare maximization

problem as well.

Proposition 3.5.3 *Consider either the revenue or the welfare maximization problem and assume that the functions $F_i(\lambda_i)$ are concave. Let \mathbf{u}_{ub}^* be the optimal solution to the upper bound problem of Theorem 3.4.1. For any $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_M) > \mathbf{0}$, consider the static policy $\mathbf{u}^\varepsilon = \mathbf{u}_{\text{ub}}^* + \boldsymbol{\varepsilon}$, and let $J^c(\mathbf{u}^\varepsilon)$ its performance. For any given $\delta > 0$, let $(c^*, \boldsymbol{\varepsilon}^*)$ be an optimal solution of the following optimization problem*

$$\begin{aligned} \min_{c, \boldsymbol{\varepsilon}} \quad & c & (3.19) \\ \text{s.t.} \quad & (1 + \delta) \sum_{i=1}^M \lambda_i(u_i^\varepsilon) u_i^\varepsilon \left(1 - \sum_{j \in \mathcal{R}_i} \exp \left\{ \inf_{\theta \geq 0} \xi_{ji}^\varepsilon(c, \theta) \right\} \right) \geq \sum_{i=1}^M \lambda_i(u_{\text{ub},i}^*) u_{\text{ub},i}^* \\ & \boldsymbol{\varepsilon} \geq \mathbf{0}, \end{aligned}$$

where $\xi_{ji}^\varepsilon(c, \theta)$ is defined in (3.12). Then, the performance of the static policy $\mathbf{u}^{\boldsymbol{\varepsilon}^*}$ in the c^* -scaled system satisfies

$$\frac{J^{*,c^*} - J^{c^*}(\mathbf{u}^{\boldsymbol{\varepsilon}^*})}{J^{c^*}(\mathbf{u}^{\boldsymbol{\varepsilon}^*})} \leq \delta. \quad (3.20)$$

Proof: Fix some $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_M) > \mathbf{0}$ and consider the static pricing policy \mathbf{u}^ε resulting in average revenue or social welfare equal to $J^c(\mathbf{u}^\varepsilon)$. Due to the suboptimality of the static policy, Theorem 3.4.1, and Proposition 3.5.1, $J^c(\mathbf{u}^\varepsilon)$ satisfies

$$\begin{aligned} \sum_{i=1}^M \lambda_i(u_{\text{ub},i}^*) u_{\text{ub},i}^* &= J_{\text{ub}}^1 = \frac{1}{c} J_{\text{ub}}^c \geq \frac{1}{c} J^c(\mathbf{u}^\varepsilon) = \\ & \sum_{i=1}^M \lambda_i(u_i^\varepsilon) u_i^\varepsilon (1 - \mathbf{P}_{\text{loss}}^{i,c}(\mathbf{u}^\varepsilon)) \geq \sum_{i=1}^M \lambda_i(u_i^\varepsilon) u_i^\varepsilon \left(1 - \sum_{j \in \mathcal{R}_i} \exp \left\{ \inf_{\theta \geq 0} \xi_{ji}^\varepsilon(c, \theta) \right\} \right). \end{aligned} \quad (3.21)$$

Using the same argument as in the proof of Theorem 3.5.2, we can first take $c \rightarrow \infty$ in the right hand side of the above and bring the blocking probabilities to zero. If we then take $\boldsymbol{\varepsilon} \rightarrow \mathbf{0}$ we conclude that the right hand side of (3.21) converges to its left hand side and the inequality is satisfied with equality. Thus for any $\delta > 0$, we can find a scaling factor c and a static policy \mathbf{u}^ε , such that $\frac{J_{\text{ub}}^c - J^c(\mathbf{u}^\varepsilon)}{J^c(\mathbf{u}^\varepsilon)} \leq \delta$, by solving the optimization problem in (3.19).

More specifically, if (c^*, ϵ^*) is an optimal solution of (3.19) we have

$$\frac{J_{\text{ub}}^{c^*} - J^{c^*}(\mathbf{u}^{\epsilon^*})}{J^{c^*}(\mathbf{u}^{\epsilon^*})} \leq \delta,$$

and the desired result follows since $J_{\text{ub}}^{c^*} \geq J^{*,c^*}$. ■

3.6 Structure of Asymptotically Optimal Static Pricing Policy

As we have seen the optimal solution of the upper bound problem of Theorem 3.4.1 provides a static pricing policy which is asymptotically optimal in the regime of many small users we considered. We will next characterize its structure. To that end, we will view the upper bound problem (3.9) as one involving optimization with respect to u_i , rather than λ_i . We will also write n_i in the form $\lambda_i(u_i)/\mu_i$. We start with the revenue maximization problem.

3.6.1 Revenue Maximization

The upper bound problem becomes

$$\begin{aligned} \max_{u_i} \quad & \sum_i \lambda_i(u_i)u_i \\ \text{s.t.} \quad & \sum_i \frac{\lambda_i(u_i)r_{ji}}{\mu_i} \leq C_j, \quad \forall j. \end{aligned} \tag{3.22}$$

Let $\mathbf{q} = (q_1, \dots, q_L) \geq 0$ be the Lagrange multiplier vector, where q_j is associated with the capacity constraint on link j . Writing the problem in (3.22) as a minimization problem, its Lagrangean function becomes

$$\mathcal{L}(\mathbf{u}, \mathbf{q}) = - \sum_{i=1}^M \lambda_i(u_i)u_i + \sum_{j=1}^L q_j \left(\sum_{i=1}^M \frac{\lambda_i(u_i)r_{ji}}{\mu_i} - C_j \right).$$

Assuming an interior solution $u_i \in (0, u_{i,\max})$, u_i should minimize

$$\left(-\lambda_i(u_i)u_i + \sum_j q_j \lambda_i(u_i) \frac{r_{ji}}{\mu_i} \right). \quad (3.23)$$

Therefore, from the first order optimality condition we obtain

$$u_i = -\frac{\lambda_i(u_i)}{d\lambda_i(u_i)/du_i} + \sum_{j=1}^L q_j \frac{r_{ji}}{\mu_i}, \quad \forall i. \quad (3.24)$$

This structure is insightful. The first term is the reciprocal of the demand elasticity, prescribing that we should charge more to classes with more inelastic demand. The second term is a usage-based charge. Notice, that by complementary slackness conditions $q_j = 0$, if the corresponding constraint is not active, which can be interpreted as link j not being congested. On the other hand, if link j is congested (i.e., the corresponding constraint is satisfied with equality at the optimal solution), we charge each class a price $q_j > 0$ per unit of volume on link j . Here, we define as volume the quantity r_{ji}/μ_i , which is the bandwidth occupied times the expected holding time. Thus, the second term in (3.24) includes a charge for volume on congested links along the route \mathcal{R}_i of class i .

This pricing structure is appealing from an implementation point of view. Large (backbone) networks might typically accommodate many service classes (number of offered services times number of origin-destination pairs), but consist of a relatively small number of links. Later on we will use this pricing structure and optimize over the shadow prices \mathbf{q} to obtain near-optimal performance even away from the limiting regime.

3.6.2 Welfare Maximization

The case of welfare maximization can be treated similarly. Using (3.4) an interior solution $u_i \in (0, u_{i,\max})$ should minimize

$$\left(-\lambda_{0,i} \int_{u_i}^{u_{i,\max}} v f_i(v) dv + \sum_j q_j \lambda_i(u_i) \frac{r_{ji}}{\mu_i} \right),$$

which is analogous to the condition in (3.23) for the revenue maximization case. Therefore, from the first order optimality condition we obtain

$$\lambda_{0,i} u_i f_i(u_i) + \sum_j q_j \frac{r_{ji}}{\mu_i} \frac{d\lambda_i(u_i)}{du_i} = 0,$$

which, by using (3.3), becomes

$$u_i = \sum_{j=1}^L q_j \frac{r_{ji}}{\mu_i}, \quad \forall i. \quad (3.25)$$

As in the revenue maximization case $q_j = 0$ for non-active constraints, thus, the pricing structure in (3.25) prescribes a usage-based charge for volume on all congested links along the route \mathcal{R}_i of class i .

3.7 Large Scale Problems

In this section we discuss how the pricing policies we have considered so far can be computed and applied to large scale systems.

Large networks consist of numerous classes (equal to the number of offered services times the origin-destination pairs) and many links with large capacities. As a result, the state space $\mathcal{S} = \{\mathbf{n} \mid \mathbf{R}\mathbf{n} \leq \mathbf{C}\}$ becomes enormous and it is intractable to compute the optimal (dynamic) policy. One could potentially leverage recent approximate dynamic programming techniques to compute an approximately optimal dynamic policy. This direction has been successfully explored in Section 2.6 and [PT00], and can be generalized in the network setting. The sheer dimensionality of the network problem though, makes the computational effort more challenging.

In this thesis we are focusing on static pricing policies because they are simpler and have significant implementation advantages over dynamic ones; we have outline those in the Introduction. As we commented in Section 3.3, computing the optimal static policy exactly is also computationally intractable. Instead we will experiment with the following

two approaches to compute effective static pricing policies:

1. **Policy from the Upper Bound.** As we have seen the optimal solution of the upper bound problem in (3.9) forms a static pricing policy for our original model of Section 3.1. we have seen that in the limiting regime of many small users this policy is asymptotically optimal. Furthermore, it is quite easy to obtain; their computation amounts to solving a nonlinear programming problem with $O(L)$ linear constraints and $O(M)$ decision variables; for which effective algorithms exist.
2. **Using the structure of asymptotically optimal static policy.** A concern with the static policy from the upper bound is that it might not perform as well away from the limiting regime. Some earlier experience with the single-link problem indicates that its asymptotically optimal structure (given in Section 3.6) is effective away from the limiting regime but the values of the various parameters might not be appropriate away from the limit. More specifically, note that the structure of the policies in Section 3.6 depends on the selection of a set of shadow prices (Lagrange multipliers) for the resources at all congested links of the network. To improve upon the policy obtained from the upper bound we seek to optimize the performance objective (revenue or social welfare) over those shadow prices. To that end, we employ a simulation-based method outlined in the following subsection.

3.7.1 A Simulation-based Method

The underlying idea is rather simple and is the basis of so-called perturbation analysis techniques (see Cassandras [Cas93], Fu and Hu [FH97] and references therein). We adopt the structure of the policies of Sections 3.6 and 4.4 and view them as functions of the Lagrange multipliers q_j , $j = 1, \dots, L$. During the course of a simulation of the system we obtain “gradient information” which is used to optimize over q_j 's. To that end, we will apply a technique developed by Marbach and Tsitsiklis [MT01]. Alternatively, we could optimize over prices directly, but the dual approach of optimizing over the Lagrange multipliers q_j is

more preferable in large networks since, typically, the number of classes is much larger than the number of links. In the remaining of this section we will focus on revenue maximization in the original model of Section 3.1. The discussion readily extends to welfare maximization.

To fix our notation for discussing the simulation-based optimization method, consider the policy structure of Equation (3.24), i.e.,

$$u_i = -\frac{\lambda_i(u_i)}{d\lambda_i(u_i)/du_i} + \sum_{j=1}^L q_j \frac{r_{ji}}{\mu_i}, \quad i = 1, \dots, M.$$

To explicitly denote that u_i is a function of $\mathbf{q} = (q_1, \dots, q_L)$, we will write $u_i(\mathbf{q})$ and be referring to this as the “ \mathbf{q} policy”. The demand function of class i also becomes a function of \mathbf{q} , we will write $\lambda_i(\mathbf{q})$. In the uniformized discrete-time Markov chain discussed in Section 3.2, the transition probability from state \mathbf{n} to state $\bar{\mathbf{n}}$ is given by

$$p_{\mathbf{n}\bar{\mathbf{n}}}(\mathbf{q}) = \mathbf{P}[\mathbf{n}(t+1) = \bar{\mathbf{n}} | \mathbf{n}(t) = \mathbf{n}; \mathbf{q}] = \begin{cases} \lambda_i(\mathbf{q})/\nu & \text{if } \bar{\mathbf{n}} = \mathbf{n} + \mathbf{e}_i, i \notin \mathcal{C}(\mathbf{n}), \\ n_i \mu_i / \nu & \text{if } \bar{\mathbf{n}} = \mathbf{n} - \mathbf{e}_i, \\ 1 - \sum_{i \notin \mathcal{C}(\mathbf{n})} \frac{\lambda_i(\mathbf{q})}{\nu} - \sum_i \frac{n_i \mu_i}{\nu} & \text{if } \bar{\mathbf{n}} = \mathbf{n}, \\ 0 & \text{otherwise,} \end{cases}$$

where $\mathbf{n}(t)$ is the state at discrete time t and $\mathcal{C}(\mathbf{n}) = \{i \mid \mathbf{R}(\mathbf{n} + \mathbf{e}_i) \not\subseteq \mathbf{C}\}$, defined in Section 3.2, is the set of classes whose calls cannot be admitted in state \mathbf{n} . The instantaneous revenue rate at state \mathbf{n} is given by

$$g_{\mathbf{n}}(\mathbf{q}) = \sum_{i \notin \mathcal{C}(\mathbf{n})} \lambda_i(\mathbf{q}) u_i(\mathbf{q}),$$

and the expected long-term average revenue under policy \mathbf{q} is

$$J(\mathbf{q}) = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbf{E} \left[\sum_{t=0}^T g_{\mathbf{n}(t)}(\mathbf{q}) \right].$$

The simulation-based optimization algorithm of [MT01] is for unconstrained problems

and requires a number of technical assumptions to guarantee convergence, including that $p_{n\bar{n}}(\mathbf{q})$ and $g_n(\mathbf{q})$ are bounded, twice differentiable, and have bounded first and second derivatives, and that $\frac{\nabla p_{n\bar{n}}(\mathbf{q})}{p_{n\bar{n}}(\mathbf{q})}$ is bounded. A demand function satisfying our Assumption C might violate these properties. Consequently, for the purposes of the simulation-based optimization we will replace $\lambda_i(u_i)$ with another (smoother) function $\hat{\lambda}_i(u_i)$ satisfying

1. $\hat{\lambda}_i(u_i) \geq 0$ for all $u_i \in \mathbb{R}$;
2. $\hat{\lambda}_i(u_i)$ is strictly decreasing, bounded, and has bounded third derivative for all $u_i \in \mathbb{R}$;
3. for any given $\epsilon > 0$, $\hat{\lambda}_i(u_i)$ satisfies $|\hat{\lambda}_i(u_i) - \lambda_i(u_i)| \leq \epsilon$ for all $u_i \geq 0$; and
4. $\hat{\lambda}_i(u_i)u_i$ and $\frac{d\hat{\lambda}_i(u_i)/du_i}{\hat{\lambda}_i(u_i)}$ are bounded for all $u_i \geq 0$.

Such a modified demand function satisfies the properties on $p_{n\bar{n}}(\mathbf{q})$ and $g_n(\mathbf{q})$ mentioned above. Note that when $u_i < 0$, we have $\hat{\lambda}_i(u_i)u_i < 0$ and incur a negative revenue, which is worse than setting $u_i = 0$. Therefore, the optimal solution of the simulation-based optimization will correspond to nonnegative prices.

To provide an example on how such a smooth demand function $\hat{\lambda}_i(u_i)$ can be constructed consider the case of linear demand, i.e.,

$$\lambda_i(u_i) = \begin{cases} \lambda_{i,0} - \lambda_{i,1}u_i, & 0 \leq u_i \leq u_{i,\max} = \frac{\lambda_{i,0}}{\lambda_{i,1}}, \\ 0 & u_i > u_{i,\max}. \end{cases} \quad (3.26)$$

This can be replaced with

$$\hat{\lambda}_i(u_i) = \lambda_{i,0} - \frac{2\lambda_{i,0}u_i}{\sqrt{\left(u_i - \frac{\lambda_{i,0}}{\lambda_{i,1}}\right)^2 + \epsilon} + \sqrt{\left(u_i + \frac{\lambda_{i,0}}{\lambda_{i,1}}\right)^2 + \epsilon}}, \quad \forall u_i \in \mathbb{R}, \quad (3.27)$$

and some $\epsilon > 0$, which satisfies all requirements 1–4 stated above. By selecting a small enough ϵ , we can make $\hat{\lambda}_i(u_i)$ arbitrarily close to $\lambda_i(u_i)$, for all $u_i \geq 0$ (see Figure 3.1 for an example).




Figure 3.1: $\lambda_i(u_i)$ and $\hat{\lambda}_i(u_i)$ of Equations (3.26) and (3.27), respectively, when $\lambda_{i,0} = 40$, $\lambda_{i,1} = 4$, and $\epsilon = 1$.

For every $\mathbf{q} \in \mathbb{R}^L$, let $P(\mathbf{q})$ be the transition probability matrix of the Markov chain with entries $p_{\mathbf{n}\bar{\mathbf{n}}}(\mathbf{q})$. Let $\mathcal{P} = \{P(\mathbf{q}) \mid \mathbf{q} \in \mathbb{R}^L\}$ be the set of all such matrices and let $\bar{\mathcal{P}}$ be its closure. Obviously, the Markov chain corresponding to every $P \in \bar{\mathcal{P}}$ is aperiodic. The empty state $\mathbf{n}^0 = \mathbf{0}$ is recurrent for every such Markov chain. Moreover, the system has finite number of states, the service rate is positive for all classes, thus, there exists a number N_0 , which is no more than the total number of states, such that for every state \mathbf{n} , and every collection of $\{P_1, \dots, P_{N_0}\}$ of N_0 matrices in $\bar{\mathcal{P}}$, we have

$$\sum_{n=1}^{N_0} \left[\prod_{i=1}^n P_i \right]_{\mathbf{n}\mathbf{n}^0} > 0,$$

where $[\mathbf{A}]_{\mathbf{n},\mathbf{n}^0}$ denotes the $(\mathbf{n}, \mathbf{n}^0)$ element of the matrix \mathbf{A} .

Given these observations and with the modified demand function $\hat{\lambda}_i(u_i)$, our setting satisfies all the assumptions in [MT01]. We will apply the following simulation-based optimization algorithm proposed there. The algorithm optimizes $J(\mathbf{q})$ over \mathbf{q} by estimating the gradient $\nabla J(\mathbf{q})$ and updating \mathbf{q} in a single sample path of the simulation. The update can be taken either at visits to the recurrent state \mathbf{n}^0 , or at every time step. We provide the algorithm that updates \mathbf{q} and \tilde{J} at every time step, where \tilde{J} is the estimate of $J(\mathbf{q})$.

Algorithm 3.7.1 ([MT01]) *At a typical time t , the state is $\mathbf{n}(t)$, and the values of $\mathbf{q}(t)$, $\mathbf{z}(t)$, and $\bar{J}(t)$ are available from the previous iteration. $\mathbf{z}(t)$ is a vector of the same dimension as \mathbf{q} . We update \mathbf{q} and \bar{J} according to*

$$\begin{aligned}\mathbf{q}(t+1) &= \mathbf{q}(t) + \gamma_t \left(\nabla g_{\mathbf{n}(t)}(\mathbf{q}) + (g_{\mathbf{n}(t)}(\mathbf{q}) - \bar{J}(t))\mathbf{z}(t) \right) \\ \bar{J}(t+1) &= \bar{J}(t) + \gamma_t (g_{\mathbf{n}(t)}(\mathbf{q}) - \bar{J}(t)),\end{aligned}$$

We simulate a transition to the next state $\mathbf{n}(t+1)$ according to the transition probabilities $p_{\mathbf{n}\bar{\mathbf{n}}}(\mathbf{q}(t+1))$, and finally update \mathbf{z} by letting

$$\mathbf{z}(t+1) = \begin{cases} 0, & \text{if } \mathbf{n}(t+1) = \mathbf{n}^0, \\ \mathbf{z}(t) + \frac{\nabla p_{\mathbf{n}(t)\mathbf{n}(t+1)}(\mathbf{q}(t))}{p_{\mathbf{n}(t)\mathbf{n}(t+1)}(\mathbf{q}(t))}, & \text{otherwise.} \end{cases}$$

The convergence of the algorithm (w.p.1) to a stationary point of $J(\mathbf{q})$ (i.e., a point where the gradient is zero) is guaranteed by selecting appropriately a stepsize γ_t (e.g., γ_t is diminishing as in $\gamma_t = 1/t$).

3.7.2 Numerical Results

In this section we tackle, numerically, some illustrative network pricing problems using the ideas discussed on the model of Section 3.1. We will present revenue maximization problems. The qualitative conclusions would not be much different in welfare maximization.

A Five-node Network

Our first example, depicted in Figure 3.2, is a network with 5 nodes, 4 links, and 12 service classes. The link capacities are shown in the figure. The parameters for the 12 classes of services, are listed in Table 3.1. The demand functions in this example are linear and have the form of Equation (3.26). Table 3.2 compares the upper bound J_{ub} (cf. Theorem 3.4.1) with the two policies proposed in Section 3.7, namely, the policy from the upper bound and




Figure 3.2: A 5-node, 4-link network with 12 service classes.

Class i	Nodes	Links	Bandwidth Requirement r_i	Demand Function $\lambda_i(u_i)$	μ_i
1	1, 0, 2	1, 2	1	$150 - 75u_1$	5
2	1, 0, 3	1, 3	1	$150 - 60u_2$	5
3	1, 0, 4	1, 4	1	$250 - 125u_3$	5
4	2, 0, 3	2, 3	1	$140 - 40u_4$	5
5	2, 0, 4	2, 4	1	$300 - 150u_5$	5
6	3, 0, 4	3, 4	1	$150 - 35u_6$	5
7	1, 0, 2	1, 2	5	$6 - u_7$	1
8	1, 0, 3	1, 3	5	$7 - 1.2u_8$	1
9	1, 0, 4	1, 4	5	$6.6 - 0.8u_9$	1
10	2, 0, 3	2, 3	5	$6 - 0.6u_{10}$	1
11	2, 0, 4	2, 4	5	$6 - 0.6u_{11}$	1
12	3, 0, 4	3, 4	5	$6 - 0.5u_{12}$	1

Table 3.1: The services provided by the network of Figure 3.2.

the policy from the simulation-based optimization approach. The corresponding prices are given in Table 3.3.

We conclude that the optimized version (via the simulation-based optimization approach) of our asymptotically optimal static pricing policy is quite close to the optimal. Note that the percentage gap in Table 3.2 is an upper bound on the suboptimality gap. It should be noted that even this rather small network has large enough state space for computing the optimal policy.

It is perhaps of interest to use Proposition 3.5.3 to compute by how much we should scale the network to achieve a given suboptimality gap. Using the notation introduced there, for

J_{ub}	$J(\mathbf{u}_{\text{ub}}^*)$	J_{sim}	$\frac{J_{\text{ub}} - J_{\text{sim}}}{J_{\text{ub}}} \times 100\%$
805.49	757.29	778.56	3.34%

Table 3.2: Comparing the various policies for the network of Figure 3.2. We use $J(\mathbf{u}_{\text{ub}}^*)$ and J_{sim} to denote the performance of the policy obtained from the upper bound problem and the simulation-based optimization approach, respectively.

Policy	u_1	u_2	u_3	u_4	u_5	u_6	u_7	u_8	u_9	u_{10}	u_{11}	u_{12}
\mathbf{u}_{ub}^*	1.03	1.28	1.03	1.75	1.00	2.14	3.82	3.74	4.95	5.00	5.00	6.00
$\mathbf{u}_{\text{sim}}^*$	1.08	1.31	1.08	1.81	1.08	2.21	5.05	4.51	6.17	6.57	7.02	7.57

Table 3.3: The prices for the network of Figure 3.2 under the policy obtained from the upper bound problem (\mathbf{u}_{ub}^*) and the policy obtained from the simulation-based optimization approach ($\mathbf{u}_{\text{sim}}^*$).

the network of Figure 3.2 we compute that a suboptimality gap of $\delta = 0.1$ is guaranteed by scaling the network by $c = 10.75$ and using policy \mathbf{u}^ϵ with $\epsilon = 0.32\mathbf{e}$, where \mathbf{e} is the vector of all ones. Similarly, $\delta = 0.05$ is achieved with $c = 25.92$ and $\epsilon = 0.23\mathbf{e}$. Finally, $\delta = 0.01$ is achieved with $c = 211.82$ and $\epsilon = 0.1\mathbf{e}$. Note that for simplicity of the calculations involved we only considered $\epsilon = \epsilon\mathbf{e}$ in the optimization problem (3.19). The results can be improved by considering arbitrary ϵ . Clearly, these guarantees come from (crude) bounds on the blocking probability and are not meant to be very tight. Our optimized policy ($\mathbf{u}_{\text{sim}}^*$), for example, would be much closer to optimal in each of those scaled systems. Nevertheless, Proposition 3.5.3 provides a simple way to quickly assess efficiency gains by scaling the system.

A Large-scale Network

The second example, depicted in Figure 3.3, is a network of a larger size (perhaps comparable to a backbone network in the U.S.). It consists of 9 nodes, 13 links, and provides 59 classes of services, the parameters of which are listed in Table 3.4.

i	Nodes (Links)	τ_i	$\lambda_i(u_i)$	μ_i	i	Nodes (Links)	τ_i	$\lambda_i(u_i)$	μ_i
1	1, 0 (1)	1	$300 - 120u_1$	2	2	2, 0 (2)	1	$400 - 120u_2$	2
3	2, 1 (8)	1	$200 - 40u_3$	2	4	3, 0 (3)	1	$250 - 50u_4$	2
5	5, 0 (5)	1	$400 - 80u_5$	2	6	5, 1 (10)	1	$300 - 60u_6$	2
7	5, 0, 3 (5, 3)	1	$100 - 16u_7$	2	8	5, 1, 4 (10, 9)	1	$200 - 50u_8$	2
9	6, 1, 0 (11, 1)	1	$200 - 40u_9$	1	10	6, 1 (11)	1	$200 - 40u_{10}$	1

i	Nodes (Links)	τ_i	$\lambda_i(u_i)$	μ_i	i	Nodes (Links)	τ_i	$\lambda_i(u_i)$	μ_i
11	6, 1, 5 (11, 10)	1	$150 - 30u_{11}$	1	12	7, 0, 1 (6, 1)	1	$300 - 40u_{12}$	1
13	8, 0 (7)	1	$300 - 40u_{13}$	1	14	8, 0, 1 (7, 1)	1	$300 - 40u_{14}$	1
15	8, 0, 2 (7, 2)	1	$200 - 40u_{15}$	1	16	8, 0, 4 (7, 4)	1	$100 - 20u_{16}$	1
17	1, 0 (1)	2	$100 - 10u_{17}$	4	18	2, 0 (2)	2	$80 - 8u_{18}$	4
19	2, 1 (8)	2	$100 - 8u_{19}$	4	20	3, 0 (3)	2	$100 - 10u_{20}$	4
21	3, 0, 1 (3, 1)	2	$80 - 10u_{21}$	4	22	4, 0 (4)	2	$120 - 12u_{22}$	2
23	4, 1 (9)	2	$120 - 10u_{23}$	2	24	4, 1, 2 (9, 8)	2	$100 - 12u_{24}$	2
25	5, 0 (5)	2	$80 - 8u_{25}$	2	26	5, 1 (10)	2	$80 - 10u_{26}$	2
27	5, 1, 2 (10, 8)	2	$100 - 8u_{27}$	2	28	5, 3 (5, 3)	2	$80 - 8u_{28}$	2
29	5, 1, 4 (10, 9)	2	$80 - 10u_{29}$	2	30	6, 1, 0 (11, 1)	2	$80 - 6u_{30}$	1
31	6, 1 (11)	2	$100 - 10u_{31}$	1	32	6, 1, 2 (11, 8)	2	$100 - 8u_{32}$	1
33	6, 8, 3 (13, 12)	2	$100 - 12u_{33}$	1	34	8, 0, 1 (7, 1)	2	$80 - 8u_{34}$	2
35	8, 0, 2 (7, 2)	2	$60 - 6u_{35}$	2	36	8, 0, 5 (7, 5)	2	$60 - 8u_{36}$	2
37	8, 6 (13)	2	$100 - 10u_{37}$	2	38	1, 0 (1)	4	$40 - 4u_{38}$	1
39	2, 0 (2)	4	$60 - 5u_{39}$	1	40	2, 1 (8)	4	$80 - 12u_{40}$	1
41	3, 0 (3)	4	$40 - 4u_{41}$	1	42	3, 0, 1 (3, 1)	4	$40 - 5u_{42}$	1
43	3, 0, 2 (3, 2)	4	$40 - 6u_{43}$	1	44	4, 0 (4)	4	$60 - 3u_{44}$	1
45	4, 1 (9)	4	$60 - 4u_{45}$	1	46	4, 1, 2 (9, 8)	4	$60 - 4u_{46}$	1
47	5, 0 (5)	4	$40 - 4u_{47}$	1	48	5, 1 (10)	4	$40 - 2u_{48}$	1
49	5, 0, 4 (5, 4)	4	$50 - 5u_{49}$	1	50	6, 1, 0 (11, 1)	4	$50 - 5u_{50}$	4
51	6, 1 (11)	4	$50 - 4u_{51}$	4	52	6, 1, 2 (11, 8)	4	$60 - 4u_{52}$	4
53	6, 8, 3 (13, 12)	4	$60 - 6u_{53}$	4	54	6, 1, 5 (11, 10)	4	$30 - 3u_{54}$	2
55	7, 0 (6)	4	$60 - 3u_{55}$	4	56	7, 0, 2 (6, 2)	4	$40 - 4u_{56}$	2
57	7, 0, 3 (6, 3)	4	$20 - 4u_{57}$	2	58	7, 0, 4 (6, 4)	4	$30 - 3u_{58}$	2
59	8, 0 (7)	4	$60 - 6u_{59}$	2					

Table 3.4: The services provided by the network of Figure 3.3.

Table 3.5 compares the upper bound J_{ub} (cf. Theorem 3.4.1) with the two policies proposed in Section 3.7, namely, the policy from the upper bound and the policy from the simulation-based optimization approach. Again, we observe that J_{sim} is quite close to the

J_{ub}	$J(\mathbf{u}_{ub}^*)$	J_{sim}	$\frac{J_{ub} - J_{sim}}{J_{ub}} \times 100\%$
12597.6	12117.4	12209.6	3.08%

Table 3.5: Comparing the various policies for the network of Figure 3.3. We maintain the notation of Table 3.2.

optimal.

As in the first example, we use Proposition 3.5.3 to compute by how much we should




Figure 3.3: A 9-node, 13-link network with 59 service classes. The link capacities are $C_1 = 840$, $C_2 = 420$, $C_3 = 420$, $C_4 = 420$, $C_5 = 420$, $C_6 = 300$, $C_7 = 420$, $C_8 = 420$, $C_9 = 420$, $C_{10} = 420$, $C_{11} = 420$, $C_{12} = 210$, and $C_{13} = 420$.

scale the network to achieve a given suboptimality gap. We obtain that a suboptimality gap of $\delta = 0.1$ is guaranteed by scaling the network by $c = 3.30$ and using policy \mathbf{u}^ϵ with $\epsilon = 0.59\mathbf{e}$. Similarly, $\delta = 0.05$ is achieved with $c = 10.74$ and $\epsilon = 0.38\mathbf{e}$. Finally, $\delta = 0.01$ is achieved with $c = 218.07$ and $\epsilon = 0.12\mathbf{e}$. The first two cases ($\delta = 0.1, 0.05$) yield scaling factors that are even smaller than the corresponding ones in the first example.

Chapter 4

Pricing in Communication Networks with Demand Substitution Effects

In this chapter we will extend the model we have considered so far to incorporate demand substitution effects. In particular, the model introduced in Section 3.1 assumes that the demand of each class $\lambda_i(u_i)$ is function of the price for that class only. We are interested in considering the situation where users might decide to use another class of service as a (non-perfect) substitute of their desired class if the latter one ends up being very expensive. Our main results extend to this situation as well. We will present a model that accounts for such substitution effects in Section 4.1. Following the development of the previous chapters, we will develop an upper bound on the optimal performance in Section 4.2, establish the asymptotic optimality of static pricing in Section 4.3, and characterize the structure of the asymptotically optimal static policy in Section 4.4. Numerical results are in Section 4.5.

4.1 The Model with Demand Substitution Effects

The model is in fact identical to the one introduced in Section 3.1, with the exception that demand for each class i , $i = 1, \dots, M$, is not only a function of u_i , but of the whole price vector \mathbf{u} , i.e., $\lambda(\mathbf{u}) = (\lambda_1(\mathbf{u}), \dots, \lambda_M(\mathbf{u}))$. We will maintain the rest of the notation that was introduced in Section 3.1. We will denote the load on link j by

$$\rho_j(\mathbf{u}) \triangleq \sum_{i=1}^M \frac{\tau_{ji} \lambda_i(\mathbf{u})}{\mu_i C_j}, \quad j = 1, \dots, L.$$

We will be making the following assumption.

Assumption D

1. If $\lambda_i(\mathbf{u}) > 0$, then $\frac{\partial \lambda_i(\mathbf{u})}{\partial u_i} < 0$, for $i = 1, \dots, M$;
2. $\frac{\partial \lambda_i(\mathbf{u})}{\partial u_k} \geq 0$, for $k \neq i, k = 1, \dots, M$;
3. if $\lambda_i(\mathbf{u}) > 0$, then $\sum_{k=1}^M \frac{\partial \lambda_k(\mathbf{u})}{\partial u_i} < 0$, for $i = 1, \dots, M$;
4. $\lim_{\mathbf{u} \rightarrow \infty} \lambda_i(\mathbf{u}) = 0$, for $i = 1, \dots, M$, where $\mathbf{u} \rightarrow \infty$ means $\min_i |r_{j,i} > 0| u_i \rightarrow \infty$.

Assumption D-1 indicates that demand for any class is a strictly decreasing function of its own price. Assumption D-2 indicates that substitution among classes can take place, in the sense that the increase of the price for a class can increase the demand for other classes. Assumption D-3 states that only a fraction of demand lost for a class appears as demand for other classes (due to substitution). Assumption D-4 expresses the condition that as all prices increase, the demand will eventually decrease to zero for all classes.

As an example, following linear demand functions with substitution effects between two classes satisfy Assumption D:

$$\begin{aligned}\lambda_1(\mathbf{u}) &= \lambda_{1,0} - \lambda_{1,1}u_1 + \lambda_{1,2}u_2, \\ \lambda_2(\mathbf{u}) &= \lambda_{2,0} + \lambda_{2,1}u_1 - \lambda_{2,2}u_2,\end{aligned}\tag{4.1}$$

for $\mathbf{u} \in \mathcal{U} = \{\mathbf{u} \mid \lambda_1(\mathbf{u}) \geq 0, \lambda_2(\mathbf{u}) \geq 0\}$, where $\lambda_{1,0}, \lambda_{2,0} > 0, \lambda_{1,1} > \lambda_{2,1} > 0, \lambda_{2,2} > \lambda_{1,2} > 0$.

Substitution effects can also be incorporated to our welfare maximization model of Section 3.1.2. The model remains identical to the one introduced there with the exception that the user utility U_i of class i is a random variable depending on the whole price vector \mathbf{u} . In particular, we will assume that it has a probability density function, denoted by $f_i(u_i \mid u_j, j = 1, \dots, M, j \neq i)$, conditional on the prices of all other classes. Potential calls decide to join the network if and only if the utility they extract exceeds the prevailing price.

Thus, the arrival rate of class i calls under price \mathbf{u} is

$$\lambda_i(\mathbf{u}) = \lambda_{i,0} \mathbf{P}[U_i \geq u_i \mid u_j, j = 1, \dots, M, j \neq i],$$

where $\lambda_{i,0}$ is the peak class i demand (corresponding to zero prices in the revenue maximization model). A class i call joining the system extracts an expected utility equal to $\mathbf{E}[U_i \mid U_i \geq u_i; u_j, j = 1, \dots, M, j \neq i]$, thus, social welfare for class i users is accumulated at a rate of

$$\lambda_i(\mathbf{u}) \mathbf{E}[U_i \mid U_i \geq u_i; u_j, j = 1, \dots, M, j \neq i].$$

Our objective remains to maximize the expected long-term average welfare rate, for which an expression can be written along the lines of (3.2).

Let us define the expected instantaneous rewards by $V_i(\mathbf{u}) \triangleq u_i$ and $V_i(\mathbf{u}) \triangleq \mathbf{E}[U_i \mid U_i \geq u_i; u_j, j = 1, \dots, M, j \neq i]$ for the case of revenue and welfare maximization, respectively. We assume that $\lambda(\mathbf{u})$ satisfies Assumption D in this case as well. Consequently, $\lambda_i(\mathbf{u})$ is non-decreasing in u_j . We will be making the following assumption for the expected rewards.

Assumption E

For all $i = 1, \dots, M$ and $\mathbf{u} \in \{\mathbf{u} \mid \lambda_i(\mathbf{u}) > 0\}$, $V_i(\mathbf{u})$ is a non-decreasing function of u_j for all $j \neq i$.

This assumption is trivially satisfied for the case of revenue maximization where $V_i(\mathbf{u}) = u_i$. For the case of welfare maximization it can be interpreted as follows. Each class i has a strong core constituency and can not be dominated by class j ($j \neq i$) customers who choose to use class i as substitute when u_j increases. These “true” class i customers perceive that they are extracting a higher utility when other services becomes relatively more expensive. Thus, $V_i(\mathbf{u})$ is non-decreasing in u_j for $j \neq i$. It also turns out $V_i(\mathbf{u})$ is non-decreasing in u_i . The next lemma establishes the result.

Lemma 4.1.1 *For all $i = 1, \dots, M$ and $\mathbf{u} \in \{\mathbf{u} \mid \lambda_i(\mathbf{u}) > 0\}$, $V_i(\mathbf{u})$ is a non-decreasing function of u_i .*

Proof: The result is trivially true for the revenue maximization case where $V_i(\mathbf{u}) = u_i$. For welfare maximization we have

$$V_i(\mathbf{u} = \mathbf{E}[U_i | U_i \geq u_i; u_j, j = 1, \dots, M, j \neq i]) = \frac{\int_{u_i}^{\infty} v f_i(v|u_j, \forall j \neq i) dv}{\int_{u_i}^{\infty} f_i(v|u_j, \forall j \neq i) dv}.$$

Taking the partial derivative we obtain

$$\frac{\partial V_i(\mathbf{u})}{\partial u_i} = \frac{f_i(u_i|u_j, \forall j \neq i) \int_{u_i}^{\infty} (v - u_i) f_i(v|u_j, \forall j \neq i) dv}{(\mathbf{P}[U_i \geq u_i | u_j, \forall j \neq i])^2},$$

which is clearly non-negative. ■

This lemma can be seen as expressing the fact that when u_i increases class i customers with relatively low utility for the service choose not to use it, thus, the ones that remain have higher utilities and drive $V_i(\mathbf{u})$ up.

4.2 An Upper Bound on Optimal Performance

Assume the demand function with substitution effects is invertible, we can express the prices as a function of arrival rates λ ; we will write $u_i = u_i(\lambda)$ for the class i price. Define also $F_i(\lambda) \triangleq \lambda_i u_i(\lambda)$ and $\tilde{F}_i(\lambda) \triangleq \lambda_i \mathbf{E}[U_i | U_i \geq u_i; u_j, j = 1, \dots, M, j \neq i]$ for the case of revenue and welfare maximization, respectively. Assume that the functions $F_i(\lambda)$ are concave functions of λ for all i . This is true, for example, for the case of linear demand functions (4.1). The following result is analogous to Theorem 3.4.1.

Theorem 4.2.1 *If F_i are concave functions of λ , an upper bound of the optimal revenue J^* is given by*

$$\begin{aligned} \max_{\lambda_i, n_i} \quad & \sum_i F_i(\lambda) = \sum_i \lambda_i u_i(\lambda) \\ \text{s.t.} \quad & \lambda_i = n_i \mu_i \\ & \sum_i n_i r_{ji} \leq C_j, \quad \forall j \end{aligned} \tag{4.2}$$

$$\lambda_i, n_i \geq 0, \quad \forall i.$$

The optimal solution of (4.2) is J_{ub} , then $J^* \leq J_{\text{ub}}$.

Proof: Consider an optimal dynamic pricing policy \mathbf{u}^* . According to the Assumption D-4, we assume that the price \mathbf{u}^* becomes large enough and the arrival rate $\lambda_i(\mathbf{u})$ is equal to zero, whenever the state is such that a class i call cannot be admitted (which means $\mathbf{R}(\mathbf{n}+\mathbf{e}_i) \notin \mathbf{C}$ or $\exists j, \sum_k n_k r_{jk} + r_{ji} > C_j$). We can view λ_i^* and n_i as random variables, and use $\mathbf{E}[\cdot]$ to indicate expectation with respect to the steady-state distribution under this particular policy. At any time, we have $\sum_i n_i r_{ji} \leq C_j, \forall j$, which implies that $\sum_i \mathbf{E}[n_i] r_{ji} \leq C_j, \forall j$. Furthermore, Little's law implies $\mathbf{E}[\lambda_i^*] = \mu_i \mathbf{E}[n_i]$. This shows that $\mathbf{E}[n_i], \mathbf{E}[\lambda_i^*], i = 1, \dots, M$, forms a feasible solution of the problem (4.2). Using the concavity of F_i and Jensen's inequality, we have

$$J_{\text{ub}} \geq \sum_i F_i(\mathbf{E}[\lambda_i^*]) \geq \sum_i \mathbf{E}[F_i(\lambda_i^*)] = J^*,$$

where the last equality used the optimality of the policy under consideration. ■

4.3 Asymptotic Optimality of Static Pricing

Consider the same limiting regime of “many small users” of Section 3.5. We scale both demand and capacity by a scaling factor $c \geq 1$, while all other quantities are held fixed. The demand function becomes

$$\boldsymbol{\lambda}^c(\mathbf{u}) = (\lambda_1^c(\mathbf{u}), \dots, \lambda_M^c(\mathbf{u})) = (c\lambda_1(\mathbf{u}), \dots, c\lambda_M(\mathbf{u})). \quad (4.3)$$

The capacity of link j is cC_j , the load of link j is $\rho_j^c(\mathbf{u}) = \sum_i \frac{r_{ji} \lambda_i^c(\mathbf{u})}{\mu_i c C_j} = \sum_i \frac{r_{ji} \lambda_i(\mathbf{u})}{\mu_i C_j} = \rho_j(\mathbf{u}), j = 1, \dots, L$. The normalized revenue or welfare maximization problem under a static

pricing policy \mathbf{u} can be formulated as:

$$\max_{\mathbf{u} \in \mathcal{U}^c} \frac{1}{c} \sum_{i=1}^M V_i(\mathbf{u}) \lambda_i^c(\mathbf{u}) \left(1 - \mathbf{P}_{\text{loss}}^{i,c}(\mathbf{u})\right) = \max_{\mathbf{u} \in \mathcal{U}} \sum_{i=1}^M V_i(\mathbf{u}) \lambda_i(\mathbf{u}) \left(1 - \mathbf{P}_{\text{loss}}^{i,c}(\mathbf{u})\right), \quad (4.4)$$

where $\mathcal{U}^c = \mathcal{U} = \{\mathbf{u} \mid u_i \geq 0, \lambda_i(\mathbf{u}) \geq 0, i = 1, \dots, M\}$ is the feasible set of \mathbf{u} .

We use following asymptotic results about loss probabilities in [Kel91]. As $c \rightarrow \infty$, under a certain static pricing policy, the loss probability of each class converges to

$$\mathbf{P}_{\text{loss}}^i(\mathbf{u}) = 1 - \prod_{j=1}^L (1 - B_j)^{r_{ji}}, i = 1, \dots, M, \quad (4.5)$$

where $B_j \in [0, 1)$, $j = 1, \dots, L$ satisfy following conditions

$$\hat{\rho}_j(\mathbf{u}) \triangleq \sum_{i=1}^M \frac{r_{ji}}{\mu_i c C_j} \lambda_i^c(\mathbf{u}) \prod_{l=1}^L (1 - B_l)^{r_{li}} = \sum_{i=1}^M \frac{r_{ji}}{\mu_i C_j} \lambda_i(\mathbf{u}) \prod_{l=1}^L (1 - B_l)^{r_{li}} \begin{cases} = 1 & \text{if } B_j > 0, \\ \leq 1 & \text{if } B_j = 0. \end{cases} \quad (4.6)$$

Following Kelly [Kel91], we will call $\hat{\rho}_j(\mathbf{u})$ the *reduced load* on link j . We can interpret these asymptotic results as follows. Calls are blocked independently at each link j in their route. In particular, class i demand is thinned by a factor of $(1 - B_j(\mathbf{u}))^{r_{ji}}$ at link j and $\prod_{j=1}^L (1 - B_j(\mathbf{u}))^{r_{ji}} = 1 - \mathbf{P}_{\text{loss}}^{i,\infty}(\mathbf{u})$ can be seen as the proportion of accepted class i calls. This results into a satisfied demand for class i equal to $\lambda_i^c(\mathbf{u}) \prod_{j=1}^L (1 - B_j(\mathbf{u}))^{r_{ji}}$. We will use Kelly's [Kel91] terminology and say that link j is *overloaded* if $B_j(\mathbf{u}) > 0$ (which implies $\hat{\rho}_j = 1$); if $B_j(\mathbf{u}) = 0$ we will say that it is *underloaded* ($\hat{\rho}_j < 1$) or *critically loaded* ($\hat{\rho}_j = 1$). We should note that although the conditions in (4.6) lead to unique values for the reduced loads $\hat{\rho}_j(\mathbf{u})$ and the loss probabilities $\mathbf{P}_{\text{loss}}^{i,\infty}(\mathbf{u})$, the parameters $B_j(\mathbf{u})$ might not have a unique value. In fact, the values of $B_j(\mathbf{u})$ are unique if the routing matrix \mathbf{R} has rank L ; otherwise, there exists a unique vector $(B_1(\mathbf{u}), \dots, B_L(\mathbf{u}))$ with maximal support, i.e., a vector that solves (4.6) and maximizes the dimensions of the set $\mathcal{B}(\mathbf{u}) \triangleq \{j \mid B_j(\mathbf{u}) > 0\}$. The following lemma states an observation that would be useful later on.

Lemma 4.3.1 *The offered loads $\rho_j(\mathbf{u})$ satisfy $\rho_j(\mathbf{u}) \leq 1$ for all links $j = 1, \dots, L$ if and only if $B_j(\mathbf{u}) = 0$ for all links $j = 1, \dots, L$.*

Proof: We will first argue that if $\rho_j(\mathbf{u}) \leq 1$, for all j , then $B_j(\mathbf{u}) = 0$ for all j . Otherwise, suppose there is a link j with $B_j(\mathbf{u}) > 0$. Due to (4.6) there is at least one i for which $r_{ji}\lambda_i(\mathbf{u}) > 0$. Moreover, (4.6) also implies that $\hat{\rho}_j(\mathbf{u}) < \rho_j(\mathbf{u})$ and $\hat{\rho}_j(\mathbf{u}) = 1$. This contradicts the initial assumption $\rho_j(\mathbf{u}) \leq 1$. For the converse, note that if all links are either underloaded or critically loaded, i.e., $B_j(\mathbf{u}) = 0$ for all j , then $\rho_j(\mathbf{u}) = \hat{\rho}_j(\mathbf{u}) \leq 1$, for all j . ■

Another interesting observation is that due to (4.5), $B_j(\mathbf{u}) = 0$, for all j , implies that $\mathbf{P}_{\text{loss}}^i(\mathbf{u}) = 0$ for all $i = 1, \dots, M$.

We next define normalized reward of class i with respect to link j , for all links j with $r_{ji} > 0$, as follows

$$\hat{V}_{i,j}(\mathbf{u}) \triangleq \frac{V_i(\mathbf{u})}{r_{ji}/\mu_i}, \quad (4.7)$$

thus,

$$V_i(\mathbf{u})\lambda_i(\mathbf{u}) = \hat{V}_{i,j}(\mathbf{u}) \cdot \frac{r_{ji}\lambda_i(\mathbf{u})}{\mu_i} = \hat{V}_{i,j}(\mathbf{u})\hat{\lambda}_{i,j}(\mathbf{u}),$$

where $\hat{\lambda}_{i,j}(\mathbf{u}) \triangleq \frac{r_{ji}\lambda_i(\mathbf{u})}{\mu_i}$ is the normalized demand of class i for link j . We can interpret $\hat{V}_{i,j}(\mathbf{u})$ as reward per volume on link j , where volume has the same interpretation as in Section 3.6, that is, resource utilization times the expected holding time. For a given static pricing policy \mathbf{u} , the normalized rewards at link j are fixed and define an ordering among classes traversing link j . In particular, for any classes i and k traversing link j (i.e., $r_{ji}, r_{ki} > 0$), we will say that i is *more valuable* than k if $\hat{V}_{i,j}(\mathbf{u}) > \hat{V}_{k,j}(\mathbf{u})$. If calls occupy the *same* resource amount at all links in their route (i.e., for all i , $r_{ji} = r_i$ for all $j \in \mathcal{R}_i$ and $r_{ji} = 0$ for all $j \notin \mathcal{R}_i$), then the priority ordering of classes is the same on all links and $\hat{V}_{i,j}(\mathbf{u}) = \frac{V_i(\mathbf{u})}{r_{ji}/\mu_i}$ define a unique priority ordering for the whole network. The following proposition is key in establishing our asymptotic optimality result.

Proposition 4.3.2 Consider either the case of revenue or welfare maximization and assume that for any class $i = 1, \dots, M$ and all links $j = 1, \dots, L$

$$r_{ji} = \begin{cases} r_i & j \in \mathcal{R}_i, \\ 0 & \text{otherwise,} \end{cases}$$

If $\mathbf{u}_{s,\infty}$ solves the limiting case of problem (4.4), i.e.,

$$\max_{\mathbf{u} \in \mathcal{U}} \lim_{c \rightarrow \infty} \frac{1}{c} \sum_{i=1}^M V_i(\mathbf{u}) \lambda_i^c(\mathbf{u}) \left(1 - \mathbf{P}_{\text{loss}}^{i,c}(\mathbf{u})\right) = \max_{\mathbf{u} \in \mathcal{U}} \sum_{i=1}^M V_i(\mathbf{u}) \lambda_i(\mathbf{u}) \prod_{j=1}^L (1 - B_j(\mathbf{u}))^{r_{ji}}, \quad (4.8)$$

then

$$\rho_j(\mathbf{u}_{s,\infty}) = \sum_i \frac{\lambda_i(\mathbf{u}_{s,\infty}) r_{ji}}{\mu_i C_j} = \sum_{i|j \in \mathcal{R}_i > 0} \frac{\lambda_i(\mathbf{u}_{s,\infty}) r_i}{\mu_i C_j} \leq 1, \quad j = 1, \dots, L.$$

Proof: The following discussion is about the case that $c \rightarrow \infty$. Let \mathbf{u} be a static pricing policy such that on some links, the offered load is greater than 1. The average reward is

$$\sum_i V_i(\mathbf{u}) \lambda_i(\mathbf{u}) \prod_l (1 - B_l(\mathbf{u}))^{r_{li}} = \sum_i \hat{V}_i(\mathbf{u}) \frac{r_i}{\mu_i} \lambda_i(\mathbf{u}) \prod_l (1 - B_l(\mathbf{u}))^{r_{li}}, \quad (4.9)$$

where $B_l(\mathbf{u}), l = 1, \dots, L$ satisfy (cf. (4.6))

$$\sum_i \frac{r_{ji}}{\mu_i} \lambda_i(\mathbf{u}) \prod_l (1 - B_l)^{r_{li}} = \sum_{i|j \in \mathcal{R}_i} \frac{r_i}{\mu_i} \lambda_i(\mathbf{u}) \prod_l (1 - B_l)^{r_{li}} \leq C_j, \quad j = 1, \dots, L. \quad (4.10)$$

Consider following linear programming (LP) problem with decision variables $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_M)$

$$\begin{aligned} \max_{\lambda_i} \quad & \sum_i \hat{V}_i(\mathbf{u}) \lambda_i & (4.11) \\ \text{s.t.} \quad & \sum_{i|j \in \mathcal{R}_i} \lambda_i = \sum_i a_{ji} \lambda_i \leq C_j, \quad j = 1, \dots, L, \\ & 0 \leq \lambda_i \leq \frac{r_i}{\mu_i} \lambda_i(\mathbf{u}), \quad i = 1, \dots, M, \end{aligned}$$

where $a_{ji} = 1$, if $j \in \mathcal{R}_i$, and zero otherwise. Let $\hat{\boldsymbol{\lambda}} = (\hat{\lambda}_1, \dots, \hat{\lambda}_M)$ denote an optimal

solution.

Consider next the network under static policy \mathbf{u} but introduce a random admission control mechanism. More specifically, class i calls are accepted with a probability equal to $\hat{\lambda}_i / \left(\frac{r_i}{\mu_i} \lambda_i(\mathbf{u}) \right)$. Thus, admitted class i calls arrive according to a Poisson process of rate $\hat{\lambda}_i \mu_i / r_i$, since their requests follow a Poisson process of rate $\lambda_i(\mathbf{u})$. We will call System S_A the original one (without admission control) with prices \mathbf{u} , and System S_B the new system (with admission control). Note that in System S_B we have

$$\rho_j^B(\mathbf{u}) = \sum_{i|j \in \mathcal{R}_i} \frac{r_i \hat{\lambda}_i \mu_i / r_i}{\mu_i C_i} = \sum_{i|j \in \mathcal{R}_i} \frac{\hat{\lambda}_i}{C_j} \leq 1, \quad \forall j,$$

due to the feasibility condition in (4.11). Thus, Lemma 4.3.1 implies that the blocking probabilities are equal to zero for all classes. Notice now that due to (4.10), $\lambda_i = \frac{r_i}{\mu_i} \lambda_i(\mathbf{u}) \prod_l (1 - B_l)^{r_l}$, $i = 1, \dots, M$, form a feasible solution to problem (4.11). Thus, the objective value at this feasible solution (given by (4.9)) cannot be more than the optimal. The optimal value of problem (4.11) is simply the average reward in System S_B , thus, it is not less the average reward in our original System S_A .

Consider next the special structure of the problem (4.11). The elements of the constraint matrix are either 1 or 0, and the coefficients in the objective function are the normalized rewards, which, as we have discussed before, form a priority ordering for the various classes. It can be seen, that an optimal solution of (4.11) can be constructed as follows. Make the $\hat{\lambda}_i$ corresponding to the highest priority class large as the capacity constraints allow, after this is set, make the $\hat{\lambda}_i$ corresponding to the next highest priority class as large as possible, and so on. Applying the randomized admission control discussed previously, the arrival rate of high priority classes is maximized and only classes with $\hat{\lambda}_i > 0$ gain admission to the network. Let us consider the implications on a link j with offered load $\rho_j(\mathbf{u}) > 1$ in System S_A . Say that classes i_1, \dots, i_K use this link, where $\hat{V}_{i_1}(\mathbf{u}) \geq \dots \geq \hat{V}_{i_K}(\mathbf{u})$. Since in System S_B the offered load on this link is less than or equal to one, some classes do not gain admission. In particular, let i_k be the least priority class that gains admission.

In System S_B , we accept as much as possible requests from classes i_1, \dots, i_{k-1} (the arrival rate of these classes may be less than their original rate if they are not a high priority class on another congested link), reject all requests from class i_{k+1}, \dots, i_K , and from class i_k , we only accept a portion of its calls such that the offered load on link j is equal to 1 in System S_B .

Let us now slightly increase the price of the lowest priority class i_K in both systems S_A and S_B . Solve the LP in (4.11) again to obtain a new $\hat{\lambda}$ corresponding to the new set of prices, and apply admission control as before to construct System S_B . As before, the average reward in System S_B is no less than the corresponding reward in System S_A . Since we increased the price of class i_K , the load of those accepted classes in System S_B could stay the same or increase slightly due to the substitution effects. Since the $\hat{V}_i(\mathbf{u})$'s do not decrease (according to Assumption E and Lemma 4.1.1) the expected total reward will not be less than before. We keep increasing u_{i_K} , unless we reach the point where $\rho_j(\mathbf{u}) = 1$ at which we stop. During this process the ordering of classes (according to their normalized rewards) might change; to avoid further complicating the notation we will use the indices i_1, \dots, i_K for the ordered set of classes using link j . As the normalized reward $\hat{V}_{i_K}(\mathbf{u})$ increases it may reach the normalized reward $\hat{V}_{i_{K-1}}(\mathbf{u})$ of the next least priority class; in this case we start increasing the prices of both these two classes. It is also possible that during this procedure of price increases the load of the accepted classes increases to the point where the “threshold” class i_k is completely pushed out of the network in System S_B ; in this case i_{k-1} becomes the “threshold” class. These price increases may also influence the load on other links. We repeat this process on other overloaded links; we may visit an overloaded link several times due to substitution effects. We stop when we arrive at a price vector $\hat{\mathbf{u}}$ at which $\rho_j(\hat{\mathbf{u}}) \leq 1$ for all link j . This is guaranteed by Assumption D-3 and D-4. Throughout this process the average reward in System S_B did not decrease and the average reward in the corresponding System S_A remained less than or equal to the average reward in System S_B .

Consider now the LP in (4.11) at $\hat{\mathbf{u}}$. Since $\rho_j(\hat{\mathbf{u}}) \leq 1$ for all j , the capacity constraints

in (4.11) are immediately satisfied and at the optimal solution the decision variables $\hat{\lambda}_i$ can become equal to their upper bounds $\frac{r_i}{\mu_i} \lambda_i(\mathbf{u})$ for all i . Thus, in System S_B all calls are accepted with probability one and System S_B becomes identical to System S_A .

To summarize, we started from an arbitrary price vector \mathbf{u} under which some links have offered loads greater than one and constructed a price vector $\bar{\mathbf{u}}$ with higher average reward and offered loads satisfying $\rho_j(\bar{\mathbf{u}}) \leq 1$ on all links j . We conclude that the optimal limiting static policy $\mathbf{u}_{s,\infty}$ must satisfy $\rho_j(\mathbf{u}_{s,\infty}) \leq 1$ for all links j . ■

Due to Lemma 4.3.1, the result of Proposition 4.3.2 implies that at the optimal static prices in the limiting regime, $\mathbf{u}_{s,\infty}$, all links in the network are underloaded or critically loaded, and all classes experience zero blocking probabilities. The following theorem is an immediate consequence of these observations and Proposition 4.3.2

Theorem 4.3.3 *Consider either the case of revenue or welfare maximization and assume for any class $i = 1, \dots, M$ and all links $j = 1, \dots, L$*

$$r_{ji} = \begin{cases} r_i & j \in \mathcal{R}_i, \\ 0 & \text{otherwise,} \end{cases}$$

The optimal static policy in the limiting regime, $\mathbf{u}_{s,\infty}$, solves, the following optimization problem:

$$\begin{aligned} \max_{\mathbf{u}} \quad & \sum_{i=1}^M V_i(\mathbf{u}) \lambda_i(\mathbf{u}) \\ \text{s.t.} \quad & \sum_{i=1}^M \frac{\lambda_i(\mathbf{u}) r_{ji}}{\mu_i} \leq C_j, \quad j = 1, \dots, L. \end{aligned} \tag{4.12}$$

The optimization problem in (4.12) is in fact the same as the upper bound problem in (4.2), with the exception that decision variables are the prices instead of the arrival rates. Thus, in the limiting regime ($c \rightarrow \infty$), the optimal static policy achieves upper bound and

it is *asymptotically optimal*.

4.4 Structure of the Asymptotically Optimal Static Policy

As in Section 3.6, where we considered the original model of Section 3.1, we will next characterize the structure of asymptotically optimal prices for the modified model of Section 4.1 that incorporates demand substitution effects.

Let us first focus on the revenue maximization case. Consider the problem in (4.12) and rewrite it as a minimization problem. Let $\mathbf{q} = (q_1, \dots, q_L) \geq 0$ be the Lagrange multiplier vector, where q_j is associated with the capacity constraints on link j . The Lagrangean function becomes

$$\mathcal{L}(\mathbf{u}, \mathbf{q}) = - \sum_{i=1}^M \lambda_i(\mathbf{u}) u_i + \sum_{j=1}^L q_j \left(\sum_{i=1}^M \frac{\lambda_i(\mathbf{u}) r_{ji}}{\mu_i} - C_j \right). \quad (4.13)$$

Assuming an interior solution, \mathbf{u} should satisfy

$$\nabla \lambda(\mathbf{u}) \mathbf{u} = -\lambda(\mathbf{u}) + \sum_{j=1}^L q_j \sum_{i=1}^M \frac{r_{ji}}{\mu_i} \nabla \lambda_i(\mathbf{u}), \quad (4.14)$$

where $\nabla \lambda(\mathbf{u})$ is the gradient of the vector function $\lambda(\mathbf{u})$, i.e., an $M \times M$ matrix with (i, j) element equal to $\frac{\partial \lambda_j(\mathbf{u})}{\partial u_i}$.

Welfare maximization can be treated similarly. One can write down the optimality conditions for the problem in (4.12) and solve them analytically for relatively simple forms of the utility density functions $f_i(\cdot)$. The structure of those conditions is rather complex, so one would have to resort to numerical solution methods for the general case.

The discussion on solving large scale problems in Section 3.7 readily extends to the model with demand substitution effects. We can use the structure in this section and simulation-based optimization approach to find a good static policy for large network with demand substitution effects.

4.5 Numerical Results

To provide an example, we consider the network in Figure 3.2, but now incorporate demand substitution effects. The network provides 12 classes of services (See Table 4.1); Class 1 and 7, 2 and 8, can be used as substitutes of each other.

Table 4.2 compares the upper bound J_{ub} (cf. Theorem 4.2.1) with the two policies proposed in Section 3.7. The corresponding prices are given in Table 4.3. We conclude that the optimized version (via the simulation-based optimization approach) of our asymptotically optimal static pricing policy is reasonably close to the optimal.

Class i	Nodes (Links)	r_i	$\lambda_i(\mathbf{u})$	μ_i
1	1, 0, 2 (1, 2)	1	$150 - 75u_1 + u_7$	5
2	1, 0, 3 (1, 3)	1	$150 - 60u_2 + u_8$	5
3	1, 0, 4 (1, 4)	1	$250 - 125u_3$	5
4	2, 0, 3 (2, 3)	1	$140 - 40u_4$	5
5	2, 0, 4 (2, 4)	1	$300 - 150u_5$	5
6	3, 0, 4 (3, 4)	1	$150 - 35u_6$	5
7	1, 0, 2 (1, 2)	5	$6 - u_7 + 0.1u_1$	1
8	1, 0, 3 (1, 3)	5	$7 - 1.2u_8 + 0.1u_2$	1
9	1, 0, 4 (1, 4)	5	$6.6 - 0.8u_9$	1
10	2, 0, 3 (2, 3)	5	$6 - 0.6u_{10}$	1
11	2, 0, 4 (2, 4)	5	$6 - 0.6u_{11}$	1
12	3, 0, 4 (3, 4)	5	$6 - 0.5u_{12}$	1

Table 4.1: The services with demand substitution effects provided by the network in Figure 3.2.

J_{ub}	$J(\mathbf{u}_{ub}^*)$	J_{sim}	$\frac{J_{ub} - J_{sim}}{J_{ub}} \times 100\%$
814.84	765.61	783.85	3.8%

Table 4.2: Comparing the various policies for the network of Section 4.5. $J_{\mathbf{u}_{ub}^*}$ and J_{sim} denote the performance of the policy obtained from the upper bound problem and the simulation-based optimization approach, respectively.

Policy	u_1	u_2	u_3	u_4	u_5	u_6	u_7	u_8	u_9	u_{10}	u_{11}	u_{12}
\mathbf{u}_{ub}^*	1.05	1.31	1.02	1.75	1.00	2.14	4.12	4.06	4.69	5.00	5.00	6.00
\mathbf{u}_{sim}^*	1.09	1.26	1.06	1.79	1.04	2.18	4.92	4.90	5.60	5.93	5.93	6.93

Table 4.3: The prices for the network of Section 4.5 under the policy obtained from the upper bound problem (\mathbf{u}_{ub}^*) and the simulation-based optimization approach (\mathbf{u}_{sim}^*).

Chapter 5

Inventory Control for Single-Stage Make-to-Stock System

In this chapter, we switch gears and consider the second instance of a stochastic network problem we address in this thesis, that is, inventory control in supply chains subject to QoS requirements. We start the investigation of such problems with the simpler case of single-stage systems. The remainder of this chapter is organized as follows. In Section 5.1 we provide the model of the single-stage, single-class, production-inventory system, introduce the base-stock policy and formulate the problems we will consider. In Section 5.2, we analyze the base-stock policy and obtain the approximations on the stockout probability and average inventory cost using large deviations techniques. In Section 5.3, we present the numerical results.

5.1 The Model

We consider the make-to-stock manufacturing system depicted in Figure 5.1. Demand is met from the *finished goods inventory (FGI)* and unsatisfied demand is backordered. We assume a discrete-time model, where time is slotted and the state of the system is examined at the beginning of each time slot n (periodic review policy), with n being in the set of integers \mathbb{Z} . Let D_n denote the demand arriving during time slot n , and B_n denote the amount of goods that the production facility can produce (capacity) during the same time slot. Let also I_n denote the available inventory at the beginning of time slot n (without




Figure 5.1: The model of a make-to-stock system.

taking into account D_n and the amount of goods that the machine produces during time slot n). We allow I_n to take values in the set of real numbers \mathbb{R} ; when nonnegative it is equal to the amount of available inventory, and when negative it is equal to the amount of backordered demand.

We assume that the demand $\{D_n; n \in \mathbb{Z}\}$ and the service $\{B_n; n \in \mathbb{Z}\}$ processes are arbitrary, stationary, and mutually independent stochastic processes, that satisfy certain mild technical conditions (a large deviations principle, see Assumption B for details). These assumptions are satisfied by a fairly large class of stochastic processes, which includes renewal processes, Markov-modulated processes (where D_n for example is a function of an underlying Markov process), and general stationary processes with mild mixing conditions. For stability purposes, we further assume that

$$\mathbf{E}[B_1] > \mathbf{E}[D_1], \quad (5.1)$$

which by stationarity carries over to all time slots n .

We will implement a *base-stock* policy which maintains a safety stock or hedging point of w . More specifically, the system produces when the inventory is below w and idles otherwise. According to this policy the inventory evolves as follows:

$$I_{n+1} = \min\{I_n - D_n + B_n, w\}. \quad (5.2)$$

We quantify customer dissatisfaction by the probability, $\mathbf{P}[I_n \leq 0]$, of not being able to meet

incoming demand immediately (*stockout probability*). Let ϵ be a desirable upper bound on the QoS level. We are interested in selecting a hedging point w that solves the following optimization problem:

$$\begin{aligned} & \text{minimize} \quad \text{Cost} = h\mathbf{E}[I_n^+] \\ & \text{subject to} \quad \mathbf{P}[I_n \leq 0] < \epsilon, \end{aligned} \tag{5.3}$$

where h is a given scalar and I_n^+ denotes $\max\{I_n, 0\}$. We will be referring to the constraint

$$\mathbf{P}[I_n \leq 0] < \epsilon$$

as the *service-level* constraint. To achieve our goal we need to compute $\mathbf{P}[I_n \leq 0]$. An exact analytic expression is impossible to obtain, especially in view of the complicated, autocorrelated, models for the demand and production processes. To that end, either simulation or asymptotic techniques can be applied (see [PLCZ01]). We will resort to an asymptotic large deviations analysis.

5.2 Large Deviations Analysis

Define the shortfall L_n as the gap between the current inventory and the hedging point, i.e.,

$$L_n \triangleq w - I_n.$$

In terms of L_n , Equation (5.2) can be written as

$$L_{n+1} = \max\{L_n + D_n - B_n, 0\}, \tag{5.4}$$

and we have the following equality

$$\mathbf{P}[I_n \leq 0] = \mathbf{P}[L_n \geq w].$$

We can interpret L_n as the queue length of a discrete-time $G/G/1$ queue with D_n arrivals and B_n server's capacity during time slot n . On a notational remark, in the sequel we will be dropping the reference to the time slot (subscript n) when referring to steady-state quantities. For example, we will be denoting by L the steady-state queue length at an arbitrary time slot.

5.2.1 Stockout Probability

The following proposition characterizes the tail of the shortfall process (see [BPT98b] for a proof in the continuous time domain).

Proposition 5.2.1 (Single-Stage) *Under Assumption B, the steady-state queue length process L satisfies*

$$\lim_{w \rightarrow \infty} \frac{1}{w} \log \mathbf{P}[L \geq w] = -\theta^*, \quad (5.5)$$

where $\theta^* > 0$ is the largest root of the equation

$$\Lambda_D(\theta) + \Lambda_B(-\theta) = 0. \quad (5.6)$$

Intuitively, for large enough w we have

$$\mathbf{P}[I \leq 0] = \mathbf{P}[L \geq w] \sim e^{-w\theta^*}, \quad (5.7)$$

therefore, the minimum w that guarantees the stockout probability to be below ϵ is

$$w = -\frac{\log \epsilon}{\theta^*}.$$

Notice that $\Lambda_D(\theta) + \Lambda_B(-\theta)$ is zero at the origin and has negative derivative at the same point (due to (5.1)). Figure 5.2 depicts the root of the equation $\Lambda_D^+(\theta) + \Lambda_B(-\theta) = 0$. In the extreme case that $\Lambda_D(\theta) + \Lambda_B(-\theta) < 0$ for all $\theta > 0$ we will say that $\theta^* = \infty$. In this case, no stockouts occur and a safety stock of zero should be maintained (*Just in Time*

(JIT) policy).




Figure 5.2: The largest root of $\Lambda_D^+(\theta) + \Lambda_B(-\theta) = 0$.

To improve the accuracy of the asymptotics, especially for fairly large ϵ 's, we can introduce a prefactor $f(w)$ and consider the approximation

$$\mathbf{P}[I \leq 0] = \mathbf{P}[L \geq w] \approx f(w) e^{-w\theta^*},$$

where $f(w)$ is in general any function that satisfies $\frac{\log f(w)}{w} \rightarrow 0$ as $w \rightarrow \infty$ (cf. Proposition 5.2.1). Note that this is true for any polynomial function of w . For renewal demand and production processes $f(w)$ is a constant ([Asm87]) and it is equal to 1 under $M/M/1$ assumptions. We will use a constant for the more general case as well. In particular, we will set $f(w) = c$ which yields the following approximation

$$\mathbf{P}[I \leq 0] = \mathbf{P}[L \geq w] \approx c e^{-w\theta^*}. \quad (5.8)$$

Thus, the hedging point satisfies

$$w = -\frac{\log(\epsilon/c)}{\theta^*}. \quad (5.9)$$

The coefficient c can be estimated by assuming that Equation 5.8 is the exact distribution of the queue length process. By matching expectations, we obtain

$$\mathbf{E}[L] = \int_0^\infty \mathbf{P}[L \geq w] dw = \int_0^\infty c e^{-w\theta^*} dw = \frac{c}{\theta^*},$$

therefore,

$$c = \theta^* \mathbf{E}[L]. \quad (5.10)$$

Note that $\mathbf{E}[L]$ is independent of w , and can be obtained either by approximations of the expected queue length in a $G/G/1$ queue (as in [BP01]) or by simulation.

5.2.2 Inventory Cost

We finally consider the inventory cost. Let $C(w)$ be the expected inventory cost, when we fix the hedging point to w , i.e.,

$$C(w) = h\mathbf{E}[I^+], \quad (5.11)$$

where h is a given constant. Using the equivalence to the make-to-order system we obtain

$$\begin{aligned} C(w) &= h\mathbf{E}[(w - L)^+] \\ &= h\mathbf{E}[\max(w - L, 0)] \\ &= h(w - \mathbf{E}[L] + \mathbf{E}[\max(L - w, 0)]). \end{aligned} \quad (5.12)$$

Using the asymptotic in (5.9) we have

$$\begin{aligned} \mathbf{E}[\max(L - w, 0)] &= \int_0^\infty \mathbf{P}[\max(L - w, 0) > x] dx \\ &= \int_0^\infty \mathbf{P}[L - w > x] dx \\ &\approx ce^{-w\theta^*} \int_0^\infty e^{-x\theta^*} dx \\ &= c \frac{e^{-w\theta^*}}{\theta^*}. \end{aligned} \quad (5.13)$$

Using (5.9) we obtain the following approximation for the expected inventory cost

$$C(w) \approx h(w - \mathbf{E}[L] + \mathbf{E}[L]e^{-w\theta^*}). \quad (5.14)$$

Summarizing the results in this chapter, large deviations techniques yield (asymptot-

ically exact) approximations for the stockout probability in (5.9) and expected inventory cost in (5.14). These expressions can be used in the optimization problem (5.3) to obtain the proper hedging point w .

5.3 Numerical Results

In this section we provide numerical results to demonstrate the accuracy of the large deviations asymptotics developed in Section 5.2. We consider only the stockout probability as a performance metric.

The demand and production processes are discrete-time Markov modulated processes (see Figure 5.3).

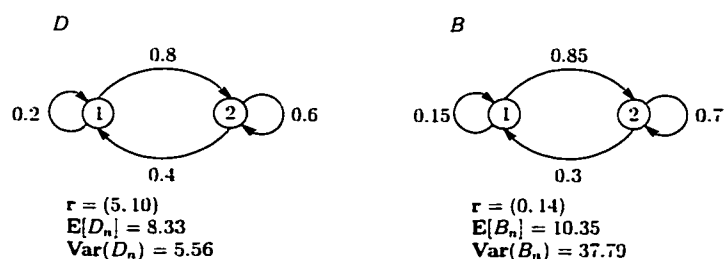


Figure 5.3: The models of the demand and production processes.

Both D and B are modulated by a two-state Markov chain. By \mathbf{r} we denote the vector of demand or production amounts at each state of the corresponding Markov chain. That is, D_n can be either 5 or 10 and B_n can either be 0 (machine down) or 14 (machine working). The load of the system is nearly 0.8. The transition probability matrices of Markov chains that modulate $\{D_n\}$ and $\{B_n\}$ are listed as follows:

$$\mathbf{P}_D = \begin{bmatrix} 0.2 & 0.8 \\ 0.4 & 0.6 \end{bmatrix}; \quad \mathbf{P}_B = \begin{bmatrix} 0.15 & 0.85 \\ 0.3 & 0.7 \end{bmatrix}.$$

We have the following LD result for this problem:

$$\theta^* = 0.120,$$

The expected shortfall from simulation is $\mathbf{E}[L_n] = 6.402$. Therefore, the analytic expression for the optimal hedging point (5.9) becomes

$$w^* = -\frac{\log \frac{\epsilon}{0.120 \cdot 6.402}}{0.120}.$$

Table 5.1 compares the results of the LD expression with the brute-force simulation. The refined LD approximation results are very accurate and the results are very close for most ϵ 's, even for large ϵ .

ϵ	LD Results	Simulation Results	
	w	w	$\mathbf{P}[X_n \leq 0]$
0.3	7.84	8	0.326
0.2	11.22	11	0.194
1.0×10^{-1}	16.99	17	0.996×10^{-1}
5.0×10^{-2}	22.77	23	4.852×10^{-2}
1.0×10^{-2}	36.18	36	1.040×10^{-2}
5.0×10^{-3}	41.96	42	5.061×10^{-3}
1.0×10^{-3}	55.37	55	1.072×10^{-3}
5.0×10^{-4}	61.14	61	5.220×10^{-4}
1.0×10^{-4}	74.56	75	0.964×10^{-4}
5.0×10^{-5}	80.33	80	5.283×10^{-5}
1.0×10^{-5}	93.74	94	0.983×10^{-5}
5.0×10^{-6}	99.52	100	4.885×10^{-6}
1.0×10^{-6}	112.93	113	1.059×10^{-6}

Table 5.1: Comparing the analytical the LD expression for the stockout probability with the simulated values. The first column lists the required stockout probabilities, ranging from 0.3 to 10^{-6} . The second column provides the hedging points obtained by the LD expression. The next two columns provide the hedging points used in the simulation and the stockout probability obtained.

Chapter 6

Inventory Control in Supply Chain Management: The Local Inventory Case

In this chapter, we consider multi-stage supply chain operating under a base-stock policy. Each stage has information about its local inventory only and we want to satisfy the service-level constraint on the finished goods of the supply chain. We propose a decomposition approach based on large deviations approximations and the results for single-stage system. The remainder of this chapter is organized as follows. In Section 6.1, we provide the model details of the multi-stage supply chain. A decomposition approach based on large deviations analysis is developed in Section 6.2.

6.1 The Model




Figure 6.1: The model of the supply chain.

Figure 6.1 depicts the supply chain model we consider in this and the following chapter. This system produces a single product class and consists of M production facilities in tandem. We will be referring to these facilities as stages of the supply chain and say that

production consists of M stages. External demand is met from the finished goods inventory maintained in front of the stage 1 production facility, and is backordered if inventory is not available. Every production facility is fed by its upstream facility; in particular, to produce one unit facility i , $i = 1, \dots, M - 1$, requires one unit of the product of facility $i - 1$. We assume that facility M is fed with an infinite supply of raw material, which implies that no material requirement constraints are in effect there. In front of every facility i , $i = 2, \dots, M$, there is an inventory buffer which holds the final product of that facility and from which facility $i - 1$ draws material for its production. We assume a periodic review policy where time is divided into time slots of equal duration. For all $i = 1, \dots, M$ and n we let B_n^i denote the amount that the facility at stage i can produce during time slot n (production capacity). We also let D_n^1 denote the amount of external orders arriving at stage 1 during time slot n . Finally, we let I_n^i , $i = 1, \dots, M$, denote the inventory in front of stage i at the beginning of time slot n . In intermediate stages $i = 2, \dots, M$, the inventory I_n^i is constrained to be nonnegative. In contrast, we allow the inventory at stage 1, I_n^1 , to take negative values to denote backordering; when I_n^1 is negative $-I_n^1$ is equal to the amount of backordered demand.

The system evolves as follows. At the beginning of time slot $n + 1$, the inventory at stage 1 is given by

$$I_{n+1}^1 = I_n^1 - D_n^1 + P_n^1,$$

where P_n^1 denotes the amount of products produced during time slot n by the facility at stage 1, which is determined by the production policy we select and confined by the production capacity B_n^1 and the available upstream inventory I_n^2 . The quantity P_n^1 can be also viewed as the demand for stage 2, which operates in a similar manner and generates demand for stage 3. Thus, the whole supply chain is driven by the external demand.

The demand process $\{D_n^1; n \in \mathbb{Z}\}$ and the production processes $\{B_n^i; n \in \mathbb{Z}\}$, $i = 1, \dots, M$, are arbitrary stationary stochastic processes that satisfy certain mild technical conditions (some form of a sample path large deviations principle). These conditions

are satisfied by renewal processes, Markov-modulated processes, and in general stationary processes with mild mixing conditions (for details see Bertsimas, Paschalidis, and Tsitsiklis [BPT98a, BPT98b, BPT99]). For stability purposes we assume that

$$\mathbf{E}[D_n^1] < \min_{i=1,\dots,M} \mathbf{E}[B_n^i], \quad (6.1)$$

which by stationarity carries over to all time slots n . Stability can be shown under both base-stock policies by using techniques from Baccelli and Liu [BL92]. For the case of an echelon base stock policy a stability proof is given in Glasserman and Tayur [GT94].

Our objective is to find a production policy within a selected class of inventory policies that minimizes expected inventory costs and guarantees that the steady-state stockout probability $\mathbf{P}[I_n^1 \leq 0]$, at some arbitrary time slot n , does not exceed a desirable small value ϵ . In this chapter, we will propose a policy for the case that each stage i has knowledge of its local inventory I_n^i only. In particular, every stage i sets a *hedging point* or *safety stock* w_i for its local inventory I_n^i and implements the production policy: produce if I_n^i falls below w_i , and idle otherwise. In the simpler single-stage ($M = 1$) this policy has been analyzed in Chapter 5 and an appropriate hedging point has been selected to maintain $\mathbf{P}[I_n^1 \leq 0] \leq \epsilon$. In a multi-stage system, however, there is strong coupling between stages since upstream inventory can constrain downstream production, which makes exact analysis particularly hard. To bypass this problem, we will use a *decomposition approach*. More specifically, we will focus in a regime where coupling between stages becomes weaker. For every stage this is the case if the safety stock in the upstream buffer is very large, implying that downstream production is rarely constrained by upstream inventory availability. In effect, each stage can be viewed as an independent single-stage system, and the results in Chapter 5 can be applied. To that end, though, we need to characterize the demand for every stage i by “propagating” the external demand through the downstream stages $1, 2, \dots, i - 1$.

We need to obtain the stockout probability in order to be able to maintain the service level constraints. Again, an exact expression is intractable, especially in view of the rather

complicated (autocorrelated) models for the demand and production processes. We will employ *large deviations* theory. In the regime of small stockout probability (or equivalently, large safety stocks) stockouts are rare events and are amenable to large deviations analysis. We will provide numerical results to demonstrate that the large deviations asymptotics are *accurate* when compared to simulations, even for fairly large stockout probabilities.

6.2 The Decomposition Approach

We propose a base-stock policy that maintains a safety stock equal to w_i for the (local) inventory of every stage i , $i = 1, \dots, M$. In particular, stage i produces until the local inventory I_n^i reaches the hedging point w_i and idles if $I_n^i \geq w_i$. The amount produced by stage i constitutes demand for the upstream stage $i + 1$, for $i = 1, \dots, M - 1$; we will denote it by D_n^{i+1} . Note that the demand for stage i , D_n^i , $i = 2, \dots, M$, is constrained by the downstream capacity B_n^{i-1} and the available inventory I_n^i .

The dynamics for the supply chain are

$$I_{n+1}^i = \min\{I_n^i - D_n^i + B_n^i, I_n^i - D_n^i + I_n^{i+1}, w_i\}, \quad i = 1, \dots, M - 1, \quad (6.2)$$

$$I_{n+1}^M = \min\{I_n^M - D_n^M + B_n^M, w_M\}. \quad (6.3)$$

The demand for stage i (or, equivalently, production of stage $i - 1$) is given by

$$D_n^i = I_{n+1}^{i-1} - I_n^{i-1} + D_n^{i-1}, \quad i = 2, \dots, M.$$

As in the single stage case, we define the inventory shortfall for stage i as follows:

$$L_n^i \triangleq w_i - I_n^i, \quad i = 1, \dots, M,$$

and the dynamics of the supply chain can be written as

$$L_{n+1}^i = \max\{L_n^i + D_n^i - B_n^i, L_n^i + D_n^i + L_n^{i+1} - w_{i+1}, 0\}, \quad i = 1, \dots, M - 1, \quad (6.4)$$

$$L_{n+1}^M = \max\{L_n^M + D_n^M - B_n^M, 0\}. \quad (6.5)$$

The demand for stage i can now be expressed as

$$D_n^i = L_n^{i-1} - L_{n+1}^{i-1} + D_n^{i-1}, \quad i = 2, \dots, M. \quad (6.6)$$

The major difficulty for analyzing this model and characterizing the stockout probabilities is that the production is constrained not only by its own capacity, but also by the upstream inventory. To bypass this difficulty we will *decouple* the various stages by ignoring the upstream inventory constraint on the downstream production. We can intuitively argue that this decomposition is in fact accurate when the inventory level of the upstream stage is high enough; then the influence of the upstream inventory constraint will be insignificant when compared to the capacity constraint. More specifically, the proposed decomposition amounts to assuming that the system operates according to a policy which satisfies

$$I_n^{i+1} \geq B_n^i, \quad i = 1, \dots, M - 1,$$

almost surely for all time slots n . As a result, the dynamics of the supply chain can be simplified as follows:

$$I_{n+1}^i = \min\{I_n^i - D_n^i + B_n^i, w_i\}, \quad i = 1, \dots, M, \quad (6.7)$$

or

$$L_{n+1}^i = \max\{L_n^i + D_n^i - B_n^i, 0\}, \quad i = 1, \dots, M. \quad (6.8)$$

That is, each stage behaves exactly as a single-stage system.

Next note that the dynamics in (6.8) are exactly the dynamics of M decoupled make-to-order $G/G/1$ queues. In particular, as in the single-stage problem discussed above, L_n^i can be interpreted as the queue length in a discrete-time $G/G/1$ queue with arrival process $\{D_n^i; n \in \mathbb{Z}\}$ and service process $\{B_n^i; n \in \mathbb{Z}\}$ (see Figure 6.2). Hence, Proposition 5.2.1

holds. To apply it, however, we need the large deviations rate functions of the processes $\{D_n^i; n \in \mathbb{Z}\}$. For $i = 1$, $\{D_n^1; n \in \mathbb{Z}\}$ is the external demand process, whose large deviations rate function is assumed known. For the remaining stages $i = 2, \dots, M$, recall that D_n^i is the demand for stage i generated by stage $i - 1$. In the equivalent make-to-order version of the system D_n^i can be interpreted as the number of departures from the stage $i - 1$ queue during time slot n . To see that consider the queue corresponding to stage $i - 1$ which has queue length equal to L_n^{i-1} at time slot n . Equation (6.6) simply states that the queue length at slot n (L_n^{i-1}) plus the number of arrivals at slot n (D_n^{i-1}) is equal to the queue length at slot $n + 1$ (L_{n+1}^{i-1}) plus the number of departures during slot n (D_n^i).




Figure 6.2: The equivalent $G/G/1$ queue of stage i , $i = 1, \dots, M$, in a decoupled multi-stage supply chain.

The following theorem characterizes the large deviations behaviour of the departure process $\{D_n^i; n \in \mathbb{Z}\}$, for all $i = 2, \dots, M$. This theorem is a corollary of a result in Bertsimas, Paschalidis, and Tsitsiklis [BPT98b] which characterizes the departure process of a $G/GI/1$ queue using a continuous-time model ¹.

Theorem 6.2.1 (Departure Process) *The partial sum of the departure process of the $G/G/1$ queue of stage $i - 1$ satisfies*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{P} \left[\sum_{j=1}^n D_j^i \geq na \right] = -\Lambda_{D^i}^{*+}(a), \quad i = 2, \dots, M, \quad (6.9)$$

where

$$\Lambda_{D^i}^{*+}(a) = \Lambda_{B^{i-1}}^{*+}(a) + \Lambda_{\Gamma^{i-1}}^{*+}(a), \quad (6.10)$$

¹In [BPT98b], the proof has been carried out for a renewal service process; we have noted though that a very similar proof applies to general, potentially autocorrelated, stationary service processes that satisfy a certain mild mixing condition.

and

$$\Lambda_{\Gamma^{i-1}}^{*+}(a) = \sup_{\{\theta \mid \Lambda_{D^{i-1}}^+(\theta) + \Lambda_{B^{i-1}}(-\theta) < 0\}} [\theta a - \Lambda_{D^{i-1}}^+(\theta)].$$

Proof: We prove the result by establishing a correspondence with a continuous-time $G/G/1$ queue and invoking the result in [BPT98b]. Consider first the queue corresponding to stage $i-1$ with queue length L_n^{i-1} at time slot n . Recall that the Lindley equation for this queue length is

$$L_{n+1}^{i-1} = \max\{L_n^{i-1} + D_n^{i-1} - B_n^{i-1}, 0\}. \quad (6.11)$$

Adding Equation (6.6) for stage i over all time slots $1, 2, \dots, n$, we obtain

$$\sum_{j=1}^n D_j^i = L_1^{i-1} - L_{n+1}^{i-1} + \sum_{j=1}^n D_j^{i-1}. \quad (6.12)$$

Consider next a continuous-time $G/G/1$ queue and let us denote by A_n the n th interarrival (interval between the arrivals of the $n-1$ st and n th customer), and by S_n the service time of the n th customer. The waiting time, W_n , of the n th customer satisfies the following Lindley equation

$$W_n = \max\{W_{n-1} + S_{n-1} - A_n, 0\}. \quad (6.13)$$

The interdeparture time, Z_n , of the n th customer (time interval between the departure times of customers $n-1$ and n) can be expressed as

$$Z_n = W_n - W_{n-1} + A_n + S_n - S_{n-1}. \quad (6.14)$$

Summing up the above over all customers $1, \dots, n$ we obtain

$$\sum_{j=1}^n Z_j = W_n - W_0 + S_n - S_0 + \sum_{j=1}^n A_j. \quad (6.15)$$

Compare now Eqs. (6.11) and (6.12) with (6.13) and (6.15) respectively. It can be seen that

by making the substitutions

$$W_j := L_{j+1}, \quad A_j := -D_j^{i-1}, \quad S_{j-1} := -B_j, \quad Z_j := -D_j^i,$$

for all j , $\sum_{j=1}^n D_j^i$ has the same large deviations behaviour as $-\sum_{j=1}^n Z_j$ ². Hence, invoking the result for the departure process of a continuous-time $G/G/1$ queue from [BPT98b] we obtain the large deviations characterization of the process $\{D_n^i, n \in \mathbb{Z}\}$ as it appears in the statement of the theorem. ■

Bertsimas, Paschalidis, and Tsitsiklis [BPT98b] also show that the departure process satisfies the exact same technical properties that the arrival process does (some form of a sample path large deviations principle). This is key because to apply Proposition 5.2.1 to every stage $i = 2, \dots, M$ in isolation we need the demand process D^i to satisfy a large deviations principle. Moreover, to derive the demand for the upstream stage $i+1$ we need to apply Theorem 6.2.1 which requires some form of a sample path large deviations principle.

We now have all the ingredients to analyze L_n^i for every stage i in isolation. The result is summarized in the following theorem.

Theorem 6.2.2 *For every stage $i = 1, \dots, M$ of the decoupled system, the steady-state queue length L^i satisfies*

$$\lim_{w_i \rightarrow \infty} \frac{1}{w_i} \log \mathbf{P}[L^i \geq w_i] = -\theta_{L,i}^*, \quad (6.16)$$

where $\theta_{L,i}^*$ is the largest root of the equation

$$\Lambda_{D^i}^+(\theta) + \Lambda_{B^i}(-\theta) = 0, \quad (6.17)$$

and $\Lambda_{D^i}^+(\theta)$, for $i = 2, \dots, M$, is the convex dual of $\Lambda_{D^i}^{*+}(a)$:

$$\Lambda_{D^i}^+(\theta) = \sup_a (\theta a - \Lambda_{D^i}^{*+}(a)), \quad (6.18)$$

²Note that for large values of n , S_n and S_0 in (6.15) are constants and do not affect the large deviations rate function.

where $\Lambda_{D_i}^{*+}(a)$ is as specified in Theorem 6.2.1.

Assume now that the stockout probability for stage 1 needs to be upper bounded by some ϵ_1 . We can then select the requirement for the stockout probability of stage i , ϵ_i , to be the same as, or an order of magnitude less than, the corresponding requirement, ϵ_{i-1} , for its downstream stage $i - 1$. Using the results of this section, we can obtain the hedging points:

$$w_i = -\frac{\log \epsilon_i}{\theta_{L,i}^*}, \quad i = 1, \dots, M.$$

We can improve the accuracy of the asymptotics, especially for fairly large ϵ 's by introducing a constant prefactor c_i for each decoupled stage i as for the single-stage systems. c_i can be estimated by (refer to Chapter 5)

$$c_i = \theta_{L,i}^* \mathbf{E}[L^i]. \quad (6.19)$$

Note that in the decoupled system $\mathbf{E}[L^i]$ is independent of w_i , and can be obtained either by approximations of the expected queue length in a $G/G/1$ queue (as in [BP01]) or by concurrent simulation (where one sample path of the stochastic processes involved is used to obtain $\mathbf{E}[L^i]$ for all stages i). Hence, the hedging point satisfies

$$w_i = -\frac{\log(\epsilon_i/c_i)}{\theta_{L,i}^*}, \quad i = 1, \dots, M. \quad (6.20)$$

Numerical results that help assess the accuracy of the large deviations asymptotics are given in Section 7.6.

Chapter 7

Inventory Control in Supply Chain Management: The Multi-echelon Approach

The policy obtained via the decomposition approach, although it maintains the service-level constraint at stage 1, might not necessarily be efficient in terms of expected inventory cost. Information of inventory availability in other stages might lead to lower such cost by giving the opportunity to trade-off inventory between different stages, i.e., lower the required safety stock in stages where inventory costs are high and compensate by increasing the safety stock in stages where costs are lower.

In this chapter we consider such a situation that each stage i has knowledge of the total downstream inventory $I_n^i + I_n^{i-1} + \dots + I_n^1$. We will implement another kind of base-stock policy. In particular, every stage i sets a hedging point or safety stock w_i for the total downstream inventory $I_n^i + I_n^{i-1} + \dots + I_n^1$ and implements the production policy: produce if $I_n^i + I_n^{i-1} + \dots + I_n^1$ falls below w_i , and idle otherwise. For the total downstream inventory from stage i we will be using the notation $X_n^i = I_n^i + I_n^{i-1} + \dots + I_n^1$, $i = 1, \dots, M$; we will be referring to this quantity as *echelon* inventory at stage i . In Section 7.1, we introduce the echelon base-stock policy. In Section 7.2, we analyze the supply chain under this policy using large deviations techniques. In Section 7.3 we propose an approach to refine the large deviations results. We formulate the optimization problem of minimizing expected inventory costs subject to given service-level constraints in Section 7.4. We discuss

extensions to the multiclass case and lost sales model in Section 7.5. Numerical results for both decomposition approach and multi-echelon approach are presented in Section 7.6.

7.1 The Multi-echelon Approach: A Global Information Case

In this section we consider the case where echelon inventory information is available at every stage $i = 1, \dots, M$. This will allow us to trade-off inventory between various stages in order to reduce expected inventory costs while maintaining the service level constraints.

We will be using the notation introduced in Section 6.1. As we noted above, X_n^i denotes the echelon inventory at time slot n and stage $i = 1, \dots, M$. We have

$$X_n^i = I_n^i + \dots + I_n^1 = I_n^i + X_n^{i-1}, \quad i = 1, \dots, M.$$

We implement an echelon base-stock production policy that maintains a hedging point or safety stock of w_i for the echelon inventory at stage i . More specifically, the facility at stage i produces until X_n^i reaches w_i and idles otherwise. Clearly, $w_1 \leq w_2 \leq \dots \leq w_M$. Figure 7.1 depicts the supply chain model and indicates the stages corresponding to the echelon safety stocks.




Figure 7.1: The supply chain model under the echelon base-stock policy.

As in Section 6.2, we define the shortfall of echelon i inventory as

$$Y_n^i \triangleq w_i - X_n^i, \quad (7.1)$$

which implies $\mathbf{P}[X_n^i \leq 0] = \mathbf{P}[Y_n^i \geq w_i]$. The dynamics of the echelon inventory are:

$$X_{n+1}^i = \min\{X_n^i - D_n^1 + B_n^i, w_i, X_n^i - D_n^1 + I_n^{i+1}\}, \quad i = 1, \dots, M-1, \quad (7.2)$$

$$X_{n+1}^M = \min\{X_n^M - D_n^1 + B_n^M, w_M\}. \quad (7.3)$$

In terms of the shortfalls the dynamics can be written as

$$Y_{n+1}^i = \max\{Y_n^i + D_n^1 - B_n^i, 0, Y_n^{i+1} + D_n^1 - (w_{i+1} - w_i)\}, \quad i = 1, \dots, M-1, \quad (7.4)$$

$$Y_{n+1}^M = \max\{Y_n^M + D_n^1 - B_n^M, 0\}. \quad (7.5)$$

7.2 Large Deviations Analysis of the Stockout Probability

Our first result, which is the main result of this section, is a large deviations result for the steady-state probability $\mathbf{P}[Y^1 \geq w_1]$, which is equal to the steady-state stockout probability $\mathbf{P}[X^1 \leq 0]$. Recall, that as in the previous section we drop the subscript n when referring to steady-state quantities.

On a notational remark, in the sequel we will be using \mathcal{O}_i to denote the i th-dimensional simplex, i.e.,

$$\mathcal{O}_i = \left\{ (\xi_1, \dots, \xi_i) \mid \xi_j \in [0, 1], \quad j = 1, \dots, i, \quad \sum_{j=1}^i \xi_j = 1 \right\}.$$

Theorem 7.2.1 *Assume the hedging points w_1, w_2, \dots, w_M in the multi-echelon system (cf. (7.4), (7.5)) satisfy*

$$w_i = \beta_{i-1} w_1, \quad i = 2, \dots, M,$$

where β_i are constants and $1 \leq \beta_1 \leq \dots \leq \beta_{M-1}$. The steady-state shortfall Y^1 of echelon

1 satisfies

$$\lim_{w_1 \rightarrow \infty} \frac{1}{w_1} \log \mathbf{P}[Y^1 \geq w_1] = -\theta_{G,1}^*, \quad (7.6)$$

where $\theta_{G,1}^*$ is determined by

$$\theta_{G,1}^* = \min \left[\inf_{a>0} \frac{1}{a} \inf_{x_0 - x_1 = a} (\Lambda_{D^1}^{*+}(x_0) + \Lambda_{B^1}^{*-}(x_1)), \right. \\ \left. \inf_{a>0} \frac{1}{a} \inf_{\substack{x_0 - \xi_1 x_1 - \xi_2 x_2 = a\beta_1 \\ (\xi_1, \xi_2) \in \mathcal{O}_2}} (\Lambda_{D^1}^{*+}(x_0) + \xi_1 \Lambda_{B^1}^{*-}(x_1) + \xi_2 \Lambda_{B^2}^{*-}(x_2)), \dots, \right. \\ \left. \inf_{a>0} \frac{1}{a} \inf_{\substack{x_0 - \xi_1 x_1 - \dots - \xi_M x_M = a\beta_{M-1} \\ (\xi_1, \dots, \xi_M) \in \mathcal{O}_M}} (\Lambda_{D^1}^{*+}(x_0) + \xi_1 \Lambda_{B^1}^{*-}(x_1) + \dots + \xi_M \Lambda_{B^M}^{*-}(x_M)) \right]. \quad (7.7)$$

To establish this result we will (i) obtain a sample path characterization of Y_n^1 , (ii) obtain a lower and an upper bound on $\mathbf{P}[Y^1 \geq w_1]$, and (iii) show that the upper and lower bounds match up to the first degree in the exponent.

We start by obtaining a sample path characterization of Y_n^1 . Suppose that at time 0, the echelon inventories are all equal to the corresponding safety stocks, i.e.,

$$Y_0^i = 0, \quad \forall i.$$

At time 1, the shortfall of the echelon 1 inventory is

$$\begin{aligned} Y_1^1 &= \max\{Y_0^1 + D_0^1 - B_0^1, 0, Y_0^2 + D_0^1 - (w_2 - w_1)\} \\ &= \max\{0, D_0^1 - B_0^1, D_0^1 - (w_2 - w_1)\} \\ &= \max\{0, D_0^1 - \min\{B_0^1, (w_2 - w_1)\}\} \\ &= \max\{0, D_0^1 - r_{1,1}^1\}, \end{aligned}$$

where $r_{1,1}^1 \triangleq \min\{B_0^1, (w_2 - w_1)\}$. Figure 7.2 depicts a certain graph in which $r_{1,1}^1$ can be interpreted as the length of the shortest path from point 1 (corresponding to stage 1) at level 0 to level 1. In general, we will use $r_{n,m}^i$ to denote the length of the shortest path among the paths with m hops for stage i , where n is the number of levels on the graph. For

the remaining stages i , $i = 2, \dots, M$, we have a similar characterization, that is,

$$Y_1^i = \max\{0, D_0^1 - r_{1,1}^i\},$$

where $r_{1,1}^i = \min\{B_0^i, (w_{i+1} - w_i)\}$, $i = 2, \dots, M - 1$, and $r_{1,1}^M = B_0^M$. In accordance with the notation we just introduced, note that in the graph of Figure 7.2 $r_{1,1}^i$, $i = 1, \dots, M$, denotes the length of the shortest path from point i at level 0 to level 1.




Figure 7.2: The paths for each stage at time slot 1 (one-level graph).

At time $n = 2$,

$$\begin{aligned} Y_2^1 &= \max\{0, Y_1^1 + D_1^1 - B_1^1, Y_1^2 + D_1^1 - (w_2 - w_1)\} \\ &= \max\{0, D_1^1 - B_1^1, D_1^1 - (w_2 - w_1), D_0^1 + D_1^1 - B_1^1 - r_{1,1}^1, \\ &\quad D_0^1 + D_1^1 - r_{1,1}^2 - (w_2 - w_1)\} \\ &= \max\{0, D_1^1 - \min\{B_1^1, (w_2 - w_1)\}, D_0^1 + D_1^1 - \min\{B_1^1 + B_0^1, \\ &\quad B_1^1 + (w_2 - w_1), (w_2 - w_1) + B_0^2 \cdot (w_2 - w_1) + (w_3 - w_2)\}\}. \\ &= \max\{0, D_1^1 - r_{2,1}^1, D_0^1 + D_1^1 - r_{2,2}^1\}, \end{aligned}$$

where $r_{2,1}^1 \triangleq \min\{B_1^1, (w_2 - w_1)\}$ and $r_{2,2}^1 \triangleq \min\{B_1^1 + B_0^1, B_1^1 + (w_2 - w_1), (w_2 - w_1) + B_0^2 \cdot (w_2 - w_1) + (w_3 - w_2)\}$. Figure 7.3 depicts a two-level graph in which $r_{2,1}^1$ denotes the length of the shortest path from point 1 at level 0 to level 1, and $r_{2,2}^1$ denotes the length of the shortest path from point 1 at level 0 to level 2. Similar results can be obtained for other stages.




Figure 7.3: The paths of each stage at time slot 2 (two-level graph).

In general, the shortfall of stage 1 at time slot n is given by

$$Y_n^1 = \max \left\{ 0, \max_{1 \leq m \leq n} \left[\sum_{j=1}^m D_{n-j}^1 - r_{n,m}^1 \right] \right\}, \quad (7.8)$$

where $r_{n,m}^1$ is equal to the length of the shortest path from point 1 at level 0 to level m in an n -level graph. A similar characterization of Y_n^1 in terms of shortest paths in a graph is given by Glasserman [Gla97], but for renewal demand and deterministic production processes. As we have argued in the Introduction, and can be seen in the sequel, stochasticity in the production processes and dependencies in all processes involved substantially complicate the picture and require a different and more involved large deviations analysis than the one in [Gla97].

Let us denote by $\{\hat{D}_n^1; n \in \mathbb{Z}\}$ the time-reversed stochastic process obtained from the demand process $\{D_n^1; n \in \mathbb{Z}\}$. In particular, for any $k \in \mathbb{Z}$, $(\hat{D}_1^1, \hat{D}_2^1, \dots, \hat{D}_k^1)$ has the same distribution as $(D_k^1, D_{k-1}^1, \dots, D_1^1)$. Similarly, let $\{\hat{B}_n^i; n \in \mathbb{Z}\}$ denote the time-reversed production process $\{B_n^i; n \in \mathbb{Z}\}$ of stage i , $i = 1, \dots, M$. Notice that due to stationarity $\sum_{j=1}^m D_{n-j}^1$ has the same distribution as $\sum_{j=1}^m \hat{D}_j^1$. More generally, $\sum_{j=k}^l D_j^1$ (or $\sum_{j=k}^l \hat{D}_j^1$) has the same distribution as $\sum_{j=1}^{l-k+1} D_j^1$ (or $\sum_{j=1}^{l-k+1} \hat{D}_j^1$), that is, the distribution of the partial sum of demands (or time-reversed demands) during a time period depends only on

the length of the period and not on the starting time. The same is true for the production processes as well. Moreover, demand and production processes are independent of each other. Using these observations, Y_n^1 has the same distribution as the right hand side of the following equation

$$Y_n^1 \stackrel{D}{=} \max \left\{ 0, \max_{1 \leq m \leq n} \left[\sum_{j=1}^m \hat{D}_j^1 - \min_{\substack{m_1+l_1+m_2+l_2+\dots+m_M=m \\ 0 \leq m_i \leq m, l_i \in \{0,1\} \\ l_i=0 \Rightarrow m_{i+1}, l_{i+1}, \dots, m_M=0}} \left(\sum_{i=1}^{m_1} \hat{B}_i^1 + l_1(w_2 - w_1) \right. \right. \right. \\ \left. \left. \left. + \sum_{i=k_1+1}^{k_1+m_2} \hat{B}_i^2 + l_2(w_3 - w_2) + \dots + \sum_{i=k_{M-1}+1}^{k_{M-1}+m_M} \hat{B}_i^M \right) \right] \right\}, \quad (7.9)$$

where " $\stackrel{D}{=}$ " denotes equality in distribution, and $k_i = i + \sum_{j=1}^i m_j$, for $i = 1, \dots, M-1$. Due to the stability condition (6.1) a steady-state distribution exists for Y_n^1 . In particular, Y_n^1 converges to Y^1 as $n \rightarrow \infty$. Therefore, using (7.9) we obtain

$$Y^1 \stackrel{D}{=} \max_{m \geq 0} \left[S_{1,m}^{\hat{D}^1} - \min_{\substack{m_1+l_1+m_2+l_2+\dots+m_M=m \\ 0 \leq m_i \leq m, l_i \in \{0,1\} \\ l_i=0 \Rightarrow m_{i+1}, l_{i+1}, \dots, m_M=0}} \left(S_{1,m_1}^{\hat{B}^1} + l_1(w_2 - w_1) \right. \right. \\ \left. \left. + S_{k_1+1, k_1+m_2}^{\hat{B}^2} + l_2(w_3 - w_2) + \dots + S_{k_{M-1}+1, k_{M-1}+m_M}^{\hat{B}^M} \right) \right], \quad (7.10)$$

where we use the notation introduced in (1.16) with the convention $\sum_{i=k+1}^k X_i = 0$ for any process $\{X_i; i \in \mathbb{Z}\}$. To facilitate handling the above expression, let us denote by G_m the argument of the maximum, i.e.,

$$Y^1 \stackrel{D}{=} \max_{m \geq 0} G_m.$$

We will proceed with establishing the large deviations result in (7.6). To that end, and in the standard large deviations methodology, we will develop a lower bound and an upper bound and show that the corresponding exponents match. We start from the lower bound. We will use the fact that for a process X and its time-reversed version \hat{X} , $\Lambda_X(\theta) = \Lambda_{\hat{X}}(\theta)$, which can be seen from (1.6). Consequently, $\Lambda_X^*(a) = \Lambda_{\hat{X}}^*(a)$.

7.2.1 Lower Bound

The lower bound result is summarized in the following proposition.

Proposition 7.2.2 *Assume the hedging points w_1, w_2, \dots, w_M in the multi-echelon system (cf. (7.4), (7.5)) satisfy*

$$w_i = \beta_{i-1} w_1, \quad i = 2, \dots, M,$$

where β_i are constants and $1 \leq \beta_1 \leq \dots \leq \beta_{M-1}$. The steady-state shortfall Y^1 of echelon 1 satisfies

$$\liminf_{w_1 \rightarrow \infty} \frac{1}{w_1} \log \mathbf{P}[Y^1 \geq w_1] \geq -\theta_{G,1}^*, \quad (7.11)$$

where $\theta_{G,1}^*$ is given in (7.7).

Proof: For any $m \geq 0$ we have

$$\frac{1}{w_1} \log \mathbf{P}[Y^1 \geq w_1] = \frac{1}{w_1} \log \mathbf{P}[\max_{m \geq 0} G_m \geq w_1] \geq \frac{1}{w_1} \log \mathbf{P}[G_m \geq w_1]. \quad (7.12)$$

Choose $a > 0$ and write $w_1 = ma$. Then $w_i - w_{i-1} = m(\beta_{i-1} - \beta_{i-2})a$ for $i = 2, \dots, M$, where $\beta_0 \triangleq 1$. Using (7.12) we obtain

$$\frac{1}{w_1} \log \mathbf{P}[Y^1 \geq w_1] \geq \frac{1}{ma} \log \mathbf{P}[G_m \geq ma], \quad (7.13)$$

and since we are interested in the regime $w_1 \rightarrow \infty$ it suffices to analyze the behaviour of the right hand side of (7.13) for large values of m . To that end, select $x_i \geq 0$, $i = 0, \dots, M$, $l_i \in \{0, 1\}$, $i = 1, \dots, M - 1$, and m_i , $i = 1, \dots, M$ such that $m_1 + l_1 + m_2 + l_2 + \dots + l_{M-1} + m_M = m$,

$$m x_0 - m_1 x_1 - l_1 m (\beta_1 - 1) a - m_2 x_2 - l_2 m (\beta_2 - \beta_1) a - \dots - m_M x_m = ma,$$

and $l_i = 0$ implies $m_{i+1}, l_{i+1}, \dots, m_M = 0$ for $i = 1, \dots, M - 1$. We have

$$\begin{aligned}
& \mathbf{P}[G_m \geq ma] \\
& \geq \mathbf{P} \left[\max_{\substack{m_1 + l_1 + \dots + m_M = m \\ 0 \leq m_i \leq m, l_i \in \{0,1\} \\ l_i = 0 \Rightarrow m_{i+1}, l_{i+1}, \dots, m_M = 0}} \left(S_{1,m}^{\hat{D}^1} - S_{1,m_1}^{\hat{B}^1} - l_1(w_2 - w_1) \right. \right. \\
& \quad \left. \left. - S_{k_1+1, k_1+m_2}^{\hat{B}^2} - l_2(w_3 - w_2) - \dots - S_{k_{M-1}+1, k_{M-1}+m_M}^{\hat{B}^M} \right) \geq ma \right] \\
& \geq \mathbf{P} \left[S_{1,m}^{\hat{D}^1} - S_{1,m_1}^{\hat{B}^1} - l_1 ma(\beta_1 - 1) - S_{k_1+1, k_1+m_2}^{\hat{B}^2} - \dots - S_{k_{M-1}+1, k_{M-1}+m_M}^{\hat{B}^M} \geq ma \right] \\
& = \mathbf{P} \left[S_{1,m}^{\hat{D}^1} - S_{1,m_1}^{\hat{B}^1} - S_{k_1+1, k_1+m_2}^{\hat{B}^2} - \dots - S_{k_{M-1}+1, k_{M-1}+m_M}^{\hat{B}^M} \right. \\
& \quad \left. \geq ma(1 + l_1(\beta_1 - 1) + l_2(\beta_2 - \beta_1) + \dots + l_{M-1}(\beta_{M-1} - \beta_{M-2})) \right]. \quad (7.14)
\end{aligned}$$

We can distinguish M cases, depending on the values we select for x_i, l_i and m_i . In particular,

Case 1: Select $l_1 = \dots = l_{M-1} = 0$ which implies $m_1 = m$ and $x_0 - x_1 = a$. Then from (7.14) we obtain

$$\begin{aligned}
\mathbf{P}[G_m \geq ma] & \geq \mathbf{P} \left[S_{1,m}^{\hat{D}^1} - S_{1,m}^{\hat{B}^1} \geq ma \right] \\
& \geq \mathbf{P}[S_{1,m}^{\hat{D}^1} \geq mx_0] \mathbf{P}[S_{1,m}^{\hat{B}^1} \leq mx_1] \geq e^{-m[\Lambda_{\hat{D}^1}^{*+}(x_0) + \Lambda_{\hat{B}^1}^{*-}(x_1) + \epsilon]}, \quad (7.15)
\end{aligned}$$

where the last inequality above is due to the LDP principle for the processes \hat{D}^1 and \hat{B}^1 (cf. (1.14) and (1.15)) and holds for large enough m and all $\epsilon > 0$. Using (7.13), taking the limit as $w_1 \rightarrow \infty$, and optimizing over x_0, x_1 and a to obtain a tighter bound, and recalling that demand and production processes have identical large deviations rate functions with their time-reversed versions, we conclude

$$\liminf_{w_1 \rightarrow \infty} \frac{1}{w_1} \log \mathbf{P}[Y^1 \geq w_1] \geq - \inf_{a > 0} \frac{1}{a} \inf_{x_0 - x_1 = a} [\Lambda_{\hat{D}^1}^{*+}(x_0) + \Lambda_{\hat{B}^1}^{*-}(x_1)]. \quad (7.16)$$

Case $i, i = 2, \dots, M$. Select $l_1, \dots, l_{i-1} = 1, l_i, \dots, l_{M-1} = 0$. This implies $x_0 - \xi_1 x_1 -$

$\dots - \xi_i x_i = a\beta_{i-1}$, where $\xi_i = \frac{m_i}{m}$, $i = 1, \dots, M$. Then from (7.14) we obtain

$$\begin{aligned} \mathbf{P}[G_m \geq ma] &\geq \mathbf{P}\left[S_{1,m}^{\hat{D}^1} - S_{1,m_1}^{\hat{B}^1} - \dots - S_{k_{i-1}+1, k_{i-1}+m_i}^{\hat{B}^i} \geq ma\beta_{i-1}\right] \\ &\geq \mathbf{P}[S_{1,m}^{\hat{D}^1} \geq mx_0] \mathbf{P}[S_{1,m_1}^{\hat{B}^1} \leq m_1 x_1] \dots \mathbf{P}[S_{k_{i-1}+1, k_{i-1}+m_i}^{\hat{B}^i} \leq m_i x_i] \\ &\geq e^{-m[\Lambda_{\hat{D}^1}^{*+}(x_0) + \xi_1 \Lambda_{\hat{B}^1}^{*-}(x_1) + \dots + \xi_i \Lambda_{\hat{B}^i}^{*-}(x_i) + \epsilon]}, \end{aligned} \quad (7.17)$$

where the last inequality above is due to the LDP principle for the processes \hat{D}^1 , $\hat{B}^1, \dots, \hat{B}^i$ (cf. (1.14) and (1.15)) and holds for large enough m and all $\epsilon > 0$. Note that since $m_1 + l_1 + m_2 + l_2 + \dots + l_{M-1} + m_M = m$, our selection of l_i 's and m_i 's implies $m_1 + m_2 + \dots + m_i = m - (i - 1)$, which by its turn implies $\xi_1 + \dots + \xi_i = 1$ as $m \rightarrow \infty$. As in case 1, we use (7.13), take the limit as $w_1 \rightarrow \infty$, and optimize over x_0, x_1, \dots, x_i and a to obtain a tighter bound, that is,

$$\begin{aligned} \liminf_{w_1 \rightarrow \infty} \frac{1}{w_1} \log \mathbf{P}[Y^1 \geq w_1] &\geq \\ &- \inf_{a > 0} \frac{1}{a} \inf_{\substack{x_0 - \xi_1 x_1 - \dots - \xi_i x_i = a\beta_{i-1} \\ (\xi_1, \dots, \xi_i) \in \mathcal{O}_i}} [\Lambda_{\hat{D}^1}^{*+}(x_0) + \xi_1 \Lambda_{\hat{B}^1}^{*-}(x_1) + \dots + \xi_i \Lambda_{\hat{B}^i}^{*-}(x_i)]. \end{aligned} \quad (7.18)$$

The tightest lower bound is obtained by summarizing (7.16) and (7.18) for all $i = 2, \dots, M$, i.e.,

$$\liminf_{w_1 \rightarrow \infty} \frac{1}{w_1} \log \mathbf{P}[Y^1 \geq w_1] \geq -\theta_{G,1}^*, \quad (7.19)$$

where $\theta_{G,1}^*$ is given by (7.7). ■

7.2.2 Upper Bound

Next we will establish an upper bound on the probability of interest. The following proposition establishes the result.

Proposition 7.2.3 *Assume the hedging points w_1, w_2, \dots, w_M in the multi-echelon system*

(cf. (7.4), (7.5)) satisfy

$$w_i = \beta_{i-1} w_1, \quad i = 2, \dots, M.$$

where β_i are constants and $1 \leq \beta_1 \leq \dots \leq \beta_{M-1}$. The steady-state shortfall Y^1 of echelon 1 satisfies

$$\limsup_{w_1 \rightarrow \infty} \frac{1}{w_1} \log \mathbf{P}[Y^1 \geq w_1] \leq -\bar{\theta}_{G,1}^*, \quad (7.20)$$

where

$$\bar{\theta}_{G,1}^* \triangleq \min(\theta_1^*, \beta_1 \theta_2^*, \dots, \beta_{M-1} \theta_M^*), \quad (7.21)$$

and where

$$\theta_i^* \triangleq \sup_{\{\theta \geq 0: \sup_{(\xi_1, \dots, \xi_i) \in \mathcal{O}_i} (\Lambda_{D^1}(\theta) + \xi_1 \Lambda_{B^1}(-\theta) + \dots + \xi_i \Lambda_{B^i}(-\theta)) < 0\}} \theta, \quad i = 1, \dots, M. \quad (7.22)$$

Proof: We have

$$\begin{aligned} \mathbf{P}[Y^1 \geq w_1] &= \mathbf{P}[\max_{m \geq 0} G_m \geq w_1] \\ &= \mathbf{P} \left[\max_{m \geq 0} \left[S_{1,m}^{\hat{D}^1} - \min_{\substack{m_1 + l_1 + m_2 + l_2 + \dots + m_M = m \\ 0 \leq m_i \leq m, l_i \in \{0,1\} \\ l_i = 0 \Rightarrow m_{i+1}, l_{i+1}, \dots, m_M = 0}} \left(S_{1,m_1}^{\hat{B}^1} + l_1(w_2 - w_1) \right. \right. \right. \\ &\quad \left. \left. \left. + \dots + S_{k_{M-1}+1, k_{M-1}+m_M}^{\hat{B}^M} \right) \right] \geq w_1 \right] \\ &= \mathbf{P} \left[\max \left\{ \max_{m \geq 0} \left(S_{1,m}^{\hat{D}^1} - S_{k_1+1, k_1+m}^{\hat{B}^1} \right), \dots, \right. \right. \\ &\quad \left. \left. \max_{m_1 + \dots + m_M = m - (M-1)} \left(S_{1,m}^{\hat{D}^1} - S_{1,m_1}^{\hat{B}^1} - (w_2 - w_1) - \dots - S_{k_{M-1}+1, k_{M-1}+m_M}^{\hat{B}^M} \right) \right\} \geq w_1 \right] \\ &\leq \mathbf{P} \left[\max_{m \geq 0} \left(S_{1,m}^{\hat{D}^1} - S_{1,m}^{\hat{B}^1} \right) \geq w_1 \right] + \dots \\ &\quad + \mathbf{P} \left[\max_{m_1 + \dots + m_M = m - (M-1)} \left(S_{1,m}^{\hat{D}^1} - S_{1,m_1}^{\hat{B}^1} - \dots - S_{k_{M-1}+1, k_{M-1}+m_M}^{\hat{B}^M} \right) \geq \beta_{M-1} w_1 \right]. \end{aligned} \quad (7.23)$$

In the second equality above we consider all possible sample paths that can lead to a value larger than w_1 . In particular, the first such sample path corresponds to $l_1 = 0$, the i th sample path corresponds to $l_1 = \dots = l_{i-1} = 1$ and $l_i = 0$, for $i = 2, \dots, M - 1$, and the M th sample path corresponds to $l_1 = \dots = l_{M-1} = 1$. The first inequality above bounds the probability of the maximum of all those sample paths by the sum of the individual probabilities. Hence, it suffices to bound each the probabilities in the right hand side of the above. We distinguish M cases:

Case 1. For the first probability in the right hand side of (7.23) and for $\theta \geq 0$ we have

$$\begin{aligned}
\mathbf{P} \left[\max_{m \geq 0} \left(S_{1,m}^{\hat{D}^1} - S_{1,m}^{\hat{B}^1} \right) \geq w_1 \right] &\leq \mathbf{E} \left[e^{\theta \max_{m \geq 0} (S_{1,m}^{\hat{D}^1} - S_{1,m}^{\hat{B}^1})} \right] e^{-\theta w_1} \\
&\leq \sum_{m \geq 0} \mathbf{E} \left[e^{\theta (S_{1,m}^{\hat{D}^1} - S_{1,m}^{\hat{B}^1})} \right] e^{-\theta w_1} \\
&\leq \left[K_1'(\theta) + \sum_{m \geq m_0} e^{m(\Lambda_{\hat{D}^1}(\theta) + \Lambda_{\hat{B}^1}(-\theta) + \epsilon_1)} \right] e^{-\theta w_1} \\
&\leq K_1(\theta, \epsilon_1) e^{-\theta w_1}, \quad \text{if } \Lambda_{\hat{D}^1}(\theta) + \Lambda_{\hat{B}^1}(-\theta) < 0,
\end{aligned} \tag{7.24}$$

where m_0 is sufficiently large and $\epsilon_1 > 0$. In the first inequality above we used the Markov inequality. In the third inequality above we have split the summation in two parts. Specifically, terms corresponding to $m = 0, \dots, m_0$ are summarized in $K_1'(\theta)$. For the remaining terms we use the existence of the limiting log-moment generating function (cf. Equation (1.6)). Finally, in the last inequality above, since the exponent is negative (for sufficiently small ϵ_1) the infinite geometric series converges to some $K_1''(\theta, \epsilon_1)$ which when combined with $K_1'(\theta)$ yields $K_1(\theta, \epsilon_1)$. Optimizing now over θ to obtain the tightest bound yields

$$\begin{aligned}
\mathbf{P} \left[\max_{m \geq 0} \left(S_{1,m}^{\hat{D}^1} - S_{1,m}^{\hat{B}^1} \right) \geq w_1 \right] &\leq \\
&\inf_{\{\theta \geq 0: \Lambda_{\hat{D}^1}(\theta) + \Lambda_{\hat{B}^1}(-\theta) < 0\}} K_1(\theta, \epsilon_1) e^{-\theta w_1} \leq K_1(\theta_1^*, \epsilon_1) e^{-\theta_1^* w_1}, \tag{7.25}
\end{aligned}$$

where

$$\theta_i^* \triangleq \sup_{\{\theta \geq 0: \Lambda_{\hat{D}^i}(\theta) + \Lambda_{\hat{B}^i}(-\theta) < 0\}} \theta. \quad (7.26)$$

Case i , $i = 2, \dots, M$. For the i th probability in the right hand side of (7.23) and for $\theta \geq 0$ we have

$$\begin{aligned} & \mathbf{P} \left[\max_{\substack{m \geq 0 \\ m_1 + m_2 + \dots + m_i = m - (i-1)}} \left(S_{1,m}^{\hat{D}^i} - S_{1,m_1}^{\hat{B}^i} - \dots - S_{k_{i-1}+1, k_{i-1}+m_i}^{\hat{B}^i} \right) \geq \beta_{i-1} w_1 \right] \\ & \leq \mathbf{E} \left[\exp \left\{ \theta \max_{\substack{m \geq 0 \\ m_1 + \dots + m_i = m - (i-1)}} \left(S_{1,m}^{\hat{D}^i} - S_{1,m_1}^{\hat{B}^i} - \dots - S_{k_{i-1}+1, k_{i-1}+m_i}^{\hat{B}^i} \right) \right\} \right] e^{-\theta \beta_{i-1} w_1} \\ & \leq \sum_{\substack{m \geq 0 \\ m_1 + \dots + m_i = m - (i-1)}} \mathbf{E} \left[e^{\theta (S_{1,m}^{\hat{D}^i} - S_{1,m_1}^{\hat{B}^i} - \dots - S_{k_{i-1}+1, k_{i-1}+m_i}^{\hat{B}^i})} \right] e^{-\theta \beta_{i-1} w_1} \\ & \leq \sum_{m \geq 0} m^i \sup_{m_1 + \dots + m_i = m - (i-1)} \mathbf{E} \left[e^{\theta (S_{1,m}^{\hat{D}^i} - S_{1,m_1}^{\hat{B}^i} - \dots - S_{k_{i-1}+1, k_{i-1}+m_i}^{\hat{B}^i})} \right] e^{-\theta \beta_{i-1} w_1} \\ & \leq \left[\sum_{m \geq m_0} m^i K_i'(\theta, \epsilon_i) \sup_{(\xi_1, \dots, \xi_i) \in \mathcal{O}_i} e^{m(\Lambda_{\hat{D}^i}(\theta) + \xi_1 \Lambda_{\hat{B}^i}(-\theta) + \dots + \xi_i \Lambda_{\hat{B}^i}(-\theta) + \epsilon_i)} \right] e^{-\theta \beta_{i-1} w_1} \\ & \leq K_i(\theta, \epsilon_i) e^{-\theta \beta_{i-1} w_1}, \quad \text{if } \sup_{(\xi_1, \dots, \xi_i) \in \mathcal{O}_i} (\Lambda_{\hat{D}^i}(\theta) + \xi_1 \Lambda_{\hat{B}^i}(-\theta) + \dots + \xi_i \Lambda_{\hat{B}^i}(-\theta)) < 0, \end{aligned} \quad (7.27)$$

where $\xi_i = m_i/m$, $i = 1, \dots, M$, m_0 is sufficiently large and $\epsilon_i > 0$. As in Case 1, in the first inequality above we used the Markov inequality, in the fourth inequality above we used the existence of the limiting log-moment generating functions, and in the last inequality above we used the fact that the infinite geometric series converges if the exponent is negative. Optimizing over θ to obtain the tightest bound

$$\begin{aligned} & \mathbf{P} \left[\max_{\substack{m \geq 0 \\ m_1 + \dots + m_i = m - (i-1)}} \left(S_{1,m}^{\hat{D}^i} - S_{1,m_1}^{\hat{B}^i} - \dots - S_{k_{i-1}+1, k_{i-1}+m_i}^{\hat{B}^i} \right) \geq \beta_{i-1} w_1 \right] \\ & \leq K_i(\theta_i^*, \epsilon_i) e^{-\theta_i^* \beta_{i-1} w_1}, \end{aligned}$$

where

$$\theta_i^* \triangleq \sup_{\{\theta \geq 0: \sup_{(\xi_1, \dots, \xi_i) \in \mathcal{O}_i} (\Lambda_{D^1}(\theta) + \xi_1 \Lambda_{B^1}(-\theta) + \dots + \xi_i \Lambda_{B^i}(-\theta)) < 0\}} \theta.$$

Summarizing Cases 1, ..., M and using (7.23) we obtain that for all $\epsilon_1, \dots, \epsilon_M > 0$ and for some $K_1(\theta_1^*, \epsilon_1), \dots, K_M(\theta_M^*, \epsilon_M)$

$$\mathbf{P}[Y^1 \geq w_1] \leq K_1(\theta_1^*, \epsilon_1) e^{-\theta_1^* w_1} + \dots + K_M(\theta_M^*, \epsilon_M) e^{-\theta_M^* \beta_{M-1} w_1}. \quad (7.28)$$

Letting $w_1 \rightarrow \infty$ we conclude

$$\limsup_{w_1 \rightarrow \infty} \frac{1}{w_1} \log \mathbf{P}[Y^1 \geq w_1] \leq -\bar{\theta}_{G,1}, \quad (7.29)$$

where

$$\bar{\theta}_{G,1} \triangleq \min(\theta_1^*, \beta_1 \theta_2^*, \dots, \beta_{M-1} \theta_M^*).$$

■

7.2.3 Upper and Lower Bounds Match

Finally, we will show that the upper bound has the same exponent as the lower bound.

Proposition 7.2.4 *It holds*

$$\theta_{G,1}^* = \bar{\theta}_{G,1},$$

where $\theta_{G,1}^*$ and $\bar{\theta}_{G,1}$ are defined in (7.7) and (7.21), respectively.

Proof: It suffices to show that

$$\inf_{a > 0} \frac{1}{a} \inf_{\substack{x_0 - \xi_1 x_1 - \dots - \xi_i x_i = a \beta_{i-1} \\ (\xi_1, \dots, \xi_i) \in \mathcal{O}_i}} (\Lambda_{D^1}^{*+}(x_0) + \xi_1 \Lambda_{B^1}^{*-}(x_1) + \dots + \xi_i \Lambda_{B^i}^{*-}(x_i)) = \beta_{i-1} \theta_i^* =$$

$$\beta_{i-1} \left(\sup_{\{\theta \geq 0: \sup_{(\xi_1, \dots, \xi_i) \in \mathcal{O}_i} (\Lambda_{D^i}(\theta) + \xi_1 \Lambda_{B^i}(-\theta) + \dots + \xi_i \Lambda_{B^i}(-\theta)) < 0\}} \theta \right),$$

for all $i = 1, \dots, M$, where $\beta_0 \triangleq 1$. To that end, notice that

$$\beta_{i-1} \left(\sup_{\{\theta \geq 0: \sup_{(\xi_1, \dots, \xi_i) \in \mathcal{O}_i} (\Lambda_{D^i}(\theta) + \xi_1 \Lambda_{B^i}(-\theta) + \dots + \xi_i \Lambda_{B^i}(-\theta)) < 0\}} \theta \right) \stackrel{\circ}{=} \sup_{\{\theta \geq 0: \sup_{(\xi_1, \dots, \xi_i) \in \mathcal{O}_i} (\Lambda_{D^i}(\theta/\beta_{i-1}) + \xi_1 \Lambda_{B^i}(-\theta/\beta_{i-1}) + \dots + \xi_i \Lambda_{B^i}(-\theta/\beta_{i-1})) < 0\}} \theta.$$

To proceed we will use the following Lemma which was shown in Bertsimas, Paschalidis, and Tsitsiklis [BPT99].

Lemma 7.2.5 ([BPT99, Lemma 6.2]) *For $\Lambda^*(\cdot)$ and $\Lambda(\cdot)$ being convex duals and assuming that $\Lambda(\theta) < 0$ for sufficiently small $\theta > 0$, it holds that*

$$\inf_{a > 0} \frac{1}{a} \Lambda^*(a) = \theta^*,$$

where θ^* is the largest root of the equation $\Lambda(\theta) = 0$.

Notice next that

$$\inf_{\substack{x_0 - \xi_1 x_1 - \dots - \xi_i x_i = a\beta_{i-1} \\ (\xi_1, \dots, \xi_i) \in \mathcal{O}_i}} (\Lambda_{D^i}^{*+}(x_0) + \xi_1 \Lambda_{B^i}^{*-}(x_1) + \dots + \xi_i \Lambda_{B^i}^{*-}(x_i)) \quad (7.30)$$

is a convex function of a as the value function of a convex optimization problem with a appearing only in the right hand side of the constraints. Moreover, it can be shown that it is lower-semicontinuous (by [BPT99, Lemma 6.3]), and thus we can apply convex duality results and use Lemma 7.2.5. Finally, for any $a_j \geq 0$, $j = 1, \dots, i$, with $\xi_1 + \dots + \xi_i = 1$, $\Lambda_{D^i}(\theta/\beta_{i-1}) + \xi_1 \Lambda_{B^i}(-\theta/\beta_{i-1}) + \dots + \xi_i \Lambda_{B^i}(-\theta/\beta_{i-1})$ is equal to zero at $\theta = 0$ and has

negative derivative by the stability condition (6.1), which implies that

$$\sup_{(\xi_1, \dots, \xi_i) \in \mathcal{O}_i} (\Lambda_{D^1}(\theta/\beta_{i-1}) + \xi_1 \Lambda_{B^1}(-\theta/\beta_{i-1}) + \dots + \xi_i \Lambda_{B^i}(-\theta/\beta_{i-1})) \quad (7.31)$$

takes negative values for sufficiently small $\theta > 0$. As a final step we show that the expression in (7.30) is the convex dual of the expression in (7.31). Indeed we have

$$\begin{aligned} & \sup_a \left\{ \theta a - \inf_{\substack{x_0 - \xi_1 x_1 - \dots - \xi_i x_i = a\beta_{i-1} \\ (\xi_1, \dots, \xi_i) \in \mathcal{O}_i}} (\Lambda_{D^1}^{*+}(x_0) + \xi_1 \Lambda_{B^1}^{*-}(x_1) + \dots + \xi_i \Lambda_{B^i}^{*-}(x_i)) \right\} \\ &= \sup_a \sup_{\substack{x_0 - \xi_1 x_1 - \dots - \xi_i x_i = a\beta_{i-1} \\ (\xi_1, \dots, \xi_i) \in \mathcal{O}_i}} \{ \theta a - \Lambda_{D^1}^{*+}(x_0) - \xi_1 \Lambda_{B^1}^{*-}(x_1) - \dots - \xi_i \Lambda_{B^i}^{*-}(x_i) \} \\ &= \sup_{\substack{x_0, x_1, \dots, x_i \\ (\xi_1, \dots, \xi_i) \in \mathcal{O}_i}} \left\{ \theta \frac{x_0 - \xi_1 x_1 - \dots - \xi_i x_i}{\beta_{i-1}} - \Lambda_{D^1}^{*+}(x_0) - \xi_1 \Lambda_{B^1}^{*-}(x_1) - \dots - \xi_i \Lambda_{B^i}^{*-}(x_i) \right\} \\ &= \sup_{(\xi_1, \dots, \xi_i) \in \mathcal{O}_i} \{ \Lambda_{D^1}(\theta/\beta_{i-1}) + \xi_1 \Lambda_{B^1}(-\theta/\beta_{i-1}) + \dots + \xi_i \Lambda_{B^i}(-\theta/\beta_{i-1}) \}. \end{aligned}$$

■

Combining Propositions 7.2.2, 7.2.3 and 7.2.4 we obtain the main result summarized in Theorem 7.2.1. Some remarks are in order.

Remarks:

1. Theorem 7.2.1 provides us with the asymptotic decay rate for the overflow probability of the shortfall, or equivalently, with the asymptotic decay rate of the stockout probability for the echelon inventory at stage 1. More intuitively, Theorem 7.2.1 asserts that

$$\mathbf{P}[X^1 \leq 0] = \mathbf{P}[Y^1 \geq w_1] \sim e^{-\theta_{G,1}^* w_1}. \quad (7.32)$$

2. The proof of Theorem 7.2.1 characterizes the most likely path that leads to stockouts and provides intuition on how they occur. Recall from the proof that we have shown

(cf. Equation (7.21))

$$\theta_{G,1}^* = \min(\theta_1^*, \beta_1 \theta_2^*, \dots, \beta_{M-1} \theta_M^*),$$

where θ_i^* , $i = 1, \dots, M$, is the largest root of the equation $\sup_{(\xi_1, \dots, \xi_i) \in \mathcal{O}_i} (\Lambda_{D^1}(\theta) + \xi_1 \Lambda_{B^1}(-\theta) + \dots + \xi_i \Lambda_{B^i}(-\theta)) = 0$ (cf. Equation (7.22)). Consider the case $\theta_{G,1}^* = \beta_{i-1} \theta_i^*$ for some $i = 1, \dots, M$, where $\beta_0 \triangleq 1$. To avoid degenerate cases assume that all production processes B^i have distinct limiting log-moment generating functions and that $1 < \beta_1 < \dots < \beta_{M-1}$. Let ξ_j^* , $j = 1, \dots, i$, be the optimal solution of the optimization problem $\sup_{(\xi_1, \dots, \xi_i) \in \mathcal{O}_i} (\Lambda_{D^1}(\theta) + \xi_1 \Lambda_{B^1}(-\theta) + \dots + \xi_i \Lambda_{B^i}(-\theta))$ at $\theta = \theta_i^*$. It can be seen that one of the ξ_j^* 's, is equal to 1. In particular, ξ_i^* is equal to 1; otherwise, i.e., if $\xi_j^* = 1$ for some $j < i$, $\theta_j^* = \theta_i^*$ and $\beta_{j-1} \theta_j^*$ will be the minimizer in the definition of $\theta_{G,1}^*$ since $\beta_{j-1} \theta_j^* < \beta_{i-1} \theta_i^*$. Therefore, θ_i^* is the largest root of the equation $\Lambda_{D^1}(\theta) + \Lambda_{B^i}(-\theta) = 0$ and the stockout probability at stage 1 behaves as the exponential

$$e^{-\beta_{i-1} \theta_i^* w_1} = e^{-\theta_i^* w_1}.$$

Considering the single stage result (cf. Proposition 5.2.1) we can say that stage i production capacity is the “bottleneck” and characterizes the stockout probability at stage 1.

3. Suppose that $\beta = (\beta_1, \dots, \beta_{M-1}) \rightarrow \infty$. Then from Equation (7.21) we have

$$\lim_{\beta \rightarrow \infty} \theta_{G,1}^* = \theta_{L,1}^*,$$

where $\theta_{L,1}^*$ is the largest root of the equation $\Lambda_{D^1}(\theta) + \Lambda_{B^1}(-\theta) = 0$. This is consistent with the result of Theorem 6.2.2. Essentially, as $\beta \rightarrow \infty$ stages decompose and stage 1 is not affected by the upstream material requirement constraint, which makes Theorem

6.2.2 accurate. In general, Equation (7.21) shows that

$$\theta_{G,1}^* \leq \theta_{L,1}^*, \quad (7.33)$$

and Theorem 6.2.2 underestimates the stockout probability.

The result of Theorem 7.2.1 can be easily generalized to yield the steady-state stock-out probability of the echelon inventory X^i at the remaining stages $i = 2, \dots, M$. More specifically, we can think of echelon inventory X^i at stage i , $i = 1, \dots, M$ as the echelon 1 inventory of an $(M + 1 - i)$ -stage supply chain starting at the i th stage of the original system. Hence, generalizing Theorem 7.2.1 we obtain the following corollary.

Corollary 7.2.6 *Assume the base-stock levels w_i, w_{i+1}, \dots, w_M , for $i = 1, \dots, M$, in the multi-echelon system satisfy*

$$w_{i+k} = \beta_{i+k-1}^i w_i, \quad k = 1, \dots, M - i,$$

where β_{i+k-1}^i are constants and $1 \leq \beta_i^i \leq \dots \leq \beta_{M-1}^i$. The steady-state shortfall Y^i of echelon i satisfies

$$\lim_{w_i \rightarrow \infty} \frac{1}{w_i} \log \mathbf{P} [Y^i \geq w_i] = -\theta_{G,i}^*, \quad (7.34)$$

where $\theta_{G,i}^*$ is determined by

$$\begin{aligned} \theta_{G,i}^* = \min & \left[\inf_{a>0} \frac{1}{a} \inf_{x_0 - x_i = a} (\Lambda_{D^i}^{*+}(x_0) + \Lambda_{B^i}^{*-}(x_i)), \right. \\ & \inf_{a>0} \frac{1}{a} \inf_{\substack{x_0 - \xi_i x_i - \xi_{i+1} x_{i+1} = a\beta_i^i \\ (\xi_i, \xi_{i+1}) \in \mathcal{O}_2}} (\Lambda_{D^i}^{*+}(x_0) + \xi_i \Lambda_{B^i}^{*-}(x_i) + \xi_{i+1} \Lambda_{B^{i+1}}^{*-}(x_{i+1})), \dots, \\ & \left. \inf_{a>0} \frac{1}{a} \inf_{\substack{x_0 - \xi_i x_i - \dots - \xi_M x_M = a\beta_{M-1}^i \\ (\xi_i, \dots, \xi_M) \in \mathcal{O}_{M-i+1}}} (\Lambda_{D^i}^{*+}(x_0) + \xi_i \Lambda_{B^i}^{*-}(x_i) + \dots + \xi_M \Lambda_{B^M}^{*-}(x_M)) \right]. \quad (7.35) \end{aligned}$$

7.3 Refining the Large Deviations Asymptotics

Next we will discuss heuristics for refining the large deviations asymptotics. Without loss of generality we will concentrate on stage 1. The discussion can be easily extended to the remaining stages based on Corollary 7.2.6.

Theorem 7.2.1 provides us with the asymptotic decay rate for the stockout probability of the echelon-1 inventory as its base-stock level goes to infinity. This leads to the following approximation

$$\mathbf{P}[Y^1 \geq w_1] \sim e^{-\theta_{G,1}^* w_1}. \quad (7.36)$$

To improve the accuracy of the approximation, especially for relatively large stockout probabilities (i.e., small safety stock w_1), we will introduce a prefactor in front of the exponential. This is in accordance with the development in Section 6.2 where we used a constant prefactor (cf. (5.8)). A constant prefactor was also used in improving the large deviations approximation in the multiclass single-stage case considered in Bertsimas and Paschalidis [BP01]. Here, instead, we will use the following refined approximation

$$\mathbf{P}[Y^1 \geq w_1] \approx f_1(w_1, \boldsymbol{\beta}) e^{-\theta_{G,1}^* w_1}, \quad (7.37)$$

where the prefactor $f_1(w_1, \boldsymbol{\beta})$ is a function of w_1 and $\boldsymbol{\beta} = (\beta_1, \dots, \beta_{M-1}) = (\frac{w_2}{w_1}, \dots, \frac{w_M}{w_1})$. As we commented in Section 7.2, as $\boldsymbol{\beta} \rightarrow \infty$, different stages decouple and the upstream material requirement constraint becomes insignificant. Thus, to be consistent with the analysis of Section 6.2 we will select a function $f_1(w_1, \boldsymbol{\beta})$ that converges to a constant as $\boldsymbol{\beta} \rightarrow \infty$. In particular, $\lim_{\boldsymbol{\beta} \rightarrow \infty} f_1(w_1, \boldsymbol{\beta}) = c_1$ and $\lim_{\boldsymbol{\beta} \rightarrow \infty} \theta_{G,1}^* = \theta_{L,1}^*$, where c_1 is equal to the constant prefactor used in (5.8).

We will be using a function which is piecewise linear in w_1 and $\boldsymbol{\beta}$ to represent $f_1(w_1, \boldsymbol{\beta})$. More specifically, we will evaluate the stockout probability $\mathbf{P}[Y^1 \geq w_1]$ at several sample points $\mathbf{w} = (w_1, \dots, w_M)$ by simulation and then find a piecewise linear function $f_1(w_1, \boldsymbol{\beta})$ so that $f_1(w_1, \boldsymbol{\beta}) e^{-\theta_{G,1}^* w_1}$ matches the true value of $\mathbf{P}[Y^1 \geq w_1]$ at those sample points.

This requires two main steps: (i) selecting appropriate sample points $\mathbf{w} = (w_1, \dots, w_M)$, or equivalently points $(w_1, \boldsymbol{\beta}) = (w_1; \beta_1, \dots, \beta_{M-1}) = (w_1; \frac{w_2}{w_1}, \dots, \frac{w_M}{w_1})$, and (ii) given a “data set” of $((w_1, \boldsymbol{\beta}); f_1(w_1, \boldsymbol{\beta}))$ pairs, “fit” a function $f_1(w_1, \boldsymbol{\beta})$ to the points in the data set. We will start from step (i).

We are interested in selecting sample points $(w_1, \boldsymbol{\beta})$ such that w_1 values are scattered in \mathbb{R}^+ , $\boldsymbol{\beta}$ are in the feasible set $\mathcal{B} = \{\boldsymbol{\beta} : 1 \leq \beta_1 \leq \dots \leq \beta_{M-1}\}$, and adequately capture the form of $f_1(w_1, \boldsymbol{\beta})$ in all parts of the feasible set. Recall from the proof of Theorem 7.2.1 that (Equation (7.21))

$$\theta_{G,1}^* = \min(\theta_1^*, \beta_1 \theta_2^*, \dots, \beta_{M-1} \theta_M^*),$$

where θ_i^* , $i = 1, \dots, M$, are defined in Equation (7.22). This characterization of $\theta_{G,1}^*$ separates the feasible set \mathcal{B} into up to M polyhedral regions depending on which term is the minimizer in (7.21) (i.e., region i contains those $\boldsymbol{\beta}$ that lead to $\theta_{G,1}^* = \beta_{i-1} \theta_i^*$, $i = 1, \dots, M$, where $\beta_0 \triangleq 1$). We expect the form of $f_1(w_1, \boldsymbol{\beta})$ to be rather different in each of those regions, hence, to achieve enough “variety” in selecting the sample points we will pick a number of points in each of those M regions. In particular, we will pick sample points on the vertices and extreme rays of the M polyhedral regions; any other point in \mathcal{B} can be expressed as a convex combination of those. We can potentially pick some sample points in the interior of each polyhedral region; this would lead to a better approximation of the prefactor. For illustrative purpose, Figure 7.4 presents an example when $M = 3$.

It should be noted that when the supply chain has a large number of stages, the number of vertices of the M polyhedral regions discussed above is large and it is computationally expensive to evaluate the stockout probability at all such vertices. Instead, we can select much fewer sample points and extrapolate to obtain $f_1(\boldsymbol{\beta})$. We provide a two-stage example in Figure 7.5.

Suppose next that we have selected a set of $K \times N$ sample points $(w_1^k, \boldsymbol{\beta}^i)$, $k = 1, \dots, K$, $i = 1, \dots, N$, $w_1^1 \leq \dots \leq w_1^K$. Let $\mathbf{w}^{k,i} = (w_1^k, w_1^k \beta_1^i, \dots, w_1^k \beta_{M-1}^i)$ be the hedging point vector corresponding to $(w_1^k, \boldsymbol{\beta}^i)$. We simulate the system with each sample point $\mathbf{w}^{k,i}$,




Figure 7.4: We depict the feasible set \mathcal{B} in a three-dimensional example. The set \mathcal{B} consists of three regions: region I, where $\theta_{G,1}^* = \theta_1^*$, region II, where $\theta_{G,1}^* = \beta_1 \theta_2^*$, and region III, where $\theta_{G,1}^* = \beta_2 \theta_3^*$. As sample points we select points $Q_1, Q_2, Q_3, Q_4, Q_5, Q_6, Q_7$, and potentially some additional points scattered in regions I–III.

$k = 1, \dots, K, i = 1, \dots, N$ and obtain the stockout probability $\mathbf{P}[Y^1 \geq w_1^k]$. We compute

$$f_1(w_1^k, \beta^i) = \frac{\mathbf{P}[Y^1 \geq w_1^k]}{e^{-\theta_{G,1}^* w_1^k}}.$$

Thus, we construct a data set consisting of pairs $((w_1^k, \beta^i); f_1(w_1^k, \beta^i))$. To make this procedure computationally efficient and reduce the required simulation time we can select sample points $\mathbf{w}^{1,1}, \dots, \mathbf{w}^{K,N}$ with relatively small safety stocks which do not lead to very small stockout probabilities (such probabilities require long simulation running times). The key point here is that we use large deviations analysis to obtain the exponent of the stockout probability; the approximation in (7.36) improves as stockouts become more rare. We only use simulation to refine the approximation and we will use sample points in the regime where stockouts are not particularly rare and thus easier to simulate. We will later provide numerical results (Section 7.6) demonstrating that the proposed procedure leads to very accurate approximations.

We will now turn our attention to step (ii) mentioned above. That is, we assume we have a “data set” consisting of $K \times N$ pairs $((w_1, \beta); f_1(w_1, \beta))$ and wish to construct a

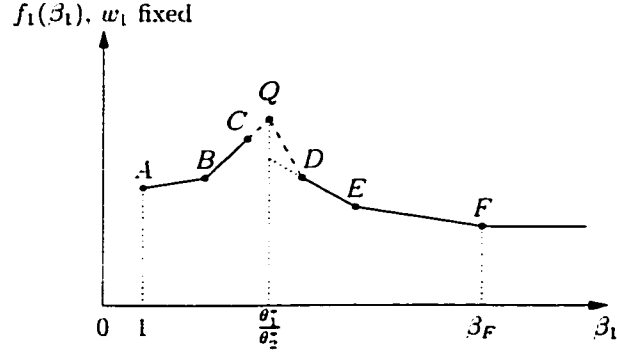


Figure 7.5: We consider a system with two stages and assume that $\theta_1^* > \theta_2^*$. For a fixed w_1 , $f_1(w_1, \beta_1)$ is a function of β_1 . A, B, C, D, E, F are sample points. Suppose that we do not know $f_1(w_1, \frac{\theta_1^*}{\theta_2^*})$. We can approximate it by extrapolating segment BC and DE as shown in the graph. To maintain the continuity of $f_1(w_1, \beta_1)$ and convexity in region $[1, \frac{\theta_1^*}{\theta_2^*}]$ and $[\frac{\theta_1^*}{\theta_2^*}, \infty)$, we select Q as $f_1(w_1, \frac{\theta_1^*}{\theta_2^*})$. For points to the right of F , we set them to be constant and equal to $f_1(w_1, \beta_F)$.

function $f_1(w_1, \beta)$ that matches the given data points. More specifically, we will be using a function $f_1(w_1, \beta)$ which is for a fixed w_1 , piecewise linear and convex function of β in each of the M polyhedral regions comprising \mathcal{B} , and for a fixed β , piecewise linear function of $w_1 \in \mathbb{R}^+$. The selection of such a function is motivated by Equation (7.36). Note that due to (7.21)

$$e^{-\theta_{G,1}^* w_1} = \max \left(e^{-\theta_1^* w_1}, e^{-\beta_1 \theta_2^* w_1}, \dots, e^{-\beta_{M-1} \theta_M^* w_1} \right),$$

which is convex in β_i in each of the M polyhedral regions, for all $i = 1, \dots, M-1$. Hence, based on the approximation in (7.36), $\mathbf{P}[Y^1 \geq w_1]$ is convex in β_i in each of those regions for large values of w_1 . Thus, we select a convex function $f_1(w_1, \beta)$ to retain convexity of $\mathbf{P}[Y^1 \geq w_1]$ for small values of w_1 as can be seen by the refined approximation in (7.37).

Let us next fix some arbitrary $\mathbf{w} = (w_1, \dots, w_M)$, or $(w_1, \beta) = (w_1; \frac{w_2}{w_1}, \dots, \frac{w_M}{w_1})$ and assume that β belongs into the polyhedral region corresponding to $\theta_{G,1}^* = \beta_{j-1} \theta_j^*$, for some $j = 1, \dots, M$. Let $((w_1^k, \beta^i); f_1(w_1^k, \beta^i))$, $k = 1, \dots, K$, $i = 1, \dots, n$ be the subset of the data set with β^i in that region (including points on the boundary and the vertices of the region derived from extrapolation). We distinguish two cases:

1. Assume that β is in the convex hull of β^1, \dots, β^n . Solve the following linear program-

ming (LP) problem for each $k = 1, \dots, K$:

$$\begin{aligned}
\min_{\zeta_i} \quad & \zeta_1 f_1(w_1^k, \beta^1) + \dots + \zeta_n f_1(w_1^k, \beta^n) \\
\text{s.t.} \quad & \zeta_1 \beta^1 + \dots + \zeta_n \beta^n = \beta \\
& \zeta_1 + \dots + \zeta_n = 1 \\
& \zeta_i \geq 0, \quad i = 1, \dots, n.
\end{aligned} \tag{7.38}$$

Note that this LP is always feasible since β is in the convex hull of β^1, \dots, β^n . Let $(\zeta_1^k, \dots, \zeta_n^k)$ be an optimal solution. Then we can set $f_1(w_1^k, \beta)$ as follows

$$f_1(w_1^k, \beta) = \sum_{i=1}^n \zeta_i^k f_1(w_1^k, \beta^i).$$

2. Assume β is not in the convex hull of β^1, \dots, β^n . Note that in this case the LP in (7.38) is infeasible, and β is in the unbounded region, e.g., the points to the right of point F in Figure 7.5, which means some components of β is fairly large, the decoupling effects take place as we discussed in Section 6.2. Therefore we can use a constant to approximate $f_1(\cdot)$. As suggested in the example of Figure 7.5 we will find the projection of β onto the convex hull of β^1, \dots, β^n , say β^\perp , and $f_1(w_1^k, \beta^\perp)$ by solving (7.38). We set $f_1(w_1^k, \beta) = f_1(w_1^k, \beta^\perp)$, e.g., $f_1(w_1, \beta_1) = f_1(w_1, \beta_F)$, for all $\beta_1 > \beta_F$ in Figure 7.5.

After we obtain $f_1(w_1^1, \beta), \dots, f_1(w_1^K, \beta)$, we can construct a piecewise linear function of w_1 based on these K sample points (using extrapolation for $w_1 < w_1^1$ or $w_1 > w_1^K$), then find the value of $f_1(w_1, \beta)$.

In summary, the procedure described above selects a number of sample points in $(\mathbb{R}^+, \mathcal{B})$ and obtains a function $f_1(w_1, \beta)$ which is piecewise linear convex in β in each of its regions and piecewise linear in w_1 . $f_1(w_1, \beta)$ will be used in Equation (7.37) to yield an analytical approximation of the stockout probability.

It should be noted that the heuristic procedure we discussed which approximates $f_1(w_1, \beta)$

is one potential approach. Alternatively, given a data set consisting of $K \times N$ sample points in $(\mathbb{R}^+, \mathcal{B})$ we can approximate $f_1(w_1, \beta)$ by some parametric form (e.g., some polynomial function or even a neural network) and then use a least squares procedure to “fit” the parametric form on the data set.

7.4 Approximating the expected inventory cost

The main motivation for analyzing the echelon inventory policy was to acquire the flexibility to reduce expected inventory costs by trading-off inventory between various stages, while, at the same time, maintaining service level constraints. To that end, we need to assess expected inventory costs.

We will assume linear inventory costs. Let h_i be the holding cost for echelon- i inventory for all $i = 1, \dots, M$. Noting that the expected echelon- i inventory is given by $\mathbf{E}[I^i] + \dots + \mathbf{E}[I^2] + \mathbf{E}[(I^1)^+]$, where $(I^1)^+ = \max(I^1, 0)$, the total expected inventory cost is given by

$$h_1 \mathbf{E}[(I^1)^+] + h_2 (\mathbf{E}[(I^1)^+] + \mathbf{E}[I^2]) + \dots + h_M (\mathbf{E}[(I^1)^+] + \mathbf{E}[I^2] + \dots + \mathbf{E}[I^M]). \quad (7.39)$$

We have

$$\begin{aligned} \mathbf{E}[(I^1)^+] &= \mathbf{E}[(w_1 - Y^1)^+] \\ &= \mathbf{E}[\max(w_1 - Y^1, 0)] \\ &= w_1 - \mathbf{E}[Y^1] + \mathbf{E}[\max(0, Y^1 - w_1)]. \end{aligned} \quad (7.40)$$

Using the tail distribution of Y^1 given in Equation (7.37), we obtain

$$\begin{aligned} \mathbf{E}[\max(0, Y^1 - w_1)] &= \int_0^\infty \mathbf{P}[\max(0, Y^1 - w_1) > y] dy \\ &= \int_0^\infty \mathbf{P}[Y^1 - w_1 > y] dy \\ &\approx f_1(w_1, \beta) e^{-\theta_{G,1}^* w_1} \int_0^\infty e^{-\theta_{G,1}^* y} dy \end{aligned}$$

$$= f_1(w_1, \beta) \frac{e^{-\theta_{G,1}^* w_1}}{\theta_{G,1}^*}. \quad (7.41)$$

For all $i \geq 2$, we have

$$I^i = (w_i - Y^i) - (w_{i-1} - Y^{i-1}),$$

which implies

$$\mathbf{E}[I^i] = (w_i - \mathbf{E}[Y^i]) - (w_{i-1} - \mathbf{E}[Y^{i-1}]). \quad (7.42)$$

Thus, combining (7.39), (7.40), (7.41), and (7.42), the total expected inventory cost can be approximated by the following expression

$$\begin{aligned} \sum_{i=1}^M h_i(w_i - \mathbf{E}[Y^i]) + (h_1 + \dots + h_M) \mathbf{E}[(Y^1 - w_1)^+] = \\ \sum_{i=1}^M h_i(w_i - \mathbf{E}[Y^i]) + (h_1 + \dots + h_M) f_1(w_1, \beta) \frac{e^{-\theta_{G,1}^* w_1}}{\theta_{G,1}^*}. \end{aligned} \quad (7.43)$$

To obtain an analytical approximation for the inventory cost we are now left with computing $\mathbf{E}[Y^i]$. This is hard to do analytically; instead we will use an approach similar to the one used in obtaining $f_1(w_1, \beta)$. We will first establish some structural properties for $\mathbf{E}[Y^i]$.

Proposition 7.4.1 *Consider the multi-echelon system (cf. (7.4), (7.5)) and let $0 \leq w_1 \leq w_2 \leq \dots \leq w_M$ be the corresponding hedging points. Define $\Delta_i \triangleq w_{i+1} - w_i$, for $i = 1, \dots, M-1$. Then $\mathbf{E}[Y^M]$ is a constant function of $(\Delta_1, \dots, \Delta_{M-1})$. Furthermore, for all $i = 1, \dots, M-1$, $\mathbf{E}[Y^i]$ is a function of $(\Delta_i, \dots, \Delta_{M-1})$, which is convex and monotonically nonincreasing in every coordinate. In addition, as $\Delta_i, \dots, \Delta_{M-1} \rightarrow \infty$, $\mathbf{E}[Y^i]$ converges to a constant.*

Proof: Recall from (7.4) and (7.5) that the shortfalls satisfy the following evolution equations:

$$Y_{n+1}^i = \max\{Y_n^i + D_n^1 - B_n^i, 0, Y_n^{i+1} + D_n^1 - \Delta_i\}, \quad i = 1, \dots, M-1, \quad (7.44)$$

$$Y_{n+1}^M = \max\{Y_n^M + D_n^1 - B_n^M, 0\}. \quad (7.45)$$

Due to the stability condition (6.1), a steady-state distribution exists for Y_n^i . In particular, Y_n^i converges as $n \rightarrow \infty$ to Y^i . From the evolution equations (7.44) and (7.45) it is clear that $\mathbf{E}[Y^M]$ is a constant function of $(\Delta_1, \dots, \Delta_{M-1})$.

Consider next the echelon inventory at stage $M - 1$ in three distinct systems A and B , and C . System A operates with hedging points satisfying $\Delta_{M-1} = \Delta_A$. System B operates with hedging points satisfying $\Delta_{M-1} = \Delta_B$. System C operates with hedging points satisfying $\Delta_C = \alpha\Delta_A + (1 - \alpha)\Delta_B$, where $0 \leq \alpha \leq 1$. Assume without loss of generality that $\Delta_A < \Delta_B$. Let $Y_{A,n}^{M-1}$, $Y_{B,n}^{M-1}$, and $Y_{C,n}^{M-1}$ be the echelon shortfall at stage $M - 1$ for systems A , B , and C , respectively, during time slot n . We define the demand and production processes for all systems A , B , and C on the same probability space so that they are driven by identical sample paths. As a result, the echelon- M shortfall in all three systems is identical for all time slots n ; we will denote it by Y_n^M . We have

$$\begin{aligned} Y_{A,n+1}^{M-1} &= \max\{Y_{A,n}^{M-1} + D_n^1 - B_n^{M-1}, 0, Y_n^M + D_n^1 - \Delta_A\}, \\ Y_{B,n+1}^{M-1} &= \max\{Y_{B,n}^{M-1} + D_n^1 - B_n^{M-1}, 0, Y_n^M + D_n^1 - \Delta_B\}, \end{aligned}$$

and

$$Y_{C,n+1}^{M-1} = \max\{Y_{C,n}^{M-1} + D_n^1 - B_n^{M-1}, 0, Y_n^M + D_n^1 - \Delta_C\}.$$

At time 0, let $Y_{A,0}^{M-1} = Y_{B,0}^{M-1} = Y_{C,0}^{M-1} = 0$, which satisfy

$$\alpha Y_{A,0}^{M-1} + (1 - \alpha) Y_{B,0}^{M-1} \geq Y_{C,0}^{M-1}.$$

Suppose at time slot n ,

$$\alpha Y_{A,n}^{M-1} + (1 - \alpha) Y_{B,n}^{M-1} \geq Y_{C,n}^{M-1}.$$

At time $n + 1$,

$$\begin{aligned} &\alpha Y_{A,n+1}^{M-1} + (1 - \alpha) Y_{B,n+1}^{M-1} \\ &= \alpha \max\{Y_{A,n}^{M-1} + D_n^1 - B_n^{M-1}, 0, Y_n^M + D_n^1 - \Delta_A\} + \end{aligned}$$

$$\begin{aligned}
& (1 - \alpha) \max\{Y_{B,n}^{M-1} + D_n^1 - B_n^{M-1}, 0, Y_n^M + D_n^1 - \Delta_B\} \\
\geq & \max\{\alpha(Y_{A,n}^{M-1} + D_n^1 - B_n^{M-1}) + (1 - \alpha)(Y_{B,n}^{M-1} + D_n^1 - B_n^{M-1}), 0, \\
& \alpha(Y_n^M + D_n^1 - \Delta_A) + (1 - \alpha)(Y_n^M + D_n^1 - \Delta_B)\} \\
= & \max\{(\alpha Y_{A,n}^{M-1} + (1 - \alpha)Y_{B,n}^{M-1}) + D_n^1 - B_n^{M-1}, 0, Y_n^M + D_n^1 - (\alpha\Delta_A + (1 - \alpha)\Delta_B)\} \\
\geq & \max\{Y_{C,n}^{M-1} + D_n^1 - B_n^{M-1}, 0, Y_n^M + D_n^1 - \Delta_C\} \\
= & Y_{C,n+1}^{M-1}.
\end{aligned}$$

Therefore, for all time slots n ,

$$\alpha Y_{A,n}^{M-1} + (1 - \alpha)Y_{B,n}^{M-1} \geq Y_{C,n}^{M-1},$$

which implies

$$\alpha \mathbf{E}[Y_A^{M-1}] + (1 - \alpha)\mathbf{E}[Y_B^{M-1}] \geq \mathbf{E}[Y_C^{M-1}].$$

Thus, $\mathbf{E}[Y^{M-1}]$ is a convex function of Δ_{M-1} . Furthermore, from (7.44) it can be easily seen that $\mathbf{E}[Y^{M-1}]$ is nonincreasing in Δ_{M-1} and as $\Delta_{M-1} \rightarrow \infty$ converges to a constant. In particular, it converges to the expected shortfall of a single-stage system with demand D^1 and capacity B^{M-1} (decoupled system).

Similarly, it can be shown that $\mathbf{E}[Y^i]$, which is a function of $(\Delta_i, \dots, \Delta_{M-1})$, is convex and nonincreasing in Δ_i and that it converges to a constant as $\Delta_i \rightarrow \infty$. Following a similar procedure, it can also be shown that for all sample paths and all time slots n , Y_{n+1}^i is a convex and nondecreasing function of Y_n^{i+1} , which by its turn is convex and nonincreasing in Δ_{i+1} , and convex and nondecreasing in Y_{n-1}^{i+2} . Therefore, $\mathbf{E}[Y^i]$ is convex and nonincreasing in Δ_{i+1} . Continuing in this fashion, we conclude that $\mathbf{E}[Y^i]$ is a function of $(\Delta_i, \dots, \Delta_{M-1})$ which is convex and nonincreasing in every coordinate. Furthermore, as $\Delta_i, \dots, \Delta_{M-1} \rightarrow \infty$, $\mathbf{E}[Y^i]$ converges to a constant. In particular, it converges to the expected shortfall of a single-stage system with demand D^1 and capacity B^i (decoupled system). ■

Motivated by these properties of $\mathbf{E}[Y^i]$ we will approximate it using a piecewise linear convex function $g_i(\Delta_i, \dots, \Delta_{M-1})$, using a similar approach to the one used in approximating $f_1(\beta)$ in Section 7.3. More specifically we will be using the following approximation

$$\mathbf{E}[Y^i] = g_i(w_{i+1} - w_i, \dots, w_M - w_{M-1}), \quad i = 1, \dots, M - 1. \quad (7.46)$$

As in Section 7.3 we can select a number of sample points \mathbf{w}^j , $j = 1, \dots, N$, and construct a piecewise linear convex function that matches $\mathbf{E}[Y^i]$ at those sample points. Note that one can evaluate $\mathbf{E}[Y^i]$ from the same simulation used to evaluate $\mathbf{P}[Y^i \geq w_i]$, thus, the same set of sample points and simulation runs can be used to construct both $g_i(\cdot)$ and $f_1(\cdot)$.

We now have all the ingredients to pose the problem of optimizing expected inventory costs subject to maintaining service level constraints. Using the approximating expression for the expected inventory cost in (7.43) we have the following optimization problem.

$$\begin{aligned} \min_{w_i} \quad & \sum_{i=1}^M h_i(w_i - g_i(w_{i+1} - w_i, \dots, w_M - w_{M-1})) \\ & + \left(\sum_{i=1}^n h_i \right) f_1(w_1; w_2/w_1, \dots, w_M/w_1) \frac{e^{-\theta_{G,1}^* w_1}}{\theta_{G,1}^*} \\ \text{s.t.} \quad & \mathbf{P}[Y^i \geq w_i] = f_i(w_i; w_{i+1}/w_i, \dots, w_M/w_i) e^{-\theta_{G,i}^* w_i} \leq \epsilon_i, \quad i = 1, \dots, M, \\ & w_M \geq \dots \geq w_2 \geq w_1 \geq 0. \end{aligned} \quad (7.47)$$

This problem can be solved analytically using standard nonlinear programming techniques. Since there are a number of approximations involved in this formulation, it is of interest to assess the accuracy of the solution when compared with “brute force” simulation. We will see in Section 7.6 that the solution predicted by the problem in (7.47) is accurate. The very significant advantage of our approach is that we can set the proper hedging points analytically, which leads to huge computational savings.

7.5 Extensions: The Multiclass Case and Lost Sales

In this section we discuss two simple extensions: (i) to a supply chain model that can accommodate multiple classes, and (ii) to a model where unsatisfied demand at stage 1 is lost instead of backlogged.

7.5.1 The Multiclass Case

In this multiclass case, a production policy consists of scheduling decisions as well. That is, at each point in time, and at each facility in the chain, we have not only to decide whether the facility will be working or not, but also decide which products it will producing, if any. Finding an optimal production policy to minimize expected inventory costs subject to service level constraints is intractable, even in a single stage system. Bertsimas and Paschalidis [BP01] have proposed production policies in the multiclass, single stage, problem by using fluid model analysis to obtain a scheduling policy and large deviations analysis for the idling policy. In this section we will use a scheduling policy that is motivated by fairness considerations and ease of analysis.

We extend the model depicted in Figure 6.1 as follows. We assume that instead of a single product class the system produces K products. We will maintain separate finished goods inventory buffers for each product class in front of stage 1. We let $I_n^{k,1}$ denote the class k inventory at stage 1 for $k = 1, \dots, K$. Similarly, we will maintain separate inventory buffers in front of each upstream stage $2, \dots, M$. We will be denoting by $I_n^{k,j}$ the class k inventory at stage j , for $k = 1, \dots, K$ and $j = 2, \dots, M$. Finally, we let $D_n^{k,1}$ denote the amount of external orders arriving at stage 1 from class k during time slot n .

We will implement a scheduling policy which allocates a constant fraction of the capacity of each facility to every class. In particular, we will let $\phi_{k,i}$ denote the fraction of the stage- i capacity B_n^i allocated to class k during time slot n , for all $k = 1, \dots, K$ and $i = 1, \dots, M$, where $\sum_{k=1}^K \phi_{k,i} = 1$. Note that $\phi_{k,i}$ is constant for all time slots. This policy will be referred to as the *generalized processor sharing* policy (GPS) and has in fact been analyzed in the

large deviations regime by Bertsimas, Paschalidis and Tsitsiklis [BPT99] for the two-class case; an approximate analysis for more than two classes can be found in Paschalidis [Pas99]. The same policy has been used by Glasserman [Gla96] in a multiclass make-to-stock system. The GPS policy is attractive because it guarantees a minimum fraction of the capacity to every class. Thus, it can be viewed as fair since the performance of a class cannot be compromised at times that other classes are congested, as might happen for example with a priority policy.

Notice, next, that according to the GPS policy the capacity allocated to a class k can be distributed to the remaining classes during times that class k has no work to be done. This allocation of the unutilized capacity can be done according to the weights $\phi_{k,i}$. As a result, classes are “coupled”, which leads to a rather involved large deviations analysis (see [BPT99]). To facilitate the analysis in our supply chain model we will make a simplifying assumption. More specifically, we will decompose the system across classes and ignore the unutilized capacity allocated to a class during times that other classes are not busy. A similar decomposition assumption has been made in [Gla96]. Hence, the multiclass supply chain is decomposed in K single class chains and the results we have developed are immediately applicable. In particular, our single class asymptotics and hedging points can be derived for each class k by using capacity $\phi_{i,k}B_n^i$ at each stage i during time slot n . The limiting log-moment generating function and the corresponding large deviations rate function of the process $\{\phi_{i,k}B_n^i; n \in \mathbb{Z}\}$ can be easily derived from $\Lambda_{B^i}(\theta)$ and $\Lambda_{B^i}^*(a)$, respectively. Of course, for stability purposes we have to assume

$$\mathbf{E}[D_n^{k,1}] < \min_{i=1,\dots,m} \phi_{i,k} \mathbf{E}[B_n^i],$$

for all $k = 1, \dots, K$.

7.5.2 A Model with Lost Sales

We next turn our attention to a model where if inventory is not available, external demand is lost and not backlogged. We will start the discussion with the multi-echelon inventory model.

Consider the supply chain model of Figure 7.1 operating under the echelon inventory policy of Section 7.2. Assume that unsatisfied demand is lost. Our notation for the lost sales system will parallel the one we used in Section 7.2. Let \tilde{X}_n^i and \tilde{Y}_n^i denote the echelon inventory and shortfall, respectively, at time slot n for stage $i = 1, \dots, M$. Let also $\mathbf{w} = (w_1, \dots, w_M)$ denote the hedging point vector. We can obtain an evolution equation for \tilde{X}_n^i (respectively, \tilde{Y}_n^i) by introducing a reflecting boundary at zero (respectively, w_i) in Equations (7.2) and (7.3) (respectively, (7.4) and (7.5)). In particular, for \tilde{Y}_n^i we have

$$\tilde{Y}_{n+1}^i = \min\{\max\{\tilde{Y}_n^i + D_n^1 - B_n^i, 0, \tilde{Y}_n^{i+1} + D_n^1 - (w_{i+1} - w_i)\}, w_i\}, \quad (7.48)$$

$$i = 1, \dots, M - 1,$$

$$\tilde{Y}_{n+1}^M = \min\{\max\{\tilde{Y}_n^M + D_n^1 - B_n^M, 0\}, w_M\}. \quad (7.49)$$

In the lost sales system the steady-state stockout probability is $\mathbf{P}[\tilde{X}^i = 0]$ or, equivalently, $\mathbf{P}[\tilde{Y}^i = w_i]$. As in Sections 7.2–7.4, our objective is to minimize the expected inventory cost subject to maintaining these probabilities below given thresholds ϵ_i for each stage i .

Note that the steady-state stockout probability at stage 1 can be interpreted as the long-term average fraction of time that the system has no stock (under ergodicity assumptions). This can be connected with the percentage of orders that are lost. Consider the case of a Bernoulli demand process (i.e., D_n^i is one with probability p and zero otherwise at each time slot n and independently of anything else in the system). Then the steady state probability that an order is lost is $p\mathbf{P}[\tilde{Y}^i = w_i]$, which is the same as the expected amount of lost sales (in product time). The same reasoning does not apply for arbitrary demand processes, since they may not see “time average”; the probability that an order is lost will depend on the distribution of the demand process. To avoid such complications and have a measure

that depends on the system we opted for steady-state stockout probability to construct service-level constraints.

The main result is that the large deviations exponent of the stockout probability in the lost sales model is the same as the one in the model with backorders considered in Section 7.2. More specifically, we will require the demand and production processes satisfy the following version of a *sample path large deviations principle (SLDP)* (see [BPT99] for an extended discussion on SPLDPs). Let $\{X_j; j \in \mathbb{Z}\}$ denote any of the demand or production processes, and let $S_{j,k}^X = \sum_{r=j}^k X_r$ denote the partial sums of the process X . We will be assuming that for all $n \geq K$ and all k_0, \dots, k_m with $1 = k_0 \leq k_1 \leq \dots \leq k_m = n$,

$$e^{-(n\epsilon_2 + \sum_{i=0}^{m-1} (k_{i+1} - k_i) \Lambda_X^*(a_i))} \leq \mathbf{P}[|S_{k_i+1, k_{i+1}}^X - (k_{i+1} - k_i) a_i| \leq \epsilon_1 n, i = 0, \dots, m-1]. \quad (7.50)$$

Intuitively, this assumption deals with the probability of sample paths that are constrained to be within a tube around a “polygonal” path made up with linear segments of slopes a_0, \dots, a_{m-1} . This assumption is satisfied by a large class of processes, including renewal, Markov-modulated, and stationary processes with mild mixing conditions (see [Cha95]). It can also be seen from the derivation of the large deviations rate function in [Cha95] that the time-reversed process \tilde{X} has the same large deviations rate function as the forward process.

We first provide an alternative expression for $\theta_{G,1}^*$ in (7.7).

Lemma 7.5.1 *It holds*

$$\theta_{G,1}^* = \min \left[\inf_{a>0} \frac{1}{a} \inf_{x_0 - x_1 = a} (\Lambda_{D^1}^*(x_0) + \Lambda_{B^1}^*(x_1)), \right. \\ \left. \inf_{a>0} \frac{1}{a} \inf_{\substack{x_0 - \xi_1 x_1 - \xi_2 x_2 = a\beta_1 \\ (\xi_1, \xi_2) \in \mathcal{O}_2}} (\Lambda_{D^1}^*(x_0) + \xi_1 \Lambda_{B^1}^*(x_1) + \xi_2 \Lambda_{B^2}^*(x_2)), \dots, \right. \\ \left. \inf_{a>0} \frac{1}{a} \inf_{\substack{x_0 - \xi_1 x_1 - \dots - \xi_M x_M = a\beta_{M-1} \\ (\xi_1, \dots, \xi_M) \in \mathcal{O}_M}} (\Lambda_{D^1}^*(x_0) + \xi_1 \Lambda_{B^1}^*(x_1) + \dots + \xi_M \Lambda_{B^M}^*(x_M)) \right]. \quad (7.51)$$

Proof: Comparing (7.7) and (7.51) it suffices to show

$$\inf_{a>0} \frac{1}{a} \inf_{\substack{x_0 - \xi_1 x_1 - \dots - \xi_i x_i = a\beta_{i-1} \\ (\xi_1, \dots, \xi_i) \in \mathcal{O}_i}} (\Lambda_{D^1}^{*+}(x_0) + \xi_1 \Lambda_{B^1}^{*-}(x_1) + \dots + \xi_i \Lambda_{B^i}^{*-}(x_i)) =$$

$$\inf_{a>0} \frac{1}{a} \inf_{\substack{x_0 - \xi_1 x_1 - \dots - \xi_i x_i = a\beta_{i-1} \\ (\xi_1, \dots, \xi_i) \in \mathcal{O}_i}} (\Lambda_{D^1}^*(x_0) + \xi_1 \Lambda_{B^1}^*(x_1) + \dots + \xi_i \Lambda_{B^i}^*(x_i))$$

for some arbitrary $i = 1, \dots, M$. We will denote by LHS (respectively, RHS) the expression on the left hand side (respectively, right hand side) of the above.

First observe that $\text{LHS} \leq \text{RHS}$, since for any X and any a we have $\Lambda_X^{*-}(a) \leq \Lambda_X^*(a)$ and $\Lambda_X^{*+}(a) \leq \Lambda_X^*(a)$ (cf. (1.10) and (1.11)).

Next consider an optimal solution $\mathbf{y} = (a^*, x_0^*, \dots, x_i^*, \xi_1^*, \dots, \xi_i^*)$ of the optimization problem in the LHS. Without loss of generality, assume that $\xi_1^*, \dots, \xi_i^* > 0$; otherwise some terms will be eliminated from the objective function and the rest of the proof carries through. Fix $\epsilon > 0$, sufficiently small. We will construct a feasible solution $\mathbf{y}'(\epsilon)$ of the LHS that is also optimal. We will distinguish several cases:

1. Suppose $x_0^* \geq \mathbf{E}[D^1]$ and $x_j^* \leq \mathbf{E}[B^j]$ for all $j = 1, \dots, i$. Then set $\mathbf{y}' = \mathbf{y}$.
2. Suppose $x_0^* < \mathbf{E}[D^1]$. Note that by feasibility $x_0^* - \xi_1^* x_1^* - \dots - \xi_i^* x_i^* \geq 0$. This implies that for some $j = 1, \dots, i$, $x_j^* < \mathbf{E}[B^j]$. Otherwise, i.e., if $x_j^* \geq \mathbf{E}[B^j]$ for all $j = 1, \dots, i$, we have

$$x_0^* - \xi_1^* x_1^* - \dots - \xi_i^* x_i^* < \mathbf{E}[D_1] - \xi_1^* \mathbf{E}[B^1] - \dots - \xi_i^* \mathbf{E}[B^i] \leq \mathbf{E}[D_1] - \min_{j \in \{1, \dots, i\}} \mathbf{E}[B^j] < 0,$$

by the stability condition (6.1). Then set

$$\mathbf{y}' = (a^*, x_0^* + \epsilon, x_1^*, \dots, x_{j-1}^*, x_j^* + \frac{\epsilon}{\xi_j^*}, x_{j+1}^*, \dots, x_i^*, \xi_1^*, \dots, \xi_i^*),$$

which is feasible. Note that since $\Lambda_{D^1}^{*+}(\cdot)$ is equal to zero below the mean $\mathbf{E}[D^1]$ and $\Lambda_{B^j}^{*-}(\cdot)$ is nonincreasing below the mean $\mathbf{E}[B^j]$ the objective value of the optimization problem in the LHS at \mathbf{y}' is no more than the corresponding value at \mathbf{y} . Hence, \mathbf{y}' is also optimal.

3. Suppose that for some $j, j = 1, \dots, i, x_j^* > \mathbf{E}[B^j]$. We distinguish two cases:

(a) Suppose $x_0^* > \mathbf{E}[D^1]$. Then set

$$\mathbf{y}' = (a^*, x_0^* - \xi_j^* \epsilon, x_1^*, \dots, x_{j-1}^*, x_j^* - \epsilon, x_{j+1}^*, \dots, x_i^*, \xi_1^*, \dots, \xi_i^*),$$

which is feasible. Note that since $\Lambda_{D^1}^{*+}(\cdot)$ is nondecreasing above the mean $\mathbf{E}[D^1]$ and $\Lambda_{B^j}^{*-}(\cdot)$ is equal to zero above the mean $\mathbf{E}[B^j]$ the objective value of the optimization problem in the LHS at \mathbf{y}' is no more than the corresponding value at \mathbf{y} . Hence, \mathbf{y}' is also optimal.

(b) Finally, suppose that $x_0^* = \mathbf{E}[D^1]$. Then the same argument as in Case 2 above establishes that for some $j' = 1, \dots, i$ not equal to j we have $x_{j'}^* < \mathbf{E}[B^{j'}]$. Then set

$$\mathbf{y}' = (a^*, x_0^*, \dots, x_{j-1}^*, x_j^* - \frac{\epsilon}{\xi_j^*}, x_{j+1}^*, \dots, x_{j'-1}^*, x_{j'}^* + \frac{\epsilon}{\xi_{j'}^*}, x_{j'+1}^*, \dots, x_i^*, \xi_1^*, \dots, \xi_i^*),$$

which is feasible. Note that since $\Lambda_{B^{j'}}^{*-}(\cdot)$ is nonincreasing below the mean $\mathbf{E}[B^{j'}]$ and $\Lambda_{B^j}^{*-}(\cdot)$ is equal to zero above the mean $\mathbf{E}[B^j]$ the objective value of the optimization problem in the LHS at \mathbf{y}' is no more than the corresponding value at \mathbf{y} . Hence, \mathbf{y}' is also optimal.

Given \mathbf{y} we keep repeating the procedure in Case 2 and 3 above until we construct a new optimal solution $\mathbf{y}' = (a', x'_0, \dots, x'_i, \xi'_1, \dots, \xi'_i)$ that satisfies $x'_0 \geq \mathbf{E}[D^1]$ and $x'_j \leq \mathbf{E}[B^j]$ for all $j = 1 \dots, i$. Such a \mathbf{y}' is feasible for the optimization problem in the RHS and achieves the same objective function value for both RHS and LHS. The optimal value of the RHS can be no worse. Thus, $\text{RHS} \leq \text{LHS}$. ■

Under somewhat more restrictive assumptions on the demand and production processes (some form of a sample path large deviations principle as we stated before) we obtain the following theorem.

Theorem 7.5.2 *Assume the hedging points w_1, w_2, \dots, w_m in the multi-echelon lost sales system satisfy*

$$w_{i+1} = \beta_i w_1, \quad i = 1, \dots, M - 1,$$

where β_i are constants and $1 \leq \beta_1 \leq \dots \leq \beta_{M-1}$. The steady-state shortfall \tilde{Y}^1 of echelon 1 satisfies

$$\lim_{w_1 \rightarrow \infty} \frac{1}{w_1} \log \mathbf{P}[\tilde{Y}^1 = w_1] = -\theta_{G,1}^*, \quad (7.52)$$

where $\theta_{G,1}^*$ is given by (7.7).

Proof: Recall that for each time slot n the lost sales system satisfies the evolution equations (7.48) and (7.49), while the system with backorders satisfies Equations (7.4) and (7.5). We define demand and production processes on the same probability space for both systems so that they are driven by identical sample paths. We observe that for all n it holds $\tilde{Y}_n^1 \leq Y_n^1$. Hence, by using Proposition 7.2.3 and 7.2.4 we obtain

$$\limsup_{w_1 \rightarrow \infty} \frac{1}{w_1} \log \mathbf{P}[\tilde{Y}^1 = w_1] \leq \limsup_{w_1 \rightarrow \infty} \frac{1}{w_1} \log \mathbf{P}[Y^1 \geq w_1] \leq -\theta_{G,1}^*.$$

For the lower bound we will mimic the proof of Proposition 7.2.2. The key of that proof is that we identified M scenarios (Case 1, \dots , M) which led to $Y^1 \geq w_1$. The probabilities of these scenarios provide a set of M lower bounds; we select the tighter by maximizing over those. Here we will establish that the scenarios provided there are also feasible scenarios in the lost sales model and lead to $\tilde{Y}^1 = w_1$. Using the same notation as in the proof of Proposition 7.2.2, let m be large enough, choose $a > 0$, set $w_1 = ma$ and consider the following M scenarios:

Scenario 1.

$$\{|S_{1,j}^{\hat{D}^1} - jx_0| \leq \epsilon_0 m, j = 1 \dots, m\}, \quad \{|S_{1,j}^{\hat{B}^1} - jx_1| \leq \epsilon_1 m, j = 1 \dots, m\},$$

where $x_0, x_1 \geq 0$, $\epsilon_0, \epsilon_1 > 0$, $x_0 - x_1 = a + \epsilon'$, and $\epsilon' = \epsilon_0 + \epsilon_1$.

Scenarios $i = 2, \dots, M$.

$$\begin{aligned} & \{|S_{1,j}^{\hat{D}^1} - jx_0| \leq \epsilon_0 m, j = 1, \dots, m\}, \\ & \{|S_{1,j}^{\hat{B}^1} - jx_1| \leq \epsilon_1 m_1, j = 1, \dots, m_1\}, \dots, \{|S_{k_{i-1}+1, k_{i-1}+j}^{\hat{B}^i} - jx_i| \leq \epsilon_i m_i, j = 1, \dots, m_i\}, \end{aligned}$$

where $x_0, \dots, x_i \geq 0$, $\epsilon_0, \dots, \epsilon_i > 0$, $x_0 - \xi_1 x_1 - \dots - \xi_i x_i = (a + \epsilon')\beta_{i-1}$, $\xi_i = m_j/m$ for $j = 1, \dots, i$, $(\xi_1, \dots, \xi_i) \in \mathcal{O}_i$, and $\epsilon' = \epsilon_0 + \dots + \epsilon_i$, and $k_{i-1} = (i-1) + \sum_{j=1}^{i-1} m_j$.

Using the same arguments in the proof of Proposition 7.2.2, according to scenario i the shortfall in the system with backorders builds up linearly with m at a rate of $a + \epsilon'$, where $\epsilon' \rightarrow 0$ as $\epsilon_0, \dots, \epsilon_i \rightarrow 0$. It reaches $m(a + \epsilon')$ in m time slots. Now from (7.48) and (7.49) note that starting from zero \tilde{Y}_n^1 and Y_n^1 follow identical sample paths until they hit w_1 . Hence, \tilde{Y}_n^1 reaches $w_1 = ma$ in m time slots. Thus, using the same notation as in the proof of Proposition 7.2.2, for every $i = 1, \dots, M$ we have

$$\begin{aligned} & \mathbf{P}[\tilde{Y}^1 = ma] \\ & \geq \mathbf{P}[\min\{G_m, ma\} = ma] \\ & \geq \mathbf{P}[|S_{1,j}^{\hat{D}^1} - jx_0| \leq \epsilon_0 m, j = 1, \dots, m] \times \mathbf{P}[|S_{1,j}^{\hat{B}^1} - jx_1| \leq \epsilon_1 m_1, j = 1, \dots, m_1] \times \dots \\ & \quad \times \mathbf{P}[|S_{k_{i-1}+1, k_{i-1}+j}^{\hat{B}^i} - jx_i| \leq \epsilon_i m_i, j = 1, \dots, m_i] \\ & \geq e^{-m(\Lambda_{D^1}^*(x_0) + \xi_1 \Lambda_{B^1}^*(x_1) + \dots + \xi_i \Lambda_{B^i}^*(x_i) + \epsilon)}, \end{aligned}$$

where m is large enough, $\epsilon \rightarrow 0$ as $\epsilon_0, \dots, \epsilon_i \rightarrow 0$, and the last inequality above is due to the SPLDP assumption in (7.50). As in the proof of Proposition 7.2.2 we optimize over all parameters of scenarios i to obtain

$$\begin{aligned} \liminf_{w_1 \rightarrow \infty} \frac{1}{w_1} \log \mathbf{P}[\tilde{Y}^1 = w_1] & \geq \\ & - \inf_{a > 0} \frac{1}{a} \inf_{\substack{x_0 - \xi_1 x_1 - \dots - \xi_i x_i = a\beta_{i-1} \\ (\xi_1, \dots, \xi_i) \in \mathcal{O}_i}} (\Lambda_{D^1}^*(x_0) + \xi_1 \Lambda_{B^1}^*(x_1) + \dots + \xi_i \Lambda_{B^i}^*(x_i)). \end{aligned}$$

By using Lemma 7.5.1 and selecting the tightest bound among all scenarios $1, \dots, M$ we obtain

$$\liminf_{w_1 \rightarrow \infty} \frac{1}{w_1} \log \mathbf{P}[\tilde{Y}^1 = w_1] \geq -\theta_{G,1}^*.$$

■

Following an analysis that parallel the one in Section 7.4 we can also obtain the total expected inventory cost in the lost system. It is given by

$$\sum_{i=1}^M h_i(w_i - \mathbf{E}[\tilde{Y}^i]).$$

Thus, to obtain the hedging point vector we can construct an optimization problem similar to the one in (7.47).

A lost sales extensions to the local inventory case, handled by our decomposition approach in Chapter 5 and Chapter 6, appears to be more involved. The single-stage result of Proposition 5.2.1 can be readily extended to the lost sales model by taking the limit $\beta = (\beta_1, \dots, \beta_{M-1}) \rightarrow \infty$ in Theorem 7.5.2 (cf. Remark 3 in Section 7.2.3). By doing this, we obtain that the decay rate of the stockout probability in a single-stage lost sales model is the largest root of $\Lambda_D(\theta) + \Lambda_B(-\theta) = 0$, where D and B denote the demand and production process, respectively. Using our decomposition technique, though, to handle the multi-stage case requires characterizing the departure process of a $G/G/1$ queue (Theorem 6.2.1). In the lost sales model, one needs to extend that result and characterizes the departure process of a $G/G/1$ queue with a finite buffer. We conjecture that this is doable along the lines of the result in [BPT98b]. A simple approximation can be easily developed by using the departure process in an infinite buffer $G/G/1$ queue to obtain a bound on the large deviations rate function of the departure process from a queue with a finite buffer. Using such an approximation one can apply our results in Chapter 6 to treat the local inventory case with lost sales.

7.6 Numerical Results

In this section, we present numerical results to evaluate the performance of the proposed large deviations approximations. We will consider a two-stage system and we will (a) use the decomposition approach developed in Section 6.2 to derive a base-stock policy for each stage under a variety of service level requirements, and (b) use the echelon base-stock policy analyzed in Section 7.1 to optimize the expected inventory cost subject to service level constraints. We will also present an example demonstrating that detailed distributional information on the demand and production processes is critical in making inventory control decisions.

Throughout this section we consider Markov-modulated demand and production processes. Figure 7.6 depicts the model of the demand and production processes in a two-stage supply-chain. We will be referring to this system as *Example 1*. Notice that according to the mean production capacities the bottleneck is the first stage. We construct *Example 2* by exchanging the order of the two production facilities. The 3-stage example we will consider will be referred to as *Example 3*.

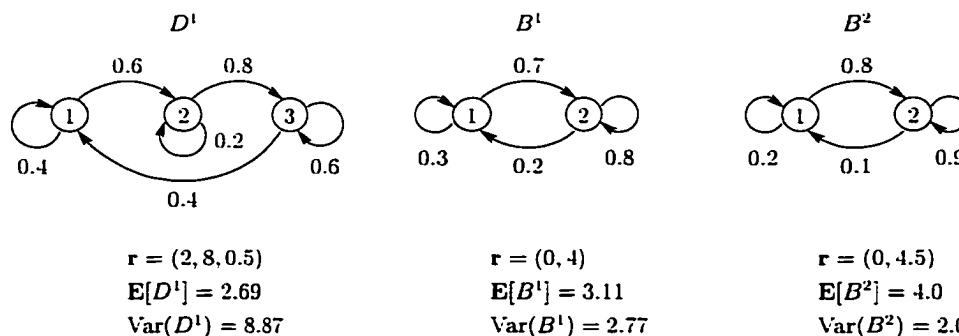


Figure 7.6: The models of demand and production processes in Example 1, a two-stage system. We denote by \mathbf{r} the vector of demand or production amounts at each state of the corresponding Markov chain.

7.6.1 A Two-Stage Supply Chain with the Decomposition Approach

For both examples we use the approach developed in Section 6.2. Using the result of Theorem 6.2.2 we compute the asymptotic decay rate of the stockout probability at each

stage, namely,

$$\theta_{L,1}^* = 0.093, \quad \theta_{L,2}^* = 0.334,$$

for Example 1 and

$$\theta_{L,1}^* = 0.258, \quad \theta_{L,2}^* = 0.093,$$

for Example 2. To compute analytically the hedging points we use the expression in (6.20). To compute the prefactor c_i (cf.(6.19)) we simulated the system to obtain the expected shortfalls, which are independent from the hedging points (since we decomposed the system). In Table 7.1 we compare the analytical results with simulation results for Example 1. Results for Example 2 are in Table 7.2. In both tables, the first column list the desired service level requirements for stages 1 (final product), and the second column list the assumed service level requirements for stage 2. The third and fourth column list the analytically computed hedging points, for stages 1 and 2, respectively. We simulated the system with these hedging points. In both tables, we report in the 5th and 6th column the simulated value of the expected inventory at stages 1 and 2, respectively. Finally, in the last two columns we report the actual simulated stockout probability at each stage.

Analytical Results				Simulation Results			
ϵ_1	ϵ_2	w_1	w_2	$\mathbf{E}[(I^1)^+]$	$\mathbf{E}[I^2]$	$\mathbf{P}[I^1 \leq 0]$	$\mathbf{P}[I^2 = 0]$
0.20	10^{-2}	15.73	10.79	8.23	9.86	0.227	1.349×10^{-2}
0.15	10^{-2}	18.82	10.79	10.70	9.86	0.175	1.349×10^{-2}
0.10	10^{-2}	23.17	10.79	14.42	9.86	0.120	1.349×10^{-2}
0.05	10^{-2}	30.61	10.79	21.20	9.86	0.0633	1.349×10^{-2}
10^{-2}	10^{-2}	47.87	10.79	37.89	9.86	1.444×10^{-2}	1.349×10^{-2}
10^{-2}	10^{-3}	47.87	17.70	38.61	16.63	1.060×10^{-2}	1.084×10^{-3}
10^{-3}	10^{-3}	72.58	17.70	63.22	16.63	1.048×10^{-3}	1.084×10^{-3}
10^{-3}	10^{-4}	72.58	24.60	63.28	23.51	0.996×10^{-3}	0.959×10^{-4}
10^{-4}	10^{-4}	97.29	24.60	87.98	23.51	1.044×10^{-4}	0.974×10^{-4}
10^{-4}	10^{-5}	97.29	31.50	87.99	30.40	1.038×10^{-4}	0.936×10^{-5}
10^{-5}	10^{-5}	121.99	31.50	112.69	30.40	1.135×10^{-5}	0.936×10^{-5}
10^{-5}	10^{-6}	121.99	38.40	112.69	37.30	1.131×10^{-5}	1.046×10^{-6}

Table 7.1: Numerical results from the decomposition approach for Example 1. The simulated values for the expected shortfalls, that are used in computing the prefactors c_i , are $\mathbf{E}[L^1] = 9.297$ and $\mathbf{E}[L^2] = 1.098$. Expected inventory costs are reported in Table 7.4.

Analytical Results				Simulation Results			
ϵ_1	ϵ_2	w_1	w_2	$\mathbf{E}[(I^1)^+]$	$\mathbf{E}[I^2]$	$\mathbf{P}[I^1 \leq 0]$	$\mathbf{P}[I^2 = 0]$
0.20	10^{-2}	4.41	46.55	2.78	38.48	0.196	0.939×10^{-2}
0.15	10^{-2}	5.53	46.55	3.69	38.48	0.156	0.939×10^{-2}
0.10	10^{-2}	7.09	46.55	5.03	38.48	0.111	0.939×10^{-2}
0.05	10^{-2}	9.78	46.55	7.49	38.48	0.0612	0.939×10^{-2}
10^{-2}	10^{-2}	16.01	46.55	13.52	38.48	1.396×10^{-2}	0.939×10^{-2}
10^{-2}	10^{-3}	16.01	71.25	13.61	63.06	1.022×10^{-2}	0.983×10^{-3}
10^{-3}	10^{-3}	24.92	71.25	22.48	63.06	1.346×10^{-3}	0.983×10^{-3}
10^{-3}	10^{-4}	24.92	95.96	22.50	87.75	1.137×10^{-3}	1.256×10^{-4}
10^{-4}	10^{-4}	33.83	95.96	31.41	87.74	1.158×10^{-4}	0.966×10^{-4}
10^{-4}	10^{-5}	33.83	120.67	31.41	112.45	1.084×10^{-4}	0.903×10^{-5}
10^{-5}	10^{-5}	42.74	120.67	40.32	112.45	1.083×10^{-5}	0.903×10^{-5}
10^{-5}	10^{-6}	42.74	145.37	40.32	137.16	1.068×10^{-5}	0.864×10^{-6}

Table 7.2: Numerical results from the decomposition approach for Example 2. The simulated values for the expected shortfalls, that are used in computing the prefactors c_i , are $\mathbf{E}[L^1] = 2.420$ and $\mathbf{E}[L^2] = 8.216$. Expected inventory costs are reported in Table 7.6.

In most cases, we select the service level requirement of the second stage to be same as, or one order of magnitude less, than ϵ_1 . For those large stockout probabilities ($\epsilon_1 > 0.01$), we set $\epsilon_2 = 0.01$. The numerical results suggest that this suffices to make the decomposition approach valid. In particular, we observe that the proposed large deviations asymptotics are fairly accurate, they capture the exponent of the stockout probability and get fairly close in the first significant digit. Of course, there are many combinations of w_1 and w_2 that would lead to the same service level. Our decomposition approach yields one possible combination. In particular, the decomposition approach minimizes the required safety stock for stage 1, w_1 , since it assumes that no upstream material requirement constraints are in effect. In the next section we explore how we can select the best such combination to minimize expected inventory costs.

7.6.2 A Two-Stage System with Multi-Echelon Approach

Next we apply the multi-echelon approach to both Examples 1 and 2 considered above.

We start with Example 1. Using the results of Theorem 7.2.1, Corollary 7.2.6 and the

characterization of $\theta_{G,1}^*$ in Equation (7.21) we obtain

$$\theta_1^* = \theta_2^* = 0.0932,$$

and

$$\theta_{G,1}^* = \min\left(\theta_1^*, \frac{w_2}{w_1}\theta_2^*\right) = 0.0932,$$

for all $w_2 \geq w_1$. For echelon 2, we obtain $\theta_{G,2}^* = 0.2584$.

We select $w_1 = 10, 30, 50$, and for each w_1 , select w_2 such that $\beta_1 = \frac{w_2}{w_1} = 1, 1.5, 2, 2.5, 3.5$. We simulated the system with those sample points (w_1, w_2) and to construct the prefactor $f_1(w_1, \beta_1)$ for the stockout probability and the approximation $g_1(\Delta_1)$ for the expected shortfall. By simulation we also obtain $\mathbf{E}[Y^2] = 2.42$, and $c_2 = \theta_{G,2}^* \mathbf{E}[Y^2] = 0.63$, which will be used as a prefactor in the echelon-2 stockout probability (as in (6.20)).

We solved the nonlinear programming problem in (7.47) for a variety of service level requirements ϵ_1 (we impose no service level requirement on stage 2, i.e., $\epsilon_2 = 1$) and holding costs $\mathbf{h} = (h_1, h_2)$ for stages 1, 2, respectively. The results are reported in Table 7.3. There are two main observations we can make:

- Our analytical approximation for the stockout probability and the expected inventory cost is very accurate. To see that compare (i) the actual stockout probability (column 5 in the table) achieved by the optimal solution \mathbf{w}_A^* of the optimization problem in (7.47) with the corresponding service level requirement ϵ_1 , and (ii) the actual inventory cost of \mathbf{w}_A^* (column 6) with its analytical approximation (column 4). Our results are accurate even for relatively large stockout probabilities, that is, away from the large deviations limiting regime.
- The performance of our analytically obtained policy \mathbf{w}_A^* is rather close to the optimal policy (obtained by simulation). In fact, our policy is within at most 2% of the optimal (difference of column 6 and 9), which drops to at most 1% if we ignore the first row

ϵ_1	\mathbf{h}	Analytical Results		Simulated Values		Simulation Results		
		\mathbf{w}_A^*	$\mathbf{E}[C]$	$\mathbf{P}[X^1 \leq 0]$	$\mathbf{E}[C]$	\mathbf{w}_S^*	$\mathbf{P}[X^1 \leq 0]$	$\mathbf{E}[C]$
0.20	(1,1)	(16.57, 20.71)	28.6	0.208	29.2	(17,21)	0.199	29.8
0.15	(1,1)	(23.41, 23.41)	34.3	0.157	34.4	(20,24)	0.150	34.7
0.10	(1,1)	(27.76, 27.76)	41.9	0.103	42.0	(26,28)	0.099	42.3
0.05	(1,1)	(35.25, 35.25)	55.8	0.052	55.9	(32,36)	0.049	56.4
10^{-2}	(1,10)	(52.51, 52.51)	541.3	$1.01 \cdot 10^{-2}$	541.3	(50,53)	$0.99 \cdot 10^{-2}$	546.2
10^{-2}	(1,1)	(52.57, 52.57)	89.6	$1.01 \cdot 10^{-2}$	89.6	(50,53)	$0.99 \cdot 10^{-2}$	90.0
10^{-2}	(5,1)	(52.50, 52.50)	246.7	$1.01 \cdot 10^{-2}$	246.6	(50,53)	$0.99 \cdot 10^{-2}$	247.2
10^{-2}	(600,1)	(47.72, 95.44)	23218	$1.03 \cdot 10^{-2}$	23211	(48,57)	$1.00 \cdot 10^{-2}$	23257
10^{-3}	(1,1)	(77.21, 77.21)	138.7	$1.03 \cdot 10^{-3}$	138.7	(74,78)	$0.99 \cdot 10^{-3}$	139.4
10^{-4}	(1,1)	(101.93, 101.93)	188.1	$9.96 \cdot 10^{-5}$	188.1	(99,102)	$9.82 \cdot 10^{-5}$	187.8

Table 7.3: Numerical results for Example 1 operated under the multi-echelon policy. We denote by \mathbf{w}_A^* (3rd column) the hedging vector obtained by solving the optimization problem in (7.47). Similarly, \mathbf{w}_S^* (7th column) denotes the hedging vector obtained by brute-force simulation over integer points. The 4th column ($\mathbf{E}[C]$) lists the optimal value of the optimization problem in (7.47), that is, our analytical approximation of the total expected inventory cost of the policy in column 3. Column 5 and 6 list the stockout probability and expected inventory cost, respectively, obtained by simulating the policy of column 3. Column 8 and 9 list the stockout probability and expected inventory cost, respectively, obtained by simulating policy of column 7.

of the table. ¹

To assess the efficiency of the analytical approach, note that to optimize the expected inventory cost by simulation we need to simulate for all possible integer combinations of w_1 and w_2 and select the one that yields the lowest cost. Moreover, simulating small stockout probabilities requires very large sample sizes. It usually takes from several hours to several days to find the optimal by brute-force simulation, depending on the length of sample paths (as dictated by the service level requirements) and the number of (w_1, w_2) points. In fact, these running times of brute-force simulation were achieved by using information from our analytical results. Brute-force simulation with no information at all would take much longer and would be computationally intractable for the smaller ϵ_1 . The nonlinear programming problem can typically be solved within one minute (for instances we considered), while “pre-processing” (i.e., obtaining the prefactors) took on the order 30min. It is evident that the

¹In some cases in Table 7.3, \mathbf{w}_A^* achieves less inventory cost than \mathbf{w}_S^* because in these instances \mathbf{w}_A^* slightly violates the service level requirements (due to the large deviations approximation).

proposed analytical approach leads to huge computational savings at a modest performance cost. It should also be noted that in the simulations we only considered integer valued hedging points w_1 and w_2 . As a result, the discrepancies between analytical results and simulation in Table 7.3 contain this quantization error and thus overestimate the actual error of the analytical approach.

To compare the solution obtained by the multi-echelon approach with the one obtained by the decomposition approach we report the latter value in Table 7.4 (9th column). Because we ignore coupling among stages the decomposition approach is not as accurate in approximating $\mathbf{P}[I^1 \leq 0]$. Thus, in some cases it leads to solutions violate the service level requirement (more in row 1–4 of Table 7.4 and slightly in some of the remaining). In terms of inventory cost, the multi-echelon approach leads, in general, to more efficient solutions (except in rows 1–4 and 8 of Table 7.4 in which we end up with less inventory cost because we violate the service level requirement).

Analytical Results				Simulation Results					
ϵ_1	ϵ_2	w_1	w_2	$\mathbf{E}[(I^1)^+]$	$\mathbf{E}[I^2]$	$\mathbf{P}[I^1 \leq 0]$	h_1	h_2	$\mathbf{E}[\text{Cost}]$
0.2	10^{-2}	15.73	10.79	8.23	9.86	0.227	1	1	26.32
0.15	10^{-2}	18.82	10.79	10.70	9.86	0.175	1	1	31.26
0.1	10^{-2}	23.17	10.79	14.42	9.86	0.120	1	1	38.70
0.05	10^{-2}	30.21	10.79	21.20	9.86	0.063	1	1	52.26
10^{-2}	10^{-3}	47.87	17.70	38.61	16.63	1.060×10^{-2}	1	10	591.01
10^{-2}	10^{-3}	47.87	17.70	38.61	16.63	1.060×10^{-2}	1	1	93.85
10^{-2}	10^{-3}	47.87	17.70	38.61	16.63	1.060×10^{-2}	5	1	248.29
10^{-2}	10^{-3}	47.87	17.70	38.61	16.63	1.060×10^{-2}	600	1	23221.24
10^{-3}	10^{-3}	72.58	17.70	63.22	16.63	1.048×10^{-3}	1	1	143.07
10^{-3}	10^{-4}	72.58	24.60	63.28	23.51	0.996×10^{-3}	1	1	150.07
10^{-4}	10^{-4}	97.29	24.60	87.98	23.51	1.044×10^{-4}	1	1	199.47
10^{-4}	10^{-5}	97.29	31.50	87.99	30.40	1.038×10^{-4}	1	1	206.38

Table 7.4: The 3rd and 4th column report the hedging points obtained by the decomposition approach. We simulated the system to obtain $\mathbf{E}[(I^1)^+]$, $\mathbf{E}[I^2]$, $\mathbf{P}[I^1 \leq 0]$, and $\mathbf{E}[\text{Cost}]$. To make comparisons with the results in Table 7.3 note that holding costs h_1 , h_2 for echelon 1 and 2, respectively, correspond to holding costs $h_1 + h_2$ and h_2 for stage 1 and 2, respectively, in the decomposition approach.

Next we consider Example 2. We compute

$$\theta_1^* = 0.2584, \quad \theta_2^* = 0.0932,$$

and

$$\theta_{G,1}^* = \min \left(\theta_1^*, \frac{w_2}{w_1} \theta_2^* \right) = \min \left(0.2584, 0.0932 \frac{w_2}{w_1} \right).$$

For echelon 2, we obtain $\theta_{G,2}^* = 0.0932$.

As in Example 1, we simulated the system with sample points (w_1, w_2) , where $w_1 = 10, 30, 50$, for each w_1 , select w_2 such that $\beta_1 = \frac{w_2}{w_1} = 1, 2, 2.5, 2.7736, 3, 4$, and construct the prefactor $f_1(w_1, \beta_1)$ for the stockout probability and the approximation $g_1(\Delta_1)$ for the expected shortfall. By simulation we also obtain $\mathbf{E}[Y^2] = 9.28$, and $c_2 = \theta_{G,2}^* \mathbf{E}[Y^2] = 0.865$, which will be used as a prefactor in the stockout probability at stage 2.

Solving the optimization problem in (7.47) we obtain the results reported in Table 7.5. The results are similar in nature to the ones we obtained for Example 1. That is, our approximations are accurate for both stockout probabilities and expected inventory costs and the analytical solution is within 2.2% of the optimal. Figure 7.7 depicts how the expected inventory cost (obtained by simulation) changes with the hedging vector for the cases $\epsilon_1 = 10^{-2}$, $h_1 = 1$, $h_2 = 1$ and $\epsilon_1 = 10^{-2}$, $h_1 = 5$, $h_2 = 1$ (rows 5 and 6 in Table 7.5). It can be seen that the policy obtained by our analytical approach is very close to optimal; deviating from \mathbf{w}_A^* can lead to significantly larger expected inventory cost, which stresses the significance of optimization. Finally, as in Example 1, we compare the multi-echelon policy with the decomposition policy in Table 7.6. As expected the multi-echelon policy leads to more economic solutions.

7.6.3 A Three-Stage System with Multi-Echelon Approach

We next consider a three-stage system (Example 3), with the Markov-modulated demand and production processes depicted in Figure 7.8.

We apply the multi-echelon approach to Example 3. Using the results of Section 7.1 (cf.

ϵ_1	\mathbf{h}	Analytical Results		Simulated Values		Simulation Results		
		\mathbf{w}_A^*	$\mathbf{E}[C]$	$\mathbf{P}[X^1 \leq 0]$	$\mathbf{E}[C]$	\mathbf{w}_S^*	$\mathbf{P}[X^1 \leq 0]$	$\mathbf{E}[C]$
0.20	(1,1)	(8.82, 22.04)	20.9	0.198	20.1	(10,21)	0.200	19.8
0.15	(1,1)	(10.06, 25.19)	24.8	0.152	24.0	(10,25)	0.150	23.8
0.10	(1,1)	(11.62, 29.91)	30.7	0.102	29.9	(13,29)	0.099	30.1
0.05	(1,1)	(14.18, 38.05)	40.9	$5.06 \cdot 10^{-2}$	40.4	(15,37)	0.05	40.0
10^{-2}	(1,1)	(21.74, 54.61)	64.7	$1.03 \cdot 10^{-2}$	64.4	(22,54)	$1.00 \cdot 10^{-2}$	64.0
10^{-2}	(5,1)	(17.75, 62.15)	129.7	$1.05 \cdot 10^{-2}$	129.1	(18,61)	$1.00 \cdot 10^{-2}$	129.1
10^{-2}	(1,10)	(26.92, 52.81)	460.3	$1.03 \cdot 10^{-2}$	459.9	(22,54)	$1.00 \cdot 10^{-2}$	467.2
10^{-3}	(1,1)	(29.26, 81.17)	98.7	$1.07 \cdot 10^{-3}$	98.7	(30,81)	$9.95 \cdot 10^{-4}$	99.2
10^{-4}	(1,1)	(38.21, 106.00)	132.5	$1.22 \cdot 10^{-4}$	132.5	(40,107)	$1.00 \cdot 10^{-4}$	135.5

Table 7.5: Numerical results for Example 2 operated under the multi-echelon policy. The notation and the structure of the table are the same as in Table 7.3.

Analytical Results				Simulation Results					
ϵ_1	ϵ_2	w_1	w_2	$\mathbf{E}[(I^1)^+]$	$\mathbf{E}[I^2]$	$\mathbf{P}[I^1 \leq 0]$	h_1	h_2	$\mathbf{E}[\text{Cost}]$
0.2	10^{-2}	4.41	46.55	2.78	38.48	0.196	1	1	44.04
0.15	10^{-2}	5.53	46.55	3.69	38.48	0.156	1	1	45.86
0.1	10^{-2}	7.09	46.55	5.03	38.48	0.111	1	1	48.54
0.05	10^{-2}	9.78	46.55	7.49	38.48	0.061	1	1	53.46
10^{-2}	10^{-3}	16.01	71.25	13.61	63.06	1.022×10^{-2}	1	1	90.28
10^{-2}	10^{-3}	16.01	71.25	13.61	63.06	1.022×10^{-2}	5	1	144.72
10^{-2}	10^{-3}	16.01	71.25	13.61	63.06	1.022×10^{-2}	1	10	780.31
10^{-3}	10^{-3}	24.92	71.25	22.48	63.06	1.346×10^{-3}	1	1	108.02
10^{-3}	10^{-4}	24.92	95.96	22.50	87.75	1.137×10^{-3}	1	1	132.75
10^{-4}	10^{-4}	33.83	95.96	31.41	87.74	1.158×10^{-4}	1	1	150.56
10^{-4}	10^{-5}	33.83	120.67	31.41	112.45	1.084×10^{-4}	1	1	175.27

Table 7.6: Comparing the multi-echelon policy with the decomposition policy in Example 2.

(7.21), (7.22)), we obtain

$$\theta_1^* = 0.1276, \theta_2^* = 0.08663, \theta_3^* = 0.08663,$$

and

$$\theta_{G,1}^* = \min \left(\theta_1^*, \frac{w_2}{w_1} \theta_2^*, \frac{w_3}{w_1} \theta_3^* \right) = \min \left(\theta_1^*, \frac{w_2}{w_1} \theta_2^* \right).$$

Therefore, the feasible set of $\beta = (\beta_1, \beta_2) = \left(\frac{w_2}{w_1}, \frac{w_3}{w_1} \right)$ has two regions determined by $\beta_1 : 1 \leq \beta_1 < \frac{\theta_1^*}{\theta_2^*}$ and $\beta_1 \geq \frac{\theta_1^*}{\theta_2^*}$. We select $w_1 = 20, 40, 60, 80$, and for each w_1 , select a set of (w_2, w_3) to include some sample points on the boundary and inside the two regions




Figure 7.7: The optimal multi-echelon policies for Example 2 derived by simulation, where the service level requirement is $\epsilon_1 = 10^{-2}$. (a) When $h_1 = 1$ and $h_2 = 1$, $w_1^* = 22$, $w_2^* = 54$, $E[\text{Cost}] = 64.04$; (b) when $h_1 = 5$ and $h_2 = 1$, $w_1^* = 18$, $w_2^* = 61$, $E[\text{Cost}] = 129.22$. In each graph, the thick curve is the boundary of the set of the feasible policies (i.e., set of vectors (w_1, w_2) satisfying the service level constraints), the optimal policy obtained by our analytical approach is marked with a circle.




Figure 7.8: The models of demand and production processes in Example 3.

of β . We simulate the system at those sample points and use the approach of Section 7.3 to obtain $f_1(w_1, \beta)$ and $g_i(\cdot)$, $i = 1, 2, 3$, and, as a result, the stockout probability and the expected inventory cost.

We solve the nonlinear programming problem in (7.47) for a variety of service level requirements ϵ_1 and holding costs $\mathbf{h} = (h_1, h_2, h_3)$. The results are reported in Table 7.7. Again, we observe that the analytical results are very close to the ones obtained by simulation. That is, our approximations are accurate for both stockout probabilities and expected inventory costs and the analytical solution is within at most 2.5% of the optimal.

ϵ_1	\mathbf{h}	Analytical Results		Simulated Values		Simulation Results		
		\mathbf{w}_A^*	$\mathbf{E}[C]$	$\mathbf{P}[X^1 \leq 0]$	$\mathbf{E}[C]$	\mathbf{w}_S^*	$\mathbf{P}[X^1 \leq 0]$	$\mathbf{E}[C]$
10^{-1}	(1,1,1)	(27.5, 53.7, 68.7)	129.8	1.14×10^{-1}	132.6	(23,51,74)	0.98×10^{-1}	136.0
7×10^{-2}	(1,1,1)	(29.3, 59.0, 74.0)	142.8	7.55×10^{-2}	145.0	(23, 56, 80)	6.82×10^{-2}	147.6
3×10^{-2}	(1,1,1)	(35.2, 68.1, 88.1)	175.1	3.36×10^{-2}	176.4	(28, 64, 92)	2.97×10^{-2}	173.4
10^{-2}	(1,1,1)	(42.5, 83.7, 103.7)	215.6	1.15×10^{-2}	215.8	(34,88,104)	1.01×10^{-2}	211.5
10^{-2}	(1,3,1)	(43.9, 78.5, 108.5)	360.2	1.20×10^{-2}	362.7	(34,85,106)	1.01×10^{-2}	364.1
10^{-3}	(1,1,1)	(59.7, 116.2, 130.0)	289.9	1.08×10^{-3}	290.1	(57,116,130)	1.01×10^{-3}	287.5
10^{-4}	(1,1,1)	(75.9, 134.4, 162.6)	362.3	0.97×10^{-4}	363.5	(69,135,166)	1.00×10^{-4}	360.6

Table 7.7: Numerical results for Example 3 operated under the multi-echelon policy. The notation and the structure of the table are the same as in Table 7.3.

7.6.4 Significance of Distributional Information

As our final example we present a two-stage supply chain model operated under the multi-echelon inventory policy. We will demonstrate that distributional information on the demand and service processes is critical in making inventory control decisions. In particular, we will try to identify the bottleneck stage that determines the stockout probability at stage 1.

The demand and production processes are all discrete-time Markov modulated processes. Letting \mathbf{P} and \mathbf{r} denote the transition probabilities and the vector of demand or production amounts in each state of the corresponding Markov chain we set:

$$\mathbf{r}_D = (5, 10), \quad \mathbf{P}_D = \begin{bmatrix} 0.2 & 0.8 \\ 0.4 & 0.6 \end{bmatrix}, \quad \mathbf{E}[D] = 8.33,$$

$$\mathbf{r}_{B^1} = (0, 25), \quad \mathbf{P}_{B^1} = \begin{bmatrix} 0.2 & 0.8 \\ 0.3 & 0.7 \end{bmatrix}, \quad \mathbf{E}[B^1] = 18.18,$$

and

$$\mathbf{r}_{B^2} = (0, 14), \quad \mathbf{P}_{B^2} = \begin{bmatrix} 0.15 & 0.85 \\ 0.05 & 0.95 \end{bmatrix}, \quad \mathbf{E}[B^2] = 13.22.$$

Applying the results of Section 7.1 (cf. (7.21), (7.22)) we obtain

$$\theta_1^* = 0.1785, \quad \theta_2^* = 0.1785.$$

Figure 7.9 depicts θ_1^* and θ_2^* as roots of the corresponding nonlinear equation.




Figure 7.9: For the example of Section 7.6.4, we plot $\Lambda_D(\theta) + \Lambda_{B^1}(-\theta)$, $\Lambda_D(\theta) + \Lambda_{B^2}(-\theta)$ and identify the corresponding largest positive roots θ_1^* and θ_2^* , respectively.

We compute

$$\theta_{G,1}^* = \min(\theta_1^*, \frac{w_2}{w_1}\theta_2^*) = \theta_1^* = 0.1785,$$

which according to the discussion in Remark 2 of Section 7.1 implies that the “bottleneck” is stage 1 in the sense that the process B^1 and not B^2 characterizes the stockout probability at stage 1. This seems to contradict naive intuition that the “bottleneck” is stage 2 since $\mathbf{E}[B^1] > \mathbf{E}[B^2]$! The conclusion that the “bottleneck” is stage 1 is explained by noting that B^1 is more bursty than B^2 .

Chapter 8

Summary and Future Research

We considered pricing and resource allocation decisions in communication and supply networks with *Quality of Service (QoS)* considerations. In particular, we focused on revenue or welfare maximization communication networks, and inventory control in supply chains subject to given QoS requirements. Complicated stochastic processes and QoS metrics make it impossible to analyze those stochastic networks exactly. Various asymptotic results were obtained in the work of this thesis.

8.1 Pricing in Multiservice Communication Networks

On pricing, we started with the problem of pricing single resource multi-service communication systems, and formulated the revenue and welfare maximization problem as a dynamic programming problem. We explored the properties of the optimal dynamic policy, and established some insightful, qualitative properties. We developed several approximation approaches, including price aggregation and approximated dynamic programming, that can handle large scale problems. We derived an upper bound for the optimal revenue, which can be used to evaluate the performance of our suboptimal policies when it is practically impossible to obtain the optimal policy. We also proposed another suboptimal policy, the static pricing policy, and we provided an algorithm for calculating the blocking probabilities for the static policy. Numerical results show that the expected revenue of those suboptimal policies are very close to the optimal revenue.

Then, we considered a general, network model and studied the problem of pricing the

use of the available resources under both revenue and welfare maximization objectives. We established that static pricing is asymptotically optimal in a regime of many small users. To that end, we established that in this limit the blocking probabilities under an appropriate static pricing policy converge to zero exponentially fast. We characterized this exponential rate of convergence, which allowed us to obtain simple estimates on the size of the network in which static pricing is within a given distance from the optimal. We characterized the structure of asymptotically optimal static prices and used this structure to obtain near-optimal policies away from the limiting regime. To that end, we employed a simulation-based optimization method that optimizes policy parameters by obtaining gradient information throughout the course of a simulation of the system. Our approach can handle large, realistic, problem sizes.

In practice, where demand is nonstationary but slowly varying, the proposed policy leads to *time-of-day* pricing. There is substantial accumulated experience with such policies in the telecommunications industry, which facilitates their actual implementation. A practical implementation would also need to be coupled with a demand estimation mechanism (in fact, only demand elasticity information is needed). The proposed simulation-based optimization approach can be driven by the actual operation of the network, instead of a simulation. In this setting, a demand estimation mechanism can be naturally be incorporated.

We also studied revenue and welfare maximization problems for networks with demand substitution effects. In particular, increasing the price of a class decreases its demand but may boost demand for other classes. We established the asymptotic optimality of the properly chosen static policy and characterized the structure of the asymptotically optimal static policy. A numerical example was presented to show how we can use this structure and a simulation-based optimization method to obtain a good static pricing policy.

Some interesting problems that we will be focus on in the future include the pricing of networks that use dynamic (instead of fixed) routing, both with and without demand substitution effects. We also plan to consider models of demand with more complicated stochastic characteristics.

8.2 Inventory Control in Supply Chains

We have developed two production policies for inventory control in a multi-stage single-class supply chain. Demand and service processes are general, potentially autocorrelated processes, which makes it possible to model complex demand scenarios and failure-prone production facilities. Both policies emphasize quality of service, which is becoming important in modern manufacturing, by maintaining desirable service level constraints. The first policy is a base-stock policy that uses only local inventory information. The second policy is an echelon-base stock policy. In both cases we relied upon large deviations techniques for analysis. This led to asymptotically tight approximations for the stockout probabilities which allows us to analytically obtain appropriate hedging points that maintain the desirable service level constraints.

Our analysis under the echelon-base stock policy provides particular insight on how stockouts occur. In particular, it identifies a “bottleneck” stage whose production capacity is “responsible” for stockouts at stage 1. But, this “bottleneck” stage is not necessarily the one with the smallest mean production capacity; it depends on the full distribution of the production processes. We provided a simple numerical example to underline this observation, which at first sight might appear counter-intuitive.

The echelon base-stock policy enables optimization among all possible hedging point vectors that satisfy the service level constraints; by solving a nonlinear optimization problem we select the one with minimum expected inventory cost. Numerical results show that the solutions obtained by analysis are very close to the ones obtained by brute force simulation. Our analytic approach for selecting appropriate hedging points leads to dramatic computational savings when compared to the time needed to obtain them by simulation.

As a simple extension to the multi-stage multi-class system, we proposed an approach to decompose the system into several multi-stage single-class systems using a processor sharing policy. Our analysis, though, did not fully take into account the fact that capacity can be reallocated from idle classes to busy ones. Further work will seek to obtain tighter large

deviations results for multi-stage systems, as for example in [BPT99]. It is also of interest to consider alternative scheduling policies.

An additional future research direction is to analyze assembly and distribution systems using large deviation techniques. For an assembly-to-order system, our result on single-stage system can be applied, and an assembly-to-stock system has similarities to the multi-stage systems we studied. For distribution systems, if we only consider the supplier side, the model is the same as for the single-stage system with aggregated demand from all retailers. Cases where retailers and suppliers are more tightly integrated needs more careful investigation.

Bibliography

- [AK86] Ramakrishna Akella and P. R. Kumar, *Optimal control of production rate in a failure prone manufacturing system*, IEEE Transactions on Automatic Control **AC-31** (1986), no. 2, 116–126.
- [Ash97] Mike Ashworth, *Revenue management*, Electric Perspectives **22** (1997), no. 6, 36.
- [Asm87] Søren Asmussen, *Applied probability and queues*, Wiley, New York, 1987.
- [Ber95] Dimitri P. Bertsekas, *Dynamic programming and stochastic control*, vol. I and II, Athena Scientific, Belmont, MA, 1995.
- [BL92] F. Baccelli and Z. Liu, *On a class of stochastic recursive sequences arising in queueing theory*, Annals of Probability **20** (1992), 350–374.
- [BM95] Gabriel R. Bitran and Susana V. Mondschein, *An application of yield management to the hotel industry considering multiple day stays*, Operations Research **43** (1995), no. 3, 427–444.
- [BP01] Dimitris Bertsimas and Ioannis Ch. Paschalidis, *Probabilistic service level guarantees in make-to-stock manufacturing systems*, Operations Research **49** (2001), no. 1, 119–133.
- [BPT98a] Dimitris Bertsimas, Ioannis Ch. Paschalidis, and John N. Tsitsiklis, *Asymptotic buffer overflow probabilities in multiclass multiplexers: An optimal control approach*, IEEE Transactions on Automatic Control **43** (1998), no. 3, 315–335.
- [BPT98b] ———, *On the large deviations behaviour of acyclic networks of $G/G/1$ queues*, The Annals of Applied Probability **8** (1998), no. 4, 1027–1069.
- [BPT99] ———, *Large deviations analysis of the generalized processor sharing policy*, Queueing Systems **32** (1999), 319–349.
- [Buc90] James A. Bucklew, *Large deviation techniques in decision, simulation, and estimation*, Wiley, New York, 1990.
- [Cas93] Christos G. Cassandras, *Discrete event systems : Modeling and performance analysis*, Irwin, 1993.

- [CG95] William J. Carroll and Richard C. Grimes, *Evolutionary change in product management: Experiences in the car rental industry*, *Interfaces* **25** (1995), no. 5, 84–105.
- [Cha95] C.S. Chang, *Sample path large deviations and intree networks*, *Queueing Systems* **20** (1995), 7–36.
- [CKW97] Costas Courcoubetis, Frank P. Kelly, and Richard Weber, *Measurement-based usage charges in communications networks*, Research report 1997–19, Statistical Laboratory, Cambridge University, 1997.
- [Cla97] David D. Clark, *An Internet model for cost allocation and pricing*, *Internet Economics* (L. McKnight and J. Bailey, eds.), MIT Press, 1997.
- [CR99] Costas A. Courcoubetis and Martin I. Reiman, *Pricing in a large single link loss system*, *Teletraffic Engineering in a Competitive World* (F.P. Key and D.G. Smith, eds.), Elsevier, June 1999, pp. 737–746.
- [Cra38] H. Cramér, *Sûr un nouveau théorème-limite de la théorie des probabilités*, *Actualités Scientifiques et Industrielles, Colloque consacré à la théorie des probabilités*, no. 736, Hermann, Paris, 1938, pp. 5–23.
- [CS60] Andrew J. Clark and Herbert Scarf, *Optimal policies for a multi-echelon inventory problem*, *Management Science* **6** (1960), 475–490.
- [CS01] F. Chen and J.-S. Song, *Optimal policies for multiechelon inventory problems with Markov-modulated demand*, *Operations Research* **49** (2001), no. 2, 226–234.
- [DE97] Paul Dupuis and Richard S. Ellis, *A weak convergence approach to the theory of large deviations*, *Wiley Series in Probability and Statistics*, Wiley, New York, 1997.
- [dVKD00] F. de Véricourt, F. Karaesmen, and Y. Dallery, *Dynamic scheduling in a make-to-stock system: A partial characterization of optimal policies*, *Operations Research* **48** (2000), no. 5, 811–819.
- [DZ95] A. Dembo and T. Zajic, *Large deviations: From empirical mean and measure to partial sums processes*, *Stochastic Processes and their Applications* **57** (1995), no. 2, 191–224.
- [DZ98] Amir Dembo and Ofer Zeitouni, *Large deviations techniques and applications*, 2nd ed., Springer, New York, 1998.
- [eco98] *A survey of manufacturing: Meet the global factory*, *The Economist* (1998), June 20th–26th.
- [Eva67] R. Evans, *Inventory control of a multiproduct system with a limited production resource*, *Naval Research Logistics* (1967), 173–184.

- [FH97] Michael Fu and Jian-Qiang Hu, *Conditional Monte Carlo: Gradient estimation and optimization applications*, Kluwer Academic Publishers, Boston, 1997.
- [FZ84] A. Federgruen and P. Zipkin, *Computational issues in an infinite horizon multi-echelon inventory model*, *Operations Research* **32** (1984), 818–836.
- [FZ86] ———, *An inventory model with limited production capacity and uncertain demands, I: The average cost criterion*, *Mathematics of Operations Research* **11** (1986), 193–207.
- [GG80] B. Gavish and S. Graves, *A one-product production/inventory problem under continuous review policy*, *Operations Research* **28** (1980), 1228–1236.
- [GK99] R.J. Gibbens and Frank P. Kelly, *Resource pricing and the evolution of congestion control*, *Automatica* **35** (1999), 1969–1985.
- [Gla96] Paul Glasserman, *Allocating production capacity among multiple products*, *Operations Research* **44** (1996), no. 5, 724–734.
- [Gla97] ———, *Bounds and asymptotics for planning critical safety stocks*, *Operations Research* **45** (1997), no. 2, 244–257.
- [GT94] Paul Glasserman and Sridhar Tayur, *The stability of a capacitated multi-echelon production-inventory system under a base-stock policy*, *Operations Research* **42** (1994), no. 5, 913–925.
- [GT95] ———, *Sensitivity analysis for base-stock levels in multiechelon production-inventory systems*, *Management Science* **41** (1995), no. 2, 263–281.
- [GvR97] Guillermo Gallego and Garrett van Ryzin, *A multiproduct dynamic pricing problem and its applications to network yield management*, *Operations Research* **45** (1997), no. 1, 24–42.
- [Ha97] A.Y. Ha, *Optimal dynamic scheduling policy for make-to-stock production system*, *Operations Research* **45** (1997), no. 1, 42–53.
- [IS01] Garud Iyengar and Karl Sigman, *Admission control policies for feed-forward networks of queues with no waiting*, INFORMS Applied Probability Society Conference (New York, NY), July 2001.
- [Kel91] Frank P. Kelly, *Loss networks*, *The Annals of Applied Probability* **1** (1991), no. 3, 319–378.
- [Kel94] ———, *On tariffs, policing and admission control for multiservice networks*, *Operations Research Letters* **15** (1994), 1–9.
- [Kel96] ———, *Notes on effective bandwidths*, *Stochastic Networks: Theory and Applications* (S. Zachary, I.B. Ziedins, and F.P. Kelly, eds.), vol. 9, Oxford University Press, 1996, pp. 141–168.

- [Kel97] ———, *Charging and rate control for elastic traffic*, *European Transactions on Telecommunications* **8** (1997), 33–37.
- [Key90] P.B. Key, *Optimal control and trunk reservation in loss networks*, *Probability in the Engineering and Information Sciences* **4** (1990), 203–242.
- [Kim89a] Sheryl E. Kimes, *The basics of yield management*, *Cornell Hotel and Restaurant Administration Quarterly* **30** (1989), no. 3, 14–20.
- [Kim89b] ———, *Yield management: A tool for capacity-constrained service firms*, *Journal of Operations Management* **8** (1989), no. 4, 348–364.
- [KMT98] Frank P. Kelly, A.K. Maulloo, and D.K.H. Tan, *Rate control for communication networks: Shadow prices, proportional fairness and stability*, *Journal of the Operational Research Society* **49** (1998), 237–252.
- [KS00] Srisankar Kunniyur and R. Srikant, *End-to-end congestion control schemes: Utility functions, random losses and ECN marks*, *Proceedings of the INFOCOM Conference (Tel Aviv, Israel)*, 2000.
- [KT98] R. Kapuscinski and S.R. Tayur, *A capacitated production-inventory model with periodic demand*, *Operations Research* **46** (1998), no. 6, 899–911.
- [KT99] ———, *Optimal policies and simulation based optimization for capacitated production inventory systems*, *Quantitative Models for Supply Chain Management* (S.R. Tayur, R. Ganeshan, and M. Magazine, eds.), Kluwer, 1999, pp. 7–40.
- [LA00] R.J. La and V. Anantharam, *Charge-sensitive TCP and rate control in the Internet*, 2000, Preprint, Available at <http://divine.eecs.berkeley.edu/ananth>.
- [Lou90] G.M. Louth, *Stochastic networks: Complexity, dependence and routing*, Ph.D. thesis, University of Cambridge, Cambridge, UK, 1990.
- [Low00] S.H. Low, *Equilibrium bandwidth and buffer allocation for elastic traffics*, *IEEE/ACM Transactions on Networking* **8** (2000), no. 3, 373–383.
- [LV93] Steven H. Low and Pravin P. Varaiya, *A new approach to service provisioning in ATM networks*, *IEEE/ACM Transactions on Networking* **1** (1993), no. 5, 547–553.
- [MB98] Lee W. McKnight and Joseph P. Bailey, *Internet economics*, MIT Press, Cambridge, MA, 1998.
- [MMV94] J.K. MacKie-Mason and H.R. Varian, *Pricing the Internet*, *Public Access to the Internet* (B. Kahin and J. Keller, eds.), Prentice-Hall, 1994.
- [MT01] Peter Marbach and John N. Tsitsiklis, *Simulation-based optimization of Markov reward processes*, *IEEE Transactions on Automatic Control* **46** (2001), no. 2, 191–209.

- [OK92] T.J. Ott and K.R. Krishnan, *Separable routing: A scheme for state-dependent routing of circuit switched telephone networks*, *Annals of Operations Research* **35** (1992), 43–68.
- [Pas99] Ioannis Ch. Paschalidis, *Class-specific quality of service guarantees in multimedia communication networks*, *Automatica*, (Special Issue on Control Methods for Communication Networks) **35** (1999), no. 12, 1951–1968, Anantharam and Walrand Eds.
- [PLCZ01] Ioannis Ch. Paschalidis, Yong Liu, Christos G. Cassandras, and Ping Zhang, *Threshold-based control for make-to-stock models: A synergy between large deviations and perturbation analysis*, *Proceedings of IEEE Conference on Decision and Control* (Orlando, FL), vol. 40, December 2001.
- [PPZ97] A. Peña-Perez and P. Zipkin, *Dynamic scheduling rules for a multiproduct make-to-stock queue*, *Operations Research* **45** (1997), no. 6, 919–930.
- [PT00] Ioannis Ch. Paschalidis and John N. Tsitsiklis, *Congestion-dependent pricing of network services*, *IEEE/ACM Transactions on Networking* **8** (2000), no. 2, 171–184.
- [Ros95] Keith W. Ross, *Multiservice loss models for broadband telecommunication networks*, Springer-Verlag, 1995.
- [SLD92] B.C. Smith, J.F. Leimkuhler, and R.M. Darrow, *Yield management at American Airlines*, *Interfaces* **22** (1992), no. 1, 8–31.
- [Sob82] M. Sobel, *The optimality of full-service policies*, *IEEE Transactions on Automatic Control* **30** (1982), 636–649.
- [SSL97] J. Subramanian, S. Stidham, and C.J. Lautenbacher, *Airline yield management with overbooking, cancellations, and no-shows*, 1997, Preprint.
- [SW95] Adam Shwartz and Alan Weiss, *Large deviations for performance analysis: Queues, communications, and computing*, Chapman & Hall, 1995.
- [VW96] M.H. Veatch and L.M. Wein, *Scheduling a make-to-stock queue: Index policies and hedging points*, *Operations Research* **44** (1996), no. 4, 634–647.
- [Wei92] L.M. Wein, *Dynamic scheduling of a multiclass make-to-stock queue*, *Operations Research* **40** (1992), 724–735.
- [ZZ90] Y.-S. Zheng and P. Zipkin, *A queueing model to analyze the value of centralized inventory information*, *Operations Research* **38** (1990), no. 2, 296–307.

Curriculum Vitae

Yong Liu

Department of Manufacturing Engineering
Boston University
Boston, MA 02215
(617) 353-9586
liuyong@bu.edu

- Education**
- BOSTON UNIVERSITY Boston, MA
09/1997–present, Department of Manufacturing Engineering
Ph.D. candidate
- XI'AN JIAOTONG UNIVERSITY Xi'an, CHINA
09/1994–06/1997, Systems Engineering Institute
M.E. in Systems Engineering
- XI'AN JIAOTONG UNIVERSITY Xi'an, CHINA
09/1990–07/1994, Department of Information and Control Engineering
B.E. in Information Engineering

- Publications**
1. I.Ch. Paschalidis and Y. Liu, “Distributed Resource Allocation in Multiservice Communication Networks Using Pricing”, submitted to the *41st IEEE Conference on Decision and Control*, 2002.
 2. I.Ch. Paschalidis, Y. Liu, C.G. Cassandras, and C. Panayiotou, “Inventory Control for Supply Chains with Service Level Constraints: A Synergy between Large Deviations and Perturbation Analysis”, Submitted to the *Annals of Operations Research: Special Volume on Stochastic Models of Production-Inventory Systems*.
 3. I.Ch. Paschalidis and Y. Liu, “Pricing in Multiservice Loss Networks: Static Pricing, Asymptotic Optimality, and Demand Substitution Effects”, To appear in *IEEE/ACM Transactions on Networking*.
 4. I.Ch. Paschalidis, Y. Liu, C.G. Cassandras, and P. Zhang “Threshold-Based Control for Make-to-Stock Models: A Synergy between Large Deviations and Perturbation Analysis”, *Proceedings of 40th IEEE Conference on Decision and Control*, December 2001, Orlando, Florida.
 5. I.Ch. Paschalidis and Y. Liu, “Revenue Management in Multi-Service Communication Networks”, *INFORMS Annual Conference*, November 2001, Miami Beach, FL.
 6. I.Ch. Paschalidis and Y. Liu, “Static Pricing in Multi-Service Communication Networks”, *The 11th INFORMS Applied Probability Society Conference*, July 2001, New York, NY.

7. I.Ch. Paschalidis and Y. Liu, "Large Deviations-Based Asymptotics for Inventory Control in Supply Chains", Accepted by *Operations Research*.
8. I.Ch. Paschalidis and Y. Liu, "Large Deviations-Based Asymptotics for Inventory Control in Supply Chains", *Proceedings of 39th IEEE Conference on Decision and Control*, pages 528–533, December 2000, Sydney, Australia.
9. I.Ch. Paschalidis and Y. Liu, "Inventory Control in Supply Chains: A Large Deviations Approach", *INFORMS Annual Conference*, November 2000, San Antonio, TX.

Awards

- Boston University Dean's Fellowship, 1997
- Xi'an Jiaotong University Scholarships, 1991–1994
- Honored Student of Xi'an Jiaotong University, 1991, 1992, 1994
- First Prize, Mathematics Contest of Xi'an Jiaotong University, 1991
- First Prize, Computer Contest of Xi'an Jiaotong University, 1993

Background

Date of Birth: January 12, 1973

Place of Birth: Xianyang, Shaanxi Province, China

Languages: English, Chinese

Affiliations: IEEE Student Member, INFORMS Student Member