THREE ESSAYS ON RESOURCE ALLOCATION: LOAD BALANCING ON HIGHLY VARIABLE SERVICE TIME NETWORKS, MANAGING DEFAULT RISK VIA SUBSIDIES AND SUPPLIER DIVERSIFICATION, AND OPTIMAL HOTEL ROOM ASSIGNMENT

by

Luz Adriana E. Caudillo Fuentes

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Industrial and Operations Engineering)
in The University of Michigan
2010

Doctoral Committee:

Assistant Professor Volodymyr Babich, Co-Chair
Professor Mark Van Oyen, Co-Chair
Professor Emeritus Stephen Pollock
Assistant Professor Damian Beil

Ad maiorem Dei gloriam.

Entonces Dios dijo: "Hágase la luz". Y la luz se hizo.

Genesis, 1, 3

"Es muy necesario que tú, personalmente vayas, ruegues, que por tu intercesión se realice, se lleve a efecto mi querer, mi voluntad."

"Señora mía, Niña, ya voy a realizar tu venerable aliento, tu venerable palabra."

Nican Mopohua

A Gustavo y Luz María.

ACKNOWLEDGEMENTS

When the moment comes to say "thank you," the mission looks next to impossible, given the non-negligible and almost inevitable risk of not naming everyone who in some way helped me accomplish the task. Yet, the task needs to be done, and to whoever reads this, I know I am indebted to you in some way, for simply browsing these pages looking for some reference and making this work worthwhile, or for having helped in major ways to complete it, or for having supported me during the challenging doctoral program.

First of all, my greatest admiration and gratitude to my parents Gustavo Caudillo Márquez and Luz María Fuentes Catemaxque, the very reason for and joy of my life and my inspiration forever: Without the example of exemplary persons who grew up in adversity never to be known in first-world countries, I would never have found enough strength to accomplish this large and complex project. A few lines to them, who, besides Spanish, speak the language of honesty and love. "Gracias papi, Gracias mami, por la vida y por darme la fuerza necesaria para este proyecto. Es por ustedes y por Dios que seguí adelante." May these sheets be a humble homage of a daughter who will never be able to say thank you in the way they deserve.

I would like to express my admiration and my most sincere appreciation to my advisor and mentor, Professor Volodymyr Babich, to whom I not only owe academic training and learning, but finding confidence in my work, the key element to make any task possible. I deeply appreciate his respectful encouragement and equally respectful challenging questions. His expertise and his natural talent to teach how to model and extract useful insights were not only essential in the completion of this work, but skills that I hope to cultivate and share in the future with others.

I am grateful to Professor Mark Lewis who smoothly took me through the comprehensive exams stage of the doctoral program. He also was my advisor while I worked on the paper on load balancing in the presence of heavy-tailed distribution services. For this work, no less important were the contributions of David Kaufman. I would like to thank Douglas Down from McMaster University and anonymous referees for their helpful comments on this work. In addition, I would like to thank the Mexican Council for Science and Technology (Consejo Nacional de Ciencia y Tecnología) for the funding I received while working on this topic. Moreover, this work was supported by the National Science Foundation under Grant Nos. CMMI-0540808 and CMMI-0826255. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

I would like to thank Yihua Li for sharing his knowledge on deterministic optimization. It has been exciting to be involved in a large, real problem, with all the inherent challenges. It has been exciting for me to contribute by helping to obtain a consistent mathematical formulation, to help in the search of solution methods, and to assist in the testing for the Room Assignment project. Having the opportunity to work with him in this and other problems has provided me experience with an optimization area that I have always found interesting, and that was one of my motivations to study Operations Research.

I am very thankful to Professor Mark Van Oyen, who accepted the role as my co-chair. Not only did he help by accepting the position, but he did actively guide me in the complex process of getting ready for the oral defense. This challenging stage would not have been possible without his feedback.

I also sincerely appreciate the time and help from each member of the Doctoral

Committee. Professor Stephen Pollock guided me through the graduate program. Professor Damian Beil accepted the challenge of assessing my research within a short period of time, and provided valuable feedback.

I am indebted to the Department of Industrial and Operations Engineering (IOE), especially to Professor Larry Seiford for his support in helping me find opportunities for both funding and learning. I thank all the professors from the department, especially those who in some way assessed me or instructed me.

I would also like to thank José Carlos Zavala Jurado, who kindly put time and effort to help me put together the different parts of my dissertation into a unified LaTeX document. No less valuable was the LaTeX help from Mark Paterson. Without his expertise, the final stage of the formatting of the dissertation would have not concluded satisfactorily nor in a timely manner. I am sincerely indebted to him and his family for their unmatched generosity and willingness to help.

Felipe Contreras Alcalá shared valuable programming knowledge in my first years in Michigan and put time and effort to assist me with the LaTeX formatting as well.

Gail Goretsky, besides providing much needed support, gave me great editorial feedback for Chapter 3.

I appreciate the support in 2008 from the community of Phinishers.org, who where kind enough to answer many TeX questions and, no less importantly, to share the road.

During this work, many challenges occurred, both academic and personal. To all my classmates in IOE, and my classmates in Financial Engineering, my sincere thanks, especially to those with whom I had the joy of sharing homework, meals, and life. My acknowledgement to my fellow Graduate Student Instructors (GSI) and Graduate Student Mentors (GSM) in IOE and the Department of Statistics.

To all my friends, I cannot thank you enough. My friends and roommates, Nancy Koch and Karina Robles, helped me survive the times when challenges were great. I cannot thank enough the many friends from St. Mary Student Parish who supported me, especially those from G.R.A.D. and from the Spanish Mass. In Florida, with no less enthusiasm, I would like to thank my friends from St. John Vianney Catholic Church, particularly the people from the Hispanic ministries and my friends from Annunciation Parish for their support in the last and challenging stages. Thank you for all the support from my friends, Anthony Musci, Claudia Torres Garibay, Sarah Heath, Guadalupe Patlán Jiménez, Monica Franco Sánchez, and Sivia Gordillo Barrios. I want to thank too Cecilia Hernández-Winchell and her family, David Medina and his family, Ms. Carmen Abreu and Nahum Cano and their family, Flora Trejo and her family, and Silvia Rabadán and her family. Thank you all... You made me strong and helped me grow... Wherever I go, you are in my heart...

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

ABSTRACT

THREE ESSAYS ON RESOURCE ALLOCATION: LOAD BALANCING ON HIGHLY VARIABLE SERVICE TIME NETWORKS, MANAGING DEFAULT RISK VIA SUBSIDIES AND SUPPLIER DIVERSIFICATION, AND OPTIMAL HOTEL ROOM ASSIGNMENT

by

Luz Adriana E. Caudillo Fuentes

Co-Chairs: Volodymyr Babich and Mark Van Oyen

The first essay considers a service center with two stations in accordance with independent Poisson processes. Service times at either station follow the same general distribution, are independent of each other and are independent of the arrival process. The system is charged station-dependent holding costs at each station per customer per unit time. At any point in time, a decision-maker may decide to move, at a cost, some number of jobs from one queue to the other. We study the problem with the purpose of providing insights into this decision-making scenario. We do so, in the important case that the service time distribution is highly variable or simply has a heavy tail. We propose that the savvy use of Markov decision processes can lead to easily implementable heuristics when features of the service time distribution can be captured by introducing multiple customer classes. The second essay studies the problem solved by a manufacturer who faces supplier disruptions. In order to understand the interactions between three strategies (subsidizing the supplier, supplier diversification, and the creation of back-up inventory), the problem is analyzed using a simple model with inventory storage costs and shortage penalties. The model allows us to derive conditions when these strategies are appropriate, either in isolation or in combination. A sensitivity analysis shows that the optimal decisions may not change monotonically when the parameters change. The third essay studies a hotel room assignment problem. The assignment is generally performed by the front desk staff on the arrival day using a lexicographic approach, but this may create empty room-nights between bookings that are hard to fill. This problem shares some features with the job shop problem and with the classroom assignment problem, both of which have been studied in the literature, but the problem itself has not been widely studied. We suggest a heuristic method to solve it, which can be run in a short time with the nightly batch operations that hotels routinely perform. The algorithm considerably improves the results from the lexicographic approach.

CHAPTER I

Introduction

The following three essays, among other techniques, apply Operations Research and to problems that arise in different industries. The nature of the three particular problems we present is complex. They cannot be solved by the simple, straightforward application of existing techniques; therefore, each constitutes an interesting research question. The first essay deals with work-load balancing for a system of queues. The second one analyzes the optimality of subsidies, backup inventory, diversification, and the combination of these strategies to deal with supplier disruption. The third one solves a hotel room assignment problem.

The first essay considers a service center with two stations in accordance with independent Poisson processes. Service times at either station follow the same general distribution, are independent of each other and are independent of the arrival process. The system is charged station-dependent holding costs at each station per customer per unit time. At any point in time, a decision-maker may decide to move, at a cost, some number of jobs in one queue to the other. The goals of this paper are twofold. We study the problem with the purpose of providing insights into this decision-making scenario. We do so, in the important case that the service time distribution is highly variable or simply has a heavy tail. Second, we propose that the savvy use of Markov decision processes can lead to easily implementable heuristics when features of the service time distribution can be captured by introducing multiple customer classes. These heuristic policies render better results than the application of common heuristics such as “no idling” or “join the shortest queue.”

The second essay studies the problem solved by a manufacturer that faces supplier disruptions. Different strategies to manage the risk are available to the manufacturer, and in order to understand the interactions between some of them (subsidizing the supplier, supplier diversification, and the creation of backup inventory), I analyze a model with costs coming from two sources: the costs of storing inventory when it exceeds demand, and the penalties incurred when not enough inventory exists to satisfy demand. The model is simple, yet it allows us to derive conditions when these strategies are appropriate, either in isolation or in combination. Using this model, I obtain some interesting insights. Under some circumstances, it may not be optimal to use a supplier exclusively, unless providing some minimum amount of subsidy. If no subsidy is provided, then it is better to diversify. Introducing time creates additional opportunities to hedge the risk: besides subsidizing the supplier so that she becomes more reliable in the future, the manufacturer also has the choice of having back-up inventory. In some cases, the three strategies can be used to manage the risk. However, changes in market conditions that affect, for example, future demand do not have an obvious pattern of change in the order size. I present a sensitivity analysis that can assist in these decisions.

The third essay studies a hotel room assignment problem. This problem, generally solved by the front desk staff on arrival day using a lexicographic approach, may create empty room-nights between bookings, which are usually hard to fill. While the day-of-arrival assignment practice has helped to implement overbooking, which is one of the most commonly used revenue management techniques, we studied the problem in order to find a better way to solve it. We suggest a heuristic method to solve it, and the testing we performed indicates that the room number assignment can be greatly improved by switching from the lexicographic approach to the algorithm we suggest. Our algorithm can be run in a short time with the nightly batch operations that hotels routinely perform. It is encouraging to obtain these results for a problem that has not received much attention. As we will see, this problem shares some features with the job shop and the classroom assignment problems which are NP-complete. These problems have been studied in the literature. However, their solution methods do not translate transparently to hotel settings, given the unique nature of this problem. Our algorithm provides the front desk staff with a tool that results in better solutions than commercial hotel operations software, which allows clerks to assign each room without any optimality check.

Before proceeding to the body of the different research topics, I would like to clarify that the use of the personal pronouns "he" or "she" in this work is not intended to be discriminatory.

CHAPTER II

A Simple Heuristic for Load Balancing in Parallel Processing Networks with Heavy-Tail Distributed Service Times

2.1 Introduction

Suppose that customers arrive to a service center (call center, web server, etc.) with two stations with independent Poisson arrival processes. Service times at either station follow the same general distribution, are independent of each other and are independent of the arrival process. The system is charged station dependent holding costs at each station per customer per unit time. At any point in time, a decision-maker may decide to move (or pass) some number of jobs in one queue to the other. It should be clear that the decision-maker's choice of the number of customers to move should depend on the number of customers at each station, the cost to move customers, the time elapsed since the service times of customers currently being processed by the server began and perhaps the number of future customers (s)he expects to arrive in the coming moments. With the exception of the elapsed service time information the control decisions seem ripe for an analysis via Markov decision processes (MDPs). Unfortunately, the continuous nature of the elapsed time variables makes the analysis more difficult. It is desirable to have a discrete state space to facilitate analysis and computation.

The goals of this paper are twofold. First, we are interested in providing insights into the above decision-making scenario. We do so, in the important case that the service time distribution is highly variable or simply has a heavy (non-exponential) tail. Second, we propose that the savvy use of Markov decision processes can lead to easily implementable heuristics when features of the service time distribution can be captured by introducing multiple customer classes.

The limitations of MDPs are well-known. As long as the state and action space descriptions (called the graph of the MDP) are multi-dimensional or consist of a large number of elements, solving the dynamic program quickly becomes intractable. In order to alleviate this problem there have been significant lines of research that study the structure of optimal policies in such areas as control of queues, manufacturing, transportation, inventory control and revenue management. For example, in the aforementioned model suppose the service time distribution is exponential. The state space is then two-dimensional. If the optimal policy can be described by a *monotone switching curve* the search for the optimal policy is reduced to finding the curve, rather than enumerating the state and action pairs throughout the decision space. Unfortunately, even in simple cases finding a structured optimal policy when the service time is generally distributed may be intractable. Consider an admission controlled $M/G/1$ queue that is used to model routing in a simple manufacturing system. If the service time distribution is exponential (so the system is an $M/M/1$), it is well-known that the optimal control policy is of *control limit* form. If the service times are generally distributed, then when a new customer arrives, the decision-maker must once again consider the time since the last service completion; the state space is uncountable. In this case, even reasonably sized discretizations of the time dimension lead to an intractable problem.

In reality, when the service time distribution is general, past experience gives the decision-maker significant information about its form. In this paper we present a heuristic that uses a multi-class queueing network with exponentially distributed service times as a proxy for a problem with general service time distributions. In essence, for the original model with general distributions, services are classified into types as a way to record partial historical information that is useful for making control decisions. The proxy network problem has a tractable MDP solution with a control structure that is relevant, in a heuristic sense, to the intractable control problem faced in the original model. We focus our attention on the important case that the general distribution has a "heavy tail" (does not decay exponentially) or is highly variable. Intuitively, we are interested in systems where long service times provide significant useful information about the remaining service time distribution. We discuss our heuristic in the context of the new load-balancing model described above. Although it has its roots in service centers, it is also applicable to supply chain management and to transhipment models in transportation networks.

We should point out that the goal is to introduce a method for approximating the load balancing **decisions** made in a parallel processing network, not to approximate the service time distributions themselves. With an eye towards tractability and solutions that are easy to describe and implement, we restrict attention in the proxy model to a hyper-exponential (mixture of exponentials) service distribution with two classes. We find that the optimal control policies for the two-class proxy model, when translated in a smart way (as discussed in Section 2.5), indeed lead to policies that perform well in the original system. Alternatively, we conjecture that one could approximate service times with an Erlang distribution with k phases, and provide a similar analysis using a Markov decision process formulation. The difficulty would then be in translating that process into an implementable control policy. Moreover, the decision problem would be intractable for k of moderate size. The optimal control for the proxy model we propose is quite simple. More sophisticated MDP models would quickly lose this feature.

This paper makes several contributions. Of course, we describe a method for determining good control policies for the otherwise intractable load balancing problem. The employed proxy model, which is also new, most likely has applications outside of this context, and we find it interesting in its own right. For the proxy model, we show that the optimal control structure is characterized by a series of “do-not-move/move-up-to levels” and that these levels are monotone. Not only do these structural results provide insight, but they also aid in computation. In particular, since the state space of the proxy model has infinite dimension, computation is facilitated by truncation of the queue lengths. Truncation often leads to policies that are not monotone near the boundaries. However, we “smooth” the policies in accordance with the theoretical results, and we find that these smoothed policies perform better. Finally, performance was measured via simulation. We display the results of the numerical study, which show that our policies perform well as compared to some alternative heuristics.

The remainder of the paper is organized as follows. In Section 2.2 we discuss related literature. Section 2.3 contains some preliminary results, a further description of the original and proxy models and the optimality criteria. We present a Markov decision formulation of the proxy model and show several monotonicity results in Section 2.4. The description of our heuristic for controlling the original load balancing problem, including the relationship to the proxy model, is given in Section 2.5. Section 2.5 also contains an implementation of the heuristic and the numerical comparison to several alternative heuristics. The paper is concluded in Section 2.6.

2.2 Literature review

The theory (and drawbacks) of MDPs is well-documented. We refer the interested reader to the now classic text of *Puterman* (1994). The literature on the control of parallel processing networks is also abundant so we do not provide a complete review here. Instead the reader is pointed to the work of *Shirazi et al.* (1995) and *Wang and Morris* (1985) and the references therein. We focus on those papers with direct relevance to the current work. For a basic introduction to heavy-tailed distributions and their properties, see *Sigman* (1999). A discussion of several alternative definitions can be found in *Heyde and S.G.* (2004).

Paxson and Floyd (1995) have found that for most of the traffic in the world wide web *session and connection* arrivals are modeled well using Poisson processes, but *packet* interarrivals are better described with heavy tailed distributions. This is further confirmed by *Crovella et al.* (1998b). In particular, the hyper-exponential distribution has proved to be useful to approximate heavy-tailed distributions. *Xu et al.* (2003) use such approximations to formulate generalized Petri nets in order to study the properties of distributed manufacturing systems. The hyper-exponential is one of the motivating factors of our two class Markov decision process formulation.

Harchol-Balter and Downey (1997) compare the reassignment of processes to a different server at the time of birth vs. reassignment once the process has already started (preemptive migration) in order to balance CPU load in a network of stations. They obtain a preemptive reassignment strategy that is more effective than remote execution even when the memory transfer cost is high. *Yum and Hua-chun* (1984) develop an adaptive rule for balancing the load on a parallel queueing system, where some customers are required to wait for a particular server or set of servers. Their rule is a combination of a majority-vote rule (where votes are issued by switchers or routers) and a *join the biased queue* rule as presented by *Yum and Schwartz* (1981). Yum and Schwartz use this term to denote a rule similar to *join the shortest queue*, but a bias term is added to the queue lengths. This rule is robust to changes in the buffer sizes and input rates, and performs well according to the criteria of lower delay and lower blocking probability. *Shimkin and Shwartz* (1989) study a system of queues that share an arrival process. Arriving customers are subject to admission and routing control. The purpose is to maximize income when there are holding costs and rewards for accepting customers. The arrival and service process parameters depend on the current state of the system. The authors prove the existence of a monotone optimal control policy.

Other research on systems with heavy-tail distributed service times includes *Crovella et al.* (1998a) who develop a policy that purposely operates the server hosts at different loads, and directs smaller tasks to the lighter-loaded hosts. *Riska et al.* (2000) present an inexpensive technique for modeling load balancing policies on a cluster of servers conditioned on the fact that the service times of arriving tasks are drawn from heavy-tailed distributions. Their results provide exact information regarding the distribution of task sizes that compose the queue on each server. *Beard and Frost* (2001) study a prioritization mechanism to alleviate overloads that result in blocking the access to service to all customers. Of course none of these studies include a Markov decision process formulation of an exponential model applied to the general model with heavy tailed distributed service times.

Our model is closely related to that in *Down and Lewis* (2006). Their work refers to a system of parallel queues, where the balancing decisions are taken at the times of arrivals or departures. They seek the optimal design and control policy for the system. There is also a close relation to the work of *Lewis* (2001) where an M/M/1 queue is controlled by two "gatekeepers" that make the decisions of acceptance or rejection of a customer at two moments: the arrival and the moment prior to service. Another study related to the control of queueing systems with exponentially distributed service times can also be found in *He and Neuts* (2002), who study policies that move a fixed amount of customers to control a system of two M/M/1 queues. Transfer of customers occurs when the difference of the queues reaches a critical level.

2.3 Preliminaries and Model Descriptions

In this section we discuss the formal definition of a parallel processing network with service times that follow a heavy-tailed distribution and the *proxy* model with exponential service times. Consider 2 parallel queues. Customers arrive to queue k according to independent Poisson processes of rate λ_k for $k = 1, 2$. The service processes of each queue are independent of each other and of both arrival streams. The n^{th} customer that is served by server k requires S_n^k time units of service where $\{S_n^k, n \geq 1,\ k = 1, 2\}$ are assumed to be i.i.d. and independent of the station to which the customer arrives. In the *general* model that motivates this study, the service times are assumed to follow a general distribution with finite mean. However, we are most interested in those service distributions that see a large proportion of short service times, but also see some very large service times; those that are highly variable. One such class of distributions is that with "heavy," non-exponential tails. We present the definition given by *Sigman* (1999).

Definition 2.3.1. A distribution function F, for random variable S, is said to be **heavy-tailed** if $\overline{F}(s) := 1 - F(s) = \mathbb{P}(S > s) > 0$, $s \geq 0$, and

$$\lim_{s\to\infty} \mathbb{P}(S > s + \delta | S > s) = \lim_{s\to\infty} \frac{\overline{F}(s + \delta)}{\overline{F}(s)} = 1,\ \delta \geq 0. \tag{2.3.1}$$

Intuitively, if S follows a heavy-tailed distribution, then if S ever exceeds a large value T_1, the probability that it will exceed any higher level approaches 1 for $T_2 > T_1$ gets large. Thus, while most times are short, a decision-maker that finds a customer whose service time is unusually long would not want to leave customers in queue behind it.

As an approximation to this model we consider a *proxy* model, where each arriving customer is of one of two classes. A customer's classification is not revealed until immediately prior to beginning service. Customers are of class j, $j = 1, 2$, with probability p_j and a class j customer requires an exponentially distributed amount of service with mean $1/\mu_j$. We assume that Class 1 are those with unusually long ("heavy") service times, while Class 2 corresponds to those with shorter ("standard") service times seen in the general model; that is, $1/\mu_1 \gg 1/\mu_2$. We will explain exactly how they are related to the general model when the heuristic is described more fully in Section 2.5.

In either model, let Π be the set of all non-anticipating policies. A policy $\pi \in \Pi$ prescribes how many customers to move from one queue to another, given the number of customers in each queue (the queue length processes), perhaps the amount of time each customer has been in service, and any other information that is required to make the (policy dependent) process Markovian. For example, in the proxy system the current "state" of the system includes the queue length processes and the classes of the customers currently in service at each queue.

There is a fixed cost for moving each customer of m units per customer. That is, if θ customers are moved, a cost of $m\theta$ is incurred. Customers currently in service (in either queue) cannot be moved; the control policy is assumed to be non-preemptive. The system also continuously incurs holding cost $h_k q_k$ per unit time that queue k contains q_k customers, including the one in service for $k = 1, 2$. Without loss of generality we assume that $h_1 \geq h_2$. We seek to find a strategy for load balancing under the infinite horizon expected discounted cost or the long-run average expected cost optimality criteria. Note here that the term "load balancing" is used somewhat loosely since the holding costs may cause the optimal policy to leave the distribution of the workload for each queue unbalanced. In some sense, perhaps "load distribution" would be more descriptive. However, having made this clarification, we will continue to refer to the control as balancing without further comment since it is common terminology.

For a fixed policy π, denote the set of decision epochs by $\mathbb{D} \equiv \{d_n, n \geq 0\}$ and the state at the n^{th} decision epoch by X_n. For example, if π depends only on the queue lengths, then $\mathbb{D}$ is the set of arrival times and service time completions. We assume that the time between decision epochs is bounded away from zero so that only a finite number of decisions can be made in a finite amount of time. That is, if the time between the n^{th} and $(n+1)^{st}$ decision epoch has distribution G_{n+1} then there exists $\delta > 0$ and $\epsilon > 0$ such that $1 - G_{n+1}(\delta) \geq \epsilon$ (cf. p. 532 of *Puterman* (1994)). Let $Q^{\pi}(t) = \{Q_1^{\pi}(t), Q_2^{\pi}(t)\}$ be the queue length process, and let θ_n represent the balancing decision taken at decision epoch n, under π. Define the total discounted expected cost up until time t as

$$v_{\beta,t}^{\pi}(x) = \mathbb{E}_x^{\pi}\Big(\sum_{n=0}^{N(t)} e^{-\beta d_n} c(X_n, \theta_n)\Big) + \int_0^t e^{-\beta u}\, \mathbb{E}_x^{\pi}[h_1 Q_1^{\pi}(u) + h_2 Q_2^{\pi}(u)]\mathrm{d}u,$$

where θ_n is the action taken at decision epoch n, $c(\cdot,\cdot)$ is the lump sum cost associated with moving customers from one queue to the other, $N(t)$ is the number of decision epochs in the first t time units, and the expectation of the system under policy π is conditioned on the initial state x. The criteria we are interested in are

$$v_{\beta}^{\pi}(x) = \lim_{t\to\infty} v_{\beta,t}^{\pi}(x), \qquad\qquad \varphi^{\pi}(x) = \limsup_{t\to\infty} \frac{v_{0,t}^{\pi}(x)}{t},$$

where $v_{\beta}^{\pi}(x)$ represents the infinite horizon $\beta-$discounted expected cost under π (the interchange of limit and expectation is justified by the monotone convergence theorem) and $\varphi^{\pi}(x)$ is called the long-run average expected cost starting in state x under policy π. The objective then is to find a policy π^* under each criterion such that $\gamma^{\pi^*}(x) \leq \gamma^{\pi}(x)$ for all states x and all policies $\pi \in \Pi$ for $\gamma = v_\beta, \varphi$. In the next section we provide results that simplify this search considerably for the proxy model. We view these results as interesting in their own right, but they are particularly useful in the implementation of our heuristic in the general model.

2.4 Optimal Control for the Proxy Model

For the proxy model all inter-arrival and service times are exponentially distributed, and the state may be described by a vector (I, y, i, j), where I represents the total number of customers in the system and y is the number of customers in queue 2 (including any customer in service). When $i\,(j) \in \{1,2\}$ it represents the class of customer currently at server $1\,(2)$; $i\,(j) = 0$ means that queue 1 (2) is empty. If $x = (I, y, i, j)$, then the possible actions set is $A_x = \{-(y-1)^+, -(y-2), \ldots, I-y-2, (I-y-1)^+\}$. That is, for $\theta \in A_x$, $\theta = 0$ means that nothing will be moved while $\theta > 0$ means θ customers are moved from queue 1 to queue 2 and $\theta < 0$ means that $|\theta|$ are moved from queue 2 to queue 1. A customer that is currently in service cannot be moved. Let $\mathcal{W} := \{(I, y, i, j) \mid I-1 \geq y \geq 1, i, j \in \{1,2\}\}$ represent the set of states such that both servers have at least one customer to serve. Similarly, define $\mathcal{I}_1 := \{(I, y, 0, j) \mid I = y \geq 1\}$ and $\mathcal{I}_2 := \{(I, y, i, 0) \mid I \geq 1, y = 0\}$, where $\mathcal{I}_k$ represents the set of states where there are no customers to serve in queue $k = 1, 2$ while the other queue is non-empty (the $\mathcal{I}$ stands for "idle"). The state space $\mathbb{X}$ can now be written

$$\mathbb{X} := \mathcal{W} \cup \mathcal{I}_1 \cup \mathcal{I}_2 \cup \{(0,0,0,0)\}.$$

We apply *uniformization* as described in *Lippman* (1975), with uniformization constant $\Psi = \lambda_1 + \lambda_2 + 2\max\{\mu_1, \mu_2\}$. Without loss of generality assume $\Psi = 1$. This allows us to consider the discrete-time *equivalent* to the continuous proxy model already described. That is to say that the stationary optimal policies in the discrete-time case are the same as that in the continuous-time case. The infinite horizon discounted cost and the long-run average costs also coincide, but only up to a multiplicative constant. The cost function for each period includes holding and switching costs and is given by:

$$C((I, y, i, j), \theta) = \ |\theta| \ m + (I - y - \theta)h_1 + (y + \theta)h_2,$$

where (I, y, i, j) and θ denote the current state and action, respectively. Let the total expected cost of a load balancing policy π over the first t (discrete) decision epochs be defined

$$v^{\pi}_{\alpha,t}(x) := \mathbb{E}^{\pi}_{x} \sum_{n=0}^{t-1} \alpha^{t} C(X_n, \theta_n),$$

where X_n and θ_n represent the state and balancing decision at decision epoch n. Furthermore, define $v_{\alpha,t}(x) = \inf_{\pi \in \Pi} v^{\pi}_{\alpha,t}(x)$, where Π is the set of all non-anticipating policies; $v_{\alpha,t}(x)$ is the optimal cost-to-go for a t-horizon problem starting in state x, under discount factor α. In the case when $t = \infty$, we write v_α instead of $v_{\alpha,\infty}$. This defines the infinite horizon expected discounted cost criterion. As we are also interested in the average case, the average cost of a fixed policy in the discrete time model equals

$$\limsup_{n \to \infty} \frac{1}{n} v^{\pi}_{1,n}(x) = \limsup_{n \to \infty} \frac{1}{n} \mathbb{E}^{\pi}_{x} \sum_{t=1}^{n} C(X_t, \theta_t).$$

In the remainder of the section, we present several structural results for the finite horizon case. We then give a stability result and show that the structural results continue to hold in the infinite horizon discounted cost and average cost cases. The first result states that the basic features of an optimal policy in *Down and Lewis* (2006) carry over to the current model.

Proposition II.1. *Under the finite or infinite horizon discounted expected cost or the long-run average expected cost criterion, there exists an optimal policy that does not move customers from queue 2 to queue 1 (since $h_1 \geq h_2$), except possibly to avoid idling.*

Proof. The proof follows precisely in the same manner as that in Theorem 4.1 of *Down and Lewis* (2006) and is omitted for brevity. ■

Suppose $b = y + \theta$ for $\theta \in A_{(I,y,i,j)}$ (i.e., the number of customers in the low cost queue after performing the control action). Let $\mu_0 = 0$. Define $w_{\alpha,t}(I, y, i, j, b)$ as the cost-to-go, starting in state (I, y, i, j), for moving up to amount b in period t, followed by optimal control in the remaining periods:

$$w_{\alpha,t}(I,y,i,j,b) = \begin{cases} m|b-y| + h_1(I-b) + h_2 b & \\ +U_{\alpha,t-1}(I,b,i,j), & \text{for } 1 \le y \le I-1, \\ m|b-y| + h_1(I-b) + h_2 b & \\ +p_1 U_{\alpha,t-1}(I,b,1,j) + p_2 U_{\alpha,t-1}(I,b,2,j) & \text{for } y = I, i = 0, b < I, \\ h_2 I + U_{\alpha,t-1}(I,b,0,j) & \text{for } y = b = I, i = 0, \\ m|b-y| + h_1(I-b) + h_2 b & \\ +p_1 U_{\alpha,t-1}(I,b,i,1) + p_2 U_{\alpha,t-1}(I,b,i,2) & \text{for } y = 0, j = 0, b > 1, \\ h_1 I + U_{\alpha,t-1}(I,b,i,0) & \text{for } y = b = 0, j = 0, \end{cases}$$

where for $(I, b, i, j) \in \mathcal{W}$

$$\begin{aligned} U_{\alpha,t}(I,b,i,j) = \alpha[&p_1\mu_i v_{\alpha,t}(I-1,b,1,j) + p_2\mu_i v_{\alpha,t}(I-1,b,2,j) \\ &+ p_1\mu_j v_{\alpha,t}(I-1,b-1,i,1) + p_2\mu_j v_{\alpha,t}(I-1,b-1,i,2) \\ &+ \lambda_1 v_{\alpha,t}(I+1,b,i,j) + \lambda_2 v_{\alpha,t}(I+1,b+1,i,j) \\ &+ (1-\lambda_1-\lambda_2-\mu_i-\mu_j) v_{\alpha,t}(I,b,i,j)], \end{aligned}$$

for $(I,b,0,j) \in \mathcal{I}_1$ ($b = I$ in this case)

$$\begin{aligned}U_{\alpha,t}(I,I,0,j) = \alpha[&p_1\mu_j v_{\alpha,t}(I-1,I-1,0,1) + p_2\mu_j v_{\alpha,t}(I-1,I-1,0,2)\\&+ \lambda_1(p_1 v_{\alpha,t}(I+1,I,1,j) + p_2 v_{\alpha,t}(I+1,I,2,j))\\&+ \lambda_2 v_{\alpha,t}(I+1,I+1,0,j) + (1-\lambda_1-\lambda_2-\mu_j)v_{\alpha,t}(I,I,i,j)],\end{aligned}$$

for $(I,b,i,0) \in \mathcal{I}_2$ ($b = 0$ in this case)

$$\begin{aligned}U_{\alpha,t}(I,0,i,j) = \alpha[&p_1\mu_i v_{\alpha,t}(I-1,0,1,0) + p_2\mu_i v_{\alpha,t}(I-1,0,2,0)\\&+ \lambda_1 v_{\alpha,t}(I+1,0,i,0) + \lambda_2(p_1 v_{\alpha,t}(I+1,1,i,1)\\&+ p_2 v_{\alpha,t}(I+1,1,i,2)) + (1-\lambda_1-\lambda_2-\mu_i)v_{\alpha,t}(I,0,i,0)],\end{aligned}$$

and for $I = b = 0$

$$\begin{aligned}U_{\alpha,t}(0,0,0,0) = \alpha[&\lambda_1(p_1 v_{\alpha,t}(1,0,1,0) + p_2 v_{\alpha,t}(1,0,2,0)) + \lambda_2(p_1 v_{\alpha,t}(1,1,0,1)\\&+ p_2 v_{\alpha,t}(1,1,0,2)) + (1-\lambda_1-\lambda_2)v_{\alpha,t}(0,0,0,0)].\end{aligned}$$

Let $A_{\mathcal{W}} := \{1,2,\ldots,I-1\}$. Similarly define $A_{\mathcal{I}_1} := \{1,2,\ldots I\}$, $A_{\mathcal{I}_2} := \{0,1,\ldots I-1\}$ and $A_{(0,0,i,j)} := \{0\}$. It is well-known that for $(I,y,i,j) \in \mathcal{W}$, $v_{\alpha,t}$ (note $v_{\alpha,0} = 0$) satisfies the following *finite horizon optimality equations* (FHOE)

$$v_{\alpha,t}(I,y,i,j) = \min_{b \in A_{\mathcal{K}}}\{w_{\alpha,t}(I,y,i,j,b)\}, \tag{2.4.1}$$

where $\mathcal{K} = \mathcal{W}, \mathcal{I}_1, \mathcal{I}_2,$ or $(0,0,0,0)$ depending on (I,y,i,j).

The next result states that there exists an optimal policy such that for each state there is a "do-not-move/ move-up-to" amount L. This means that if $y < L$ we move enough customers to have L customers in queue 2, and if $y \geq L$, we move no customers.

Proposition II.2. *Suppose the current types at servers* 1 *and* 2 *are* i *and* j*, respectively. Then,*

1. *there exists a level* $L^t_{I,i,j} < I$ *such that for each* $t \geq 1$*,* $I \geq 2$ *and* $(I, y, i, j) \in \mathcal{W} \cup \mathcal{I}_2$*, the optimal policy is to bring the number of customers in queue 2 up to* $L^t_{I,i,j}$ *if* $y < L^t_{I,i,j}$ *and to move no customers if* $y \geq L^t_{I,i,j}$*, and*

2. $v_{\alpha,t}(I, y, i, j) - v_{\alpha,t}(I, y+1, i, j)$ *is non-decreasing in* y *(i.e.,* $v_{\alpha,t}$ *is convex in* y*) for all* i, j*,* $t \geq 0$*,* $I \geq 3$*, and* $y \leq I - 3$*.*

Proof. By induction on t. Recall that $v_{\alpha,0}(\cdot) = 0$. So that Statement 2 holds trivially for $t = 0$. Consider $t = 1$ and assume $I \geq 2$. We have

$$w_{\alpha,1}(I, y, i, j, b) = m|b - y| + h_1(I - b) + h_2 b. \tag{2.4.2}$$

Since $(I, y, i, j) \in \mathcal{W} \cup \mathcal{I}_2$ queue 1 is non-empty. Recall from Proposition II.1 that it is not optimal to move customers from queue 2 (the low-cost queue) to queue 1 (the high-cost queue) unless possibly if queue 1 is empty. That is, it suffices to consider $b \geq y$. When we restrict attention to the set $\{y, y + 1, ..., I - 1\}$, $w_{\alpha,1}$ is a linear function of b (since on this set $|b - y| = b - y$). Depending on the direction of the inequality $m - h_1 + h_2 \geq \ (\leq)\, 0$ the optimal action is either not to move customers or to move all of the customers to the low cost queue (except for the one currently receiving service at queue 1). That is to say either letting $L^1_{I,i,j} = 0$ or $I - 1$ is optimal. This proves the first statement for $t = 1$.

Assume Statement 2 holds for $t - 1$. To prove Statement 1 at time t recall

$$w_{\alpha,t}(I, y, i, j, b) = m|b - y| + h_1(I - b) + h_2 b + U_{\alpha,t-1}(I, b, i, j).$$

From Statement 2 at epoch $t - 1$, and from the definition of $U_{\alpha,t}$, $U_{\alpha,t-1}(I, b, i, j)$ is a convex combination of convex functions. Thus, $w_{\alpha,t}(I, y, i, j, b)$ is convex in b.

Let $L^t_{I,i,j}$ be the minimal (smallest) element of the set $\text{argmin}_{b \in A_{\mathcal{W}}}\{w_{\alpha,t}(I,1,i,j,b)\}$. Note that by convexity $L^t_{I,i,j}$ is also in the $\text{argmin}_{b\in\{y,y+1,\ldots,I-1\}}\{w_{\alpha,t}(I,y,i,j,b)\}$ for all $1 \le y \le L^t_{I,i,j}$. The convexity of $w_{\alpha,t}$ together with Proposition II.1 yield that it is not optimal to move customers for $y > L^t_{I,i,j}$. This proves Statement 1.

To complete the proof it remains to show that the preceding arguments imply that Statement 2 holds for time t and all (i,j). There are several cases to consider. Suppose for now that $(I,y,i,j) \in \mathcal{W}$.

Case 1: $L^t_{I,i,j} \le y$. In this case, it is optimal not to move in the three states (I,y,i,j), $(I,y+1,i,j)$, $(I,y+2,i,j)$. Thus,

$$\begin{aligned}&v_{\alpha,t}(I,y,i,j) - 2v_{\alpha,t}(I,y+1,i,j) + v_{\alpha,t}(I,y+2,i,j)\\ &\quad = U_{\alpha,t-1}(I,y,i,j) - 2U_{\alpha,t-1}(I,y+1,i,j) + U_{\alpha,t-1}(I,y+2,i,j) \ge 0,\end{aligned}$$

where the inequality follows from the inductive hypothesis that the second statement holds at $t-1$.

Case 2: $y \le L^t_{I,i,j} - 2$. The optimal action in the three states (I,y,i,j), $(I,y+1,i,j)$, $(I,y+2,i,j)$ is to allocate $I - L^t_{I,i,j}$ customers in queue 1 and $L^t_{I,i,j}$ in queue 2. Thus,

$$v_{\alpha,t}(I,y,i,j) - v_{\alpha,t}(I,y+1,i,j) = v_{\alpha,t}(I,y+1,i,j) - v_{\alpha,t}(I,y+2,i,j) = m.$$

Case 3: $y = L^t_{I,i,j} - 1$. First note $w_{\alpha,t}(I,y+1,i,j,y+2) \ge v_{\alpha,t}(I,y+1,i,j)$; moving one customer is a (potentially) suboptimal action for state $(I,y+1,i,j)$). Thus,

$$\begin{aligned}&v_{\alpha,t}(I,y,i,j) - v_{\alpha,t}(I,y+1,i,j) - [v_{\alpha,t}(I,y+1,i,j) - v_{\alpha,t}(I,y+2,i,j)]\\ &\quad \ge v_{\alpha,t}(I,y,i,j) - w_{\alpha,t}(I,y+1,i,j,y+1) - w_{\alpha,t}(I,y+1,i,j,y+2) + v_{\alpha,t}(I,y+2,i,j)\end{aligned}$$

Note that the last expression is equal to $m + U_{\alpha,t-1}(I, y+1, i, j) - U_{\alpha,t-1}(I, y+1, i, j) - m - U_{\alpha,t-1}(I, y+2, i, j) + U_{\alpha,t-1}(I, y+2, i, j) = 0$ as desired.

Suppose $(I, y, i, j) \in \mathcal{I}_2$.

Let $L^t_{I,i,0}$ be the minimal element of the set $\text{argmin}_{b \in A_{\mathcal{I}_2}}\{w_{\alpha,t}(I, 0, i, j, b)\}$. Thus, $v_{\alpha,t}(I, 0, i, j) = w_t(I, 0, i, j, L^t_{I,i,0})$. Assuming, $L^t_{I,i,2} \geq 2$

$$
\begin{aligned}
&v_{\alpha,t}(I, 0, i, j) - v_{\alpha,t}(I, 1, i, j) - [v_{\alpha,t}(I, 1, i, j) - v_{\alpha,t}(I, 2, i, j)] \\
&\quad \geq w_{\alpha,t}(I, 0, i, j, L^t_{I,i,0}) - w_{\alpha,t}(I, 1, i, j, L^t_{I,i,0}) - [w_{\alpha,t}(I, 0, i, j, L^t_{I,i,2}) - w_{\alpha,t}(I, 2, i, j, L^t_{I,i,2})] \\
&\quad = m - m = 0.
\end{aligned}
$$

Similarly for $L^t_{I,i,2} < 2$. Since for all $(I, y, i, j) \in \mathcal{I}_1 \cup \{(0, 0, 0, 0)\}$, $y = I$ there is no convexity requirement on $\mathcal{I}_1 \cup \{(0, 0, 0, 0)\}$. ■

We remark that the previous result states the existence of an optimal "move-up-to" level for each (I, i, j). We next characterize these levels as monotone, non-decreasing in I. This result not only lends insight into the structure of the optimal policy, but it is also convenient both for implementation and to simplify its computation. Moreover, it is used to implement the load balancing heuristic presented in Section 2.5. Before proving the result, it is useful to recall the definition of submodularity:

Definition 2.4.1. A function $g(j, k)$ is said to be **submodular** if and only if the difference $g(j, k) - g(j, k+1)$ is non-decreasing in j; that is, $g(j, k) - g(j, k+1) \leq g(j+1, k) - g(j+1, k+1)$.

Proposition II.3. *Let $I \geq 3$, $y \in \{0, ..., I-1\}$. Suppose the current types at servers 1 and 2 are i and j respectively. The following hold:*

1. *For $t \geq 1$ there exists optimal move-up-to levels $L^t_{I+1,i,j}$ and $L^t_{I,i,j}$ such that $L^t_{I+1,i,j} \geq L^t_{I,i,j}$.*

2. *$v_{\alpha,t}(I, y, i, j) - v_{\alpha,t}(I, y+1, i, j)$ is non-decreasing in I (i.e., $v_{\alpha,t}$ is submodular in (I, y)) for all $t \geq 0$ and $1 \leq y \leq I - 3$.*

Proof. By induction on t. For $t = 0$ Statement 2 holds trivially since $v_{\alpha,0} = 0$. At $t = 1$ we have $w^1(I, y, i, j, b) = m|b - y| + h_1(I - b) + h_2 b$. As in the proof of Proposition II.2 it is optimal either **(a)** not to move any customers, or **(b)** to move all the customers to the low cost queue (except for the one currently receiving service at queue 1). That is to say the optimal move up to level is $L^1_{I,i,j} = y$ or $I - 1$ depending on the direction of the inequality $m - h_1 + h_2 \geq (\leq) 0$. Similarly, for state $(I + 1, y, i, j)$, the optimal move up to level is $L^1_{I+1,i,j} = y$ or I (depending on the same inequality. Thus, $L^1_{I+1,i,j} \geq L^1_{I,i,j}$ as desired.

Assume now that Statement 1 holds for t and Statement 2 for $t - 1$. There are 4 cases to consider to prove Statement 2 holds at time t. In each of the first three cases we take advantage of the fact that $w_{\alpha,t} \geq v_{\alpha,t}$.

Case 1: $y + 1 < L^t_{I,i,j}$ and $y < L^t_{I+1,i,j}$. Then,

$$\begin{aligned} &w_{\alpha,t}(I, y, i, j, L^t_{I,i,j}) - v_{\alpha,t}(I, y + 1, i, j) - v_{\alpha,t}(I + 1, y, i, j) \\ &\quad + w_{\alpha,t}(I + 1, y + 1, i, j, L^t_{I+1,i,j}) = m - m = 0. \end{aligned}$$

Case 2: $y + 1 \geq L^t_{I,i,j}$ but $y < L^t_{I+1,i,j}$. Then,

$$\begin{aligned} &w_{\alpha,t}(I, y, i, j, y + 1) - v_{\alpha,t}(I, y + 1, i, j) - v_{\alpha,t}(I + 1, y, i, j) \\ &\quad + w_{\alpha,t}(I + 1, y + 1, i, j, L^t_{I+1,i,j}) = m - m = 0. \end{aligned}$$

Case 3: $y + 1 \geq L^t_{I,i,j}$ and $y \geq L^t_{I+1,i,j}$. In this case we have

$$\begin{aligned} &w_{\alpha,t}(I, y, i, j, y) - v_{\alpha,t}(I, y + 1, i, j) - v_{\alpha,t}(I + 1, y, i, j) + w_{\alpha,t}(I + 1, y + 1, i, j, y + 1) \\ &\quad = h_1 - h_2 - (h_1 - h_2) + U_{\alpha,t-1}(I, y, i, j) - U_{\alpha,t-1}(I, y + 1, i, j) - U_{\alpha,t-1}(I + 1, y, i, j) \\ &\qquad + U_{\alpha,t-1}(I + 1, y + 1, i, j). \end{aligned}$$

Each of preceding 3 cases imply submodularity for $v_{\alpha,t}$ since $w_{\alpha,t} \geq v_{\alpha,t}$ with the last one also using the inductive hypothesis ($U_{\alpha,t-1}$ is a linear combination of $v_{\alpha,t-1}$).

Case 4: $y+1 < L^t_{I,i,j}$ and $y \geq L^t_{I+1,y,i,j}$. Note that since $y+1 < L^t_{I,i,j}$ we have $y < L^t_{I,i,j} \leq L^t_{I+1,i,j} \leq y$, where the second inequality follows from the inductive assumption. Since Case 4 leads to a contradiction it cannot occur.

It remains to show that $L^{t+1}_{I+1,i,j} \geq L^{t+1}_{I,i,j}$. First note that if $L^{t+1}_{I,i,j} = 1$ the result holds trivially. Assume that $L^{t+1}_{I,i,j} \geq 2$. Note that the submodularity of $v_{\alpha,t}$ implies submodularity of $U_{\alpha,t}(\cdot, y, \cdot, \cdot)$ for $y \geq 2$. Suppose the result does not hold so that $L^{t+1}_{I+1,i,j} < L^{t+1}_{I,i,j}$. Fix $L^{t+1}_{I+1,i,j} < y \leq L^{t+1}_{I,i,j}$ so that the optimal action in $(I+1, y, i, j)$ is to do nothing, while the optimal action in state (I, y, i, j) is to move the number of customers in queue 2 to $L^{t+1}_{I,i,j}$. By using $L^{t+1}_{I,i,j}$ in state $(I+1, y, i, j)$, the optimality equations imply

$$\begin{aligned} v_{\alpha,t+1}(I+1, y, i, j) &= h_1(I+1-y) + h_2(y) + U_{\alpha,t}(I+1, y, i, j) \\ &< m(L^{t+1}_{I,i,j} - y) + h_1(I+1-L^{t+1}_{I,i,j}) \\ &\quad + h_2(L^{t+1}_{I,i,j}) + U_{\alpha,t}(I+1, L^{t+1}_{I,i,j}, i, j). \end{aligned}$$

A little algebra yields

$$\begin{aligned} &m(L^t_{I,i,j} - y) + h_1(y - L^{t+1}_{I,i,j}) - h_2(L^t_{I,i,j} - y) \\ &\quad > U_{\alpha,t}(I+1, y, i, j) - U_{\alpha,t}(I+1, L^t_{I,i,j}, i, j). \end{aligned} \tag{2.4.3}$$

Similarly (by considering the action "do nothing" in state (I, y, i, j))

$$\begin{aligned} &m(L^t_{I,i,j} - y) + h_1(y - L^{t+1}_{I,i,j}) - h_2(L^t_{I,i,j} - y) \\ &\quad < U_{\alpha,t}(I, y, i, j) - U_{\alpha,t}(I, L^t_{I,i,j}, i, j). \end{aligned} \tag{2.4.4}$$

Combining (2.4.3) and (2.4.4) yields

$$U_{\alpha,t}(I, y, i, j) - U_{\alpha,t}(I, L^t_{I,i,j}, i, j) - [U_{\alpha,t}(I+1, y, i, j) - U_{\alpha,t}(I+1, L^t_{I,i,j}, i, j)] > 0,$$

which contradicts submodularity and the result is proven. ∎

2.4.1 The Infinite Horizon Discounted Cost and Average Cost Cases

In this section we note that the results from the previous section extend to the infinite horizon models. While the infinite horizon discounted cost case follows almost immediately, the average cost case is slightly more subtle and requires a stability result.

Proposition II.4. *For the proxy model, under any stationary, non-idling policy the system is stable if and only if*

$$(\lambda_1 + \lambda_2)\left(\frac{p_1}{\mu_1} + \frac{p_2}{\mu_2}\right) < 2. \qquad (2.4.5)$$

That is, there exists a stationary distribution.

Proof. To prove sufficiency we fix an arbitrary, stationary, non-idling policy π, find a *Lyapunov* function and apply *Foster's criterion* cf. (*Kulkarni*, 1999, Theorem 3.7). This guarantees that all recurrent states are positive recurrent. To this end, consider the Markov chain induced by π. Denote this chain, with state space $\mathbb{X}$, by $\{X_n, n \geq 0\}$. Note that since π is non-idling, $(0,0,0,0)$ is accessible from every state in $\mathbb{X}$; any recurrent states must communicate with the distinguished state $(0,0,0,0)$. Denote the chain restricted to only those states that communicate with $(0,0,0,0)$ by $\{Z_n, n \geq 0\}$ and its state space by $\mathbb{X}^0$. Let $G = \{(I, y, i, j) \in \mathbb{X} | I \leq 1\}$. Let

$\mu = \left(\frac{p_1}{\mu_1} + \frac{p_2}{\mu_2}\right)^{-1}$ and define

$$\mathcal{L}(I, y, i, j) \;=\; I/\mu + 1/\mu_i + 1/\mu_j,\ (I, y, i, j) \in \mathbb{X}.$$

For any action chosen and for $(I, y, i, j) \notin G$

$$\begin{aligned}
\mathbb{E}\left[\mathcal{L}(Z_{n+1}) - \mathcal{L}(Z_n) | Z_n = (I, y, i, j)\right] \;&=\; (\lambda_1 + \lambda_2 - \mu_i - \mu_j)/\mu \\
&\quad +\mu_i(p_1/\mu_1 + p_2/\mu_2 - 1/\mu_i) \\
&\quad +\mu_j(p_1/\mu_1 + p_2/\mu_2 - 1/\mu_j) \\
&= (\lambda_1 + \lambda_2)/\mu - 2. \qquad (2.4.6)
\end{aligned}$$

Since G is a finite subset of the irreducible set $\mathbb{X}^0$, we may now apply (*Kulkarni*, 1999, Theorem 3.7) to $\{Z_n\}$ to get that all states in $\mathbb{X}^0$ are positive recurrent when the right-hand side of (2.4.6) is strictly negative: when (2.4.5) holds. Furthermore, since (2.4.6) also applies for states outside of $\mathbb{X}^0$, applying Proposition C.1.5 of *Sennott* (1999) to $\{X_n\}$ yields that the expected time to reach $\mathbb{X}^0$ is finite. It follows that a stationary distribution exists cf. (*Sennott*, 1999, p. 294).

To show necessity of the inequality, we note that (*Meyn and Tweedie*, 1993, Theorem 11.5.1) implies that when $(\lambda_1 + \lambda_2)/\mu \geq 2$ the expected time to reach G from outside of G is infinite, and thus a stationary distribution cannot exist. ∎

Let U_α and w_α be the obvious infinite horizon analogues to $U_{\alpha,t}$ and $w_{\alpha,t}$, respectively. The following are called the *discounted cost optimality equations* (DCOE):

$$v_\alpha(I, y, i, j) = \min_{b \in A_\mathcal{K}} \{w_\alpha(I, y, i, j, b)\}, \qquad (2.4.7)$$

where $\mathcal{K} = \mathcal{W}, \mathcal{I}_1, \mathcal{I}_2,$ or $(0, 0, 0, 0)$ depending on (I, y, i, j). It is well-known that v_α satisfies the DCOE and a policy made up of actions that achieve the minimum on the right hand side of (2.4.7) is discounted cost optimal. Similarly, if we replace v_α in the definition of the DCOE by some function on $\mathbb{X}$, say ψ, and define U and w as the obvious average cost analogues to U_α and w_α, then the following are called the *average cost optimality equations* (ACOE):

$$g + \psi(I, y, i, j) = \min_{b \in A_{\mathcal{K}}} \{w(I, y, i, j, b)\}, \tag{2.4.8}$$

where $\mathcal{K} = \mathcal{W}, \mathcal{I}_1, \mathcal{I}_2,$ or $(0, 0, 0, 0)$ depending on (I, y, i, j).

Proposition II.5. *For the proxy model, the following hold:*

1. *For the discounted cost model*

   (a) *The quantity $v_{\alpha,t}$ is non-decreasing in t and $\lim_{t\to\infty} v_{\alpha,t} = v_\alpha$.*

   (b) *Any limit point of an optimal $t-$horizon policy is infinite horizon discounted cost optimal.*

   (c) *In particular, the results of Propositions II.2 and II.3 hold in the infinite horizon discounted cost case.*

2. *For the average cost model, suppose $(\lambda_1 + \lambda_2)\left(\frac{p_1}{\mu_1} + \frac{p_2}{\mu_2}\right) < 2$.*

   (a) *The policy that moves customers only to avoid idling has finite average cost.*

   (b) *The optimal average cost may be computed as $g = \lim_{\alpha\uparrow 1} v_\alpha(x)$ for any $x \in \mathbb{X}$.*

   (c) *Any limit point of a $\alpha-$discounted cost optimal policy is average cost optimal.*

   (d) *There exists a limit function, say ψ of $\psi_\alpha(x) = v_\alpha(x) - v_\alpha(0, 0, 0, 0)$ for $x \in \mathbb{X}$ such that (g, ψ) satisfy the ACOE.*

(e) *In particular, the results of Propositions II.2 and II.3 hold in the average cost case.*

Proof. Since the state space is countable, the cost function is non-negative, and the action set in each state finite, the first two results in the discounted cost case follow from Proposition 4.3.1 of *Sennott* (1999). Taking limits in the value functions and along a subsequence in the policies yields the last discounted cost result.

In the average cost case condition **P1'** in *Down and Lewis* (2006) holds for the non-idling policy described (see Example 3.3 in *Down and Lewis* (2006)). Applying Theorem 3.6 therein yields the first result in the average cost case. Corollary 7.5.10 and Theorems 7.2.3 and 7.5.6 of *Sennott* (1999) (collectively) yield Statements 2 (b), (c), and (d). Taking limits in the value functions and along a subsequence in the policies yields the last average cost result. ■

2.5 The Load Balancing Heuristic and Numerics

In this section we present the *load balancing* (LB) heuristic for control of the original system with heavy-tailed distributed service times. The LB heuristic requires a mapping from (partial) state information of the original system to the state of the proxy model.

Modeling this problem, it is tempting to consider the possibility of using a a partially observable Markov decision process (POMDP). For example, at any point of time t at which a customers service is not yet completed, the state space of the POMDP could include the number of customers in each queue, and the probabilities $p(t)$ and $q(t)$ that the service distribution for queues 1 and 2 are type 1. However, the continuous nature of p and q can create, once again, the potential for intractability. Instead of the POMDP, we chose to create the proxy model using a "trigger," as we will define it below.

For the original system the information of interest is the vector $(I(t), y(t), \eta_1(t), \eta_2(t))$, where $I(t)$ represents the total number of customers in the system at time t, $y(t)$ is the number of customers in the low cost queue (queue 2), and $\eta_k(t)$ denotes the time elapsed since the customer at station k began service. The classification of the customers as "standard" or "heavy" in service depends on a "trigger" denoted τ. The trigger indicates when the service time $\eta_k(t)$ for the original system is deemed long enough to treat the customer at station k as a heavy type customer. We suggest a method of determining τ below. Define

$$z_k(\eta) = \begin{cases} 1 & \text{if } \eta > \tau, \\ 2 & \text{if } \eta \leq \tau. \end{cases}$$

The original system is observed continuously and controlled at times of arrivals, times of departures, and whenever a customer in service reaches τ units of time at the station. At such times, when the original system vector is $(I(t), y(t), \eta_1(t), \eta_2(t))$, the LB heuristic uses the optimal action for the proxy model with state $(I(t), y(t), z_1(\eta_1(t)), z_2(\eta_2(t)))$, i.e., it uses the move-up-to policy with level $L_{(I(t), z_1(\eta_1(t)), z_2(\eta_2(t)))}$.

The parameters for the proxy model should be chosen to resemble those of the original system. The arrival rates equal those of the original system. If a service time for the original system, S, with distribution $F(\cdot)$, has mean $\mathbb{E}[S] = 1/\mu$, then we suggest setting the average service time of the proxy model, unconditional on the customer type (class), also equal to $1/\mu$. That is, $p_1/\mu_1 + p_2/\mu_2 = 1/\mu$. Note that stability, condition (2.4.5), is guaranteed if $(\lambda_1 + \lambda_2)/\mu < 2$. We will consider one, of perhaps many, ways to determine p_1 and p_2. Let $p_2 = F(\tau)$ and $p_1 = 1 - F(\tau)$, where τ is predetermined. Under these settings, the service times conditioned on customer type are set to have the same means: $1/\mu_2 = \int_0^\tau s\mathrm{d}F(s)/p_2$ and $1/\mu_1 = \int_\tau^\infty s\mathrm{d}F(s)/p_1$. To determine τ, we suggest a quantile-based method. For a given probability a, define

ϕ_a as the the a^{th} quantile: $\mathbb{P}(S \leq \phi_a) = a$. Define z as the probability that ϕ_a is reached given that the trigger τ has been reached: $z = \mathbb{P}(S > \phi_a | S > \tau)$. Given a and z, exogenously, τ can then be calculated from $F(\cdot)$. For the numerical study presented in Section 2.5.1, $a = 0.8$ and $z = 0.75$. Once τ is calculated, all of the parameters of the proxy model can be determined from the relationships above.

There is one last consideration to make regarding the implementation of the LB heuristic. The state space for the proxy model is (countably) infinite. Instead of solving this problem, the optimal policy is approximated using a truncated state space. That is, the policy is calculated for a system in which each queue has a maximum capacity of B customers. (For the numerical study it was assumed that arrivals to a full queue were lost, at no cost.) The truncated problem can be solved by well known algorithms, including policy iteration and value iteration; cf. *Puterman* (1994).

One question remains: How can we apply a truncated policy to the original, non-truncated system? For example, suppose that the capacity $B = 35$ but the original system reaches $(I(t), y(t), z_1(\eta_1(t)), z_2(\eta_2(t))) = (41, 3, i, j)$. When $B = 35$, there are no proxy calculations for queue lengths $(q_1, q_2) = (38, 3)$. One possibility is to simply use the actions associated with $(q_1, q_2) = (\min\{I(t) - y(t), B\}, \min\{y(t), B\})$; for the example, $(q_1, q_2) = (35, 3)$. A downside of this approach is that truncation may result in a non-monotone policy. In short, the computed move-up-to levels increase in I except near the capacity limits, where the levels may suddenly drop off. As an alternative, we suggest a "smoothing" of move-up-to levels, in accordance with Propositions II.2 and II.3. Denote the optimal number of customers to move in state (I, y, i, j), for the truncated proxy model, as $\theta^*(I, y, i, j)$; let $\theta^*(I, y, i, j) = 0$ for $(I - y) > B$ or $y > B$. Define $b^+(I, y, i, j) = y + \theta^*(I, y, i, j)$, if $\theta^*(I, y, i, j) > 0$; $b^+(I, y, i, j) = 0$, otherwise. We first approximate move-up-to levels $\tilde{L}_{(I,i,j)} = \max_y b^+(I, y, i, j)$. These levels are not monotone in I; they decrease and are zero for $I > 2B$. So, in a second step we smooth the move-up-to levels to guarantee that they are monotone and positive for a large total number of customers. For (I,i,j) we choose $\hat{L}_{(I,i,j)} = \max_{\{\ell:0\leq\ell\leq I\}} \tilde{L}_{(\ell,i,j)}$. Then, in the original system we implement a move-up-to policy with level $\hat{L}_{(I(t),z_1(\eta_1(t)),z_2(\eta_2(t)))}$. This *smoothing* approach was compared to the above *no-smoothing* approach (that ignores move-up-to levels explicitly) in the numerical study, and the smoothing approach performed better.

2.5.1 Numerical Study

We tested the LB heuristic against four heuristics: *do nothing* (DN), *no idling* (NI), *join the shortest queue* (JSQ), and *modified join the shortest queue* (ModJSQ). The DN policy never moves customers. The NI policy moves exactly 1 customer, if available, from one queue to the other if and only if the other server is idle. The JSQ policy only moves customers at times of arrival, by moving an arriving customer to the other queue if the other queue is shorter. The ModJSQ policy moves new arrivals to the other queue if the other queue is accruing total holding costs at a lower rate. That is, a customer arriving to queue 1 (2) is moved to queue 2 (1) if $h_1q_1 > h_2q_2$ $(h_1q_1 < h_2q_2)$.

In our experiment the service time S is distributed according to a *bounded* and *shifted* Pareto distribution. A standard (not bounded or shifted) Pareto distribution is a heavy-tailed, power-law distribution with two parameters, α and κ. It has support $[\kappa,\infty)$, and has infinite variance for $\alpha \leq 2$. A bounded (but not shifted) Pareto distribution is similar to a standard Pareto distribution except that it has support $[\kappa,\kappa_2)$; bounded above by κ_2. A bounded and shifted Pareto is a translation of the bounded Pareto to the origin. Its probability density function is

$$f(s) = \begin{cases} \frac{\alpha\ \kappa^\alpha}{1-(\kappa/\kappa_2)^\alpha}(s+\kappa)^{(-\alpha-1)} & \text{if } 0 \leq s \leq \kappa_2-\kappa,\ 0<\kappa\leq\kappa_2, \\ 0 & \text{otherwise.} \end{cases}$$

Its mean is $\mathbb{E}[S] = \frac{\kappa_2^{-\alpha}\left(\kappa^\alpha \kappa_2 \alpha - \kappa\ \kappa_2^\alpha \alpha\right)}{((\kappa/\kappa_2)^\alpha - 1)(\alpha-1)} - K$ and its variance is

$$\operatorname{Var}[S] = \frac{\kappa_2^{-2\alpha}\left(-\frac{((\kappa/\kappa_2)^\alpha-1)\left(\kappa^2\kappa_2^\alpha - \kappa^\alpha\kappa_2^2\right)\alpha\kappa_2^\alpha}{\alpha-2} - \frac{\left(\kappa^\alpha\kappa_2\alpha - \kappa\ \kappa_2^\alpha\alpha\right)^2}{(\alpha-1)^2}\right)}{\left((\kappa/\kappa_2)^\alpha - 1\right)^2}.$$

The design of experiment is as follows: We fix $\kappa = 0.1$ and vary the utilization $\rho := 1/\mu = \mathbb{E}[S]$ and $\operatorname{Var}[S]$, which in turn determine α and κ_2. We also fix $\lambda_1 = \lambda_2 = 1$, in which case the stability condition (2.4.5) becomes $\rho < 1$. We considered all combinations of $\rho \in \{0.5, 0.6, 0.7, 0.8, 0.85, 0.9, 0.95, 0.99\}$ and $\operatorname{Var}[S] \in \{1, 3, 6, 9, 12\}$. As noted above, for LB $a = 0.8$ and $z = 0.75$, which, together with the distribution function, determine τ, p_1, and p_2. For the costs, h_2 is fixed at 1 and $h_1 \in \{1.25, 1.5, 2\}$ and $m \in \{0.75, 1.5, 2.5\}$. In total, the factorial consists of 360 combinations of parameter settings.

The policies were evaluated under the average cost criterion. For each policy and parameter setting, a simulation was developed consisting of 60 (consecutive) runs, plus an additional run at the the beginning as a warm-up period. Each run had a length of 100,000 time units. The optimal policy for the proxy model was determined using a truncated state space with capacity $B = 35$ customers per queue. Unless stated otherwise, LB refers to the smoothing approach.

Of the 360 cases, LB policy performs best in 298 (82.8%), the NI policy performs best in 29 (8.1%), the JSQ policy performs best in 12 (3.3%), and the ModJSQ policy performs best in 21 (5.8%). The DN policy has the highest costs in every case. In terms of costs, on average over the 360 cases, LB improves upon DN by 58.9%, NI improves upon DN by 55.8%, ModJSQ improves up DN by 53.1%, and JSQ improves upon DN by 51.9%. The standard deviations (s.d.) of these percent improvements upon DN are 6.4%, 6.1%, 6.3%, and 5.8% for LB, NI, ModJSQ, and JSQ, respectively. This implies that the use of a control policy to balance the load of the system is worthwhile.

The average reduction in total costs for LB over the alternative policies are 7.0% (s.d. 6.4%) for NI, 14.5% (s.d. 9.1%) for JSQ, and 12.3% (s.d. 9.3%) for ModJSQ. In the cases for which LB is the best policy, LB averages 7.4% (s.d. 5.3%) lower costs than the best alternative. On the other hand, when LB is not the best policy, it has 6.1% (s.d. 9.1%) higher costs than the best. An important observation to make is the fact that in most of the cases that LB is outperformed the utilization is high. When $\rho = .99$ is omitted LB is the best policy in 89.5% of the cases, and when both $\rho = .95$ and $\rho = .95$ are omitted, LB is the best policy in 93.7% of the cases. As it turns out, these high utilization are subject to appreciable simulation error. We calculated 95% confidence intervals for the time average costs of a simulation run (assuming runs are independent) and found that we cannot be confident that the LB policy is actually outperformed in any of the cases with $\rho \geq 0.95$. Generally speaking, it can be problematic to simulate queues with truncated Pareto service time distributions; see *Gross et al.* (2000). Therefore, we compared simulated costs for the DN policy against the exact costs calculated using the Pollaczek-Khintchine formula. While the average absolute percent deviation between the simulated and exact DN costs is only 0.8% for $\rho \leq .9$, it is 12.0% for $\rho = .99$ and 2.2% for $\rho = .95$. In light of this, from here on we restrict attention to $\rho \leq .9$.

Averaged over the remaining 270 cases, LB reduces costs by 7.6% (s.d. 5.7%) over NI, 18.0% (s.d. 6.5%) over JSQ, and 15.4% (s.d. 5.5%) over ModJSQ. Table 2.1 displays the percent differences in costs for $\rho = .85, .9$; LB is not outperformed in **any** of these cases. The best alternative to LB is NI, and in fact NI is the only policy to outperform LB in the 270 cases – in 17 cases (6.3%); see Table 2.2. LB tends to be outperformed at the lower utilizations, when moving costs are higher, and the difference in holding cost rates between the queues is lower. As indicated in Table 2.1, for Var(S) = 1 the performances of LB and NI are closest when $h_1 = 1.25$ and $m = 2.5$; though the trend does not hold for all variances. Compared to NI, LB has lower costs by an average of 8.6% (s.d. 5.3%) when $\rho = .85$ and by 7.3% (s.d. 4.9%) when $\rho = .9$. For all 270 cases, where LB outperforms NI it does so by an average of 8.2% (s.d. 5.4%). On the other hand, when NI is better LB has only 0.8% (s.d. 0.6%) higher costs. This is true in general; the average difference in total costs between LB and the best policy is very small compared to the difference between the best policy and the other heuristics; see Table 2.3. In terms of moving costs, LB has higher moving costs than NI in 255 (94.4%) of the cases. The moving costs for LB are 56.1% (s.d. 47.4%) higher than those of NI. In the 15 cases where LB has lower moving costs, they are 2.6% (s.d. 2.5%) lower. The LB policy is more aggressive in moving customers.

We also ran simulations for the no-smoothing approach to the LB heuristic. On average, the smoothing approach outperformed the no-smoothing approach by 1.7% (s.d. 2.2%). To further test the effects of truncation, we ran simulations for the smoothing approach with queue capacity $B = 20$ and $\rho = .85, .9$. As compared to the smoothing approach when $B = 35$, the costs increased by 1.02% (s.d. 3.78%). (For $\rho = .85, .9$, the costs of the no-smoothing approach for $B = 35$ over the smoothing approach for $B = 35$ are higher by 0.7% (s.d. 1.2%).) So, truncation in the proxy model has an effect on the performance of the LB heuristic, and the smoothing approach does a better job than the no-smoothing approach at mitigating the effect.

In summary, LB (with smoothing) performs well. As indicated by the poor performance of DN, moving customers can greatly decrease holding costs. The simulation outputs for $\rho \geq .95$ are noisy and inconclusive; otherwise, LB outperforms ModJSQ and JSQ. The best alternative to LB is NI. The performances of LB and NI are closest when moving costs are high and the difference in holding cost rates is low. The NI policy can outperform LB when utilization is low, but LB is not outperformed by much and there are larger gains in the other cases. Finally, we should note that in choosing the parameters for the proxy model we only considered $a = 0.8$ and $z = 0.75$. There is room for improvement by optimizing these settings. We leave such an exercise for future research.

Table 2.1: The percent decrease in total long-run average costs for LB compared to NI, JSQ, and ModJSQ.

ρ	h_1	m	NI Var[S] 1	3	6	9	12	JSQ Var[S] 1	3	6	9	12
0.85	1.25	0.75	4.52	4.46	4.73	7.84	4.17	11.56	11.13	11.13	12.16	9.46
		1.5	3.56	3.86	2.39	5.76	2.95	12.07	12.82	7.56	9.02	9.57
		2.5	0.90	1.64	2.85	2.41	2.86	10.98	10.51	10.47	7.06	9.63
	1.5	0.75	10.45	10.96	8.90	8.02	6.64	17.19	16.71	14.60	12.56	11.75
		1.5	7.78	10.20	7.38	8.17	4.61	17.28	16.63	14.08	14.03	10.22
		2.5	4.56	6.75	4.94	5.75	5.82	15.77	14.38	12.47	12.00	11.41
	2	0.75	21.05	20.31	15.85	14.50	13.94	27.59	24.88	20.11	19.93	18.87
		1.5	16.83	16.70	15.58	15.52	14.29	23.48	23.56	21.76	19.10	19.68
		2.5	13.59	16.20	11.76	11.09	10.65	23.36	22.57	18.30	17.28	15.36
0.9	1.25	0.75	4.96	4.66	3.55	1.77	3.10	10.63	9.01	9.47	6.57	6.45
		1.5	2.88	2.59	4.80	0.42	0.36	11.02	8.22	6.80	5.96	5.80
		2.5	2.24	2.64	2.25	1.18	1.41	10.80	9.77	7.13	6.41	5.60
	1.5	0.75	10.99	9.68	7.62	5.66	7.18	15.96	14.29	9.80	9.94	10.85
		1.5	9.18	8.10	6.98	2.80	4.95	15.78	12.56	11.56	8.13	8.61
		2.5	7.32	6.73	5.45	2.72	4.48	15.61	12.42	12.14	8.88	8.75
	2	0.75	20.93	15.98	13.21	8.92	9.88	25.69	20.07	17.09	13.01	14.80
		1.5	17.92	15.28	12.63	10.33	8.87	23.75	20.15	18.06	14.39	12.01
		2.5	15.65	13.17	11.26	6.58	7.49	22.26	18.37	13.78	11.96	12.25

ρ	h_1	m	ModJSQ Var[S] 1	3	6	9	12
0.85	1.25	0.75	13.44	11.36	10.78	10.55	5.85
		1.5	11.65	12.67	8.58	8.94	7.41
		2.5	13.79	11.00	12.60	6.86	8.14
	1.5	0.75	15.88	14.38	11.72	9.14	7.56
		1.5	14.63	14.01	11.23	11.32	8.99
		2.5	15.33	15.11	9.77	10.40	7.47
	2	0.75	21.89	18.97	15.73	10.77	7.11
		1.5	19.00	17.40	14.11	10.93	10.29
		2.5	16.89	16.94	11.76	10.03	7.40
0.9	1.25	0.75	11.33	8.49	7.41	4.95	7.42
		1.5	12.32	10.95	6.18	4.50	2.66
		2.5	12.61	8.50	9.67	7.40	5.49
	1.5	0.75	16.76	13.23	9.35	6.49	7.12
		1.5	14.75	12.68	7.98	8.39	8.12
		2.5	15.55	11.73	7.42	7.61	9.12
	2	0.75	19.86	14.75	8.35	4.22	0.99
		1.5	18.71	14.28	9.93	6.29	2.40
		2.5	16.49	11.04	7.33	1.86	1.60

Table 2.2: The number of times each policy is optimal (among this group of heuristics).

ρ	# of times optimal/total LB	NI	JSQ	ModJSQ
0.5 – 0.6	79/90	11/90	0/90	0/90
0.7 – 0.8	84/90	6/90	0/90	0/90
0.85 – 0.9	90/90	0/90	0/90	0/90
0.95	29/45	7/45	3/45	6/45
0.99	16/45	5/45	9/45	15/45

Table 2.3: The average percent above optimal costs (among this group of heuristics).

ρ	LB	NI	JSQ	ModJSQ
0.5	0.1%	5.9%	30.5%	26.4%
0.6	0.1%	10.0%	29.9%	25.6%
0.7	0.1%	9.3%	24.2%	20.1%
0.8	0.0%	9.0%	19.9%	15.9%
0.85	0.0%	10.0%	18.3%	13.8%
0.9	0.0%	8.1%	14.4%	10.4%

2.6 Conclusion

In conclusion, we have introduced a new method for load balancing in the case for highly variable service distributions. The method introduced is robust to changes in the parameter settings even in the case where it is not adjusted to optimize the implementation. The most reasonable alternative to our heuristic appears to be a non-idling heuristic. In this case, the question is simply, is the consistency and savings worth the difficulty of implementing our heuristic. In many cases we believe more than 8.5% savings is worth the time to implement our heuristic.

At the same time, we have shown that the use of Markov decision processes can mitigate the challenges of a general service time distribution. We believe that the ideas described here can lead to insights for other queueing models. The example of admission controlled $M/G/1$ has already been alluded. Exactly the same intuition holds for service rate control in a $G/M/1$. Of course, these are just the building blocks for more sophisticated models. We note that an extension of the current work is to consider a larger network of queues, and we conjecture that the *two pairing* heuristics described in *Down and Lewis* (2006)) would be useful. We leave this for future research.

CHAPTER III

Managing supplier default risk via subsidies and supplier diversification

3.1 Introduction

Every firm experiences supplier default risk to different degrees, the severity of it measured by the magnitude of the default probabilities and/or of the possible losses. Different firms adopt different alternatives to manage this risk. As *Siggelkow* (2001) points out, some firms, for example Liz Claiborne, revert to supply chain strategies where they rely on multiple suppliers to deliver the parts (diversification of supplier default risk), while others, for example Ford—see *Babich* (2008)—, strive to establish a close relationship with one "dedicated" supplier and take steps to ensure the well-being of such a supplier. Measures to increase the probability of the supplier's survival include financial subsidies, guaranteed order quantities, sharing of managerial expertise, etc. The "dedicated supplier" strategy (also called single-sourcing) and the "supplier diversification" strategy (called multi-sourcing) are employed by firms for multiple reasons related to quality control, product design costs, etc. Among all of those reasons, we focus on how the supplier default risk affects the choice between these two strategies.

Furthermore, whether the firm chooses to have an exclusive supplier or to have a diversified pool of suppliers, in order to have a better control of the available supplies in the future, they may choose to have "backup inventory." We would like to analyze when the firm uses backup inventory, and if it would use more than one supplier to create the backup inventory or a single source to create it.

Both strategies (supplier diversification and backup inventory) are considered *operational* mitigation strategies. To increase the probability of successful delivery, manufacturers can also give financial subsidies to their supplier, such as the well-known case of Ford Motor Co., which agreed to pay between $1.6 billion and $1.8 billion to help with the restructuring of Visteon, as documented by *White* (2005). This ability to provide a subsidy to the suppliers can affect the decision to multi- or single-source, as well as the amount of backup inventory to create. It is not obvious which of these strategies is optimal for the manufacturer. This is one of the questions we are trying to analyze.

This problem is one of the multiple considerations to make when choosing a firm-sourcing/risk mitigation strategy. The body of research relating to the risk management problem is extensive. The following are some references that touch on the subject. *Tomlin and Wang* (2005) study the convenience of using dedicated vs. flexible resources when using an exclusive supplier vs. diversifying the suppliers. *Babich* (2006) uses game theory to analyze the value of the manufacturer's option to defer ordering and the supplier's option to defer pricing decisions when the suppliers have different lead times. *Tomlin* (2009) argues that the assumption that a firm knows exactly probability distribution for a supplier's yield (reliability) is not appropriate, so he studies the supplier diversification strategies and backup inventory strategies when the yield distribution is not known, but updated over time using a Bayesian model. *Yu et al.* (2009) suggest a method to decide on single or dual sourcing based on the disruption probability, where one of the suppliers is reliable, but more expensive, and the alternative supplier has a better price, but is less reliable. *Schmitt* *and Snyder* (2009) analyze a manufacturer that simultaneously experiences uncertainty in the supply from both random yield and multiperiod disruptions and warn that multiperiod disruptions cannot be optimized separately without incurring large errors.

A number of papers on risk mitigation appeared after September 11, 2001, an event that highlighted how global risks such as terrorism are important challenges to supply chains. These papers include the works from *Jüttner et al.* (2003), *Sheffi* (2001), *Rice and Caniato* (2003a), *Rice and Caniato* (2003b), among others.

In this work, we will consider settings where the manufacturer is a significant portion of the supplier's market, so that the manufacturer has influence on the supplier and can effectively increase the expected delivery probability. At the same time, we will consider settings in which the suppliers are in perfect competition, and therefore all of them sell the supplies they produce at the same price by supply and demand forces . Therefore, when the manufacturer is considering which supplier to select, his decision is based on the costs of subsidizing and on the expected costs of penalty/holding according to the corresponding delivery probability distribution (in other words, the price of the supplies is not a factor in the supplier selection). Under these assumptions, we can think of a large manufacturing firm with one or more dedicated suppliers.

We do not assume that the manufacturer is the only customer of the supplier, however. Therefore, the amount of inventory of the supplier is not directly linked to her probability of default.

Babich (2008) finds that the optimal subsidy and optimal ordering amount are independent of each other when the uncertainty on the supplier output comes from *random capacity*. We focus on the case when the uncertainty comes from *random yield*. Even in the simplest case (Bernoulli-distributed proportion, i.e., receive all or nothing) we find that with random yield, the subsidy and ordering amount are not independent of each other.

Our work is related to the work by *Wang et al.* (2009). We focus on the case of random yield, and we assume that the manufacturer will be in production while the subsidies are granted, even if the effects will be delayed for the next period. The formulation by Wang covers both random yield and random capacity, but for one production season with two distinct stages: On the first one, the manufacturer decides on the amount of effort (subsidy, in our context), and in the second, she decides on the amounts to order. Besides other insights, they find that a mixed strategy can be useful when the suppliers are very unreliable.

Yano and Lee (1995) present a survey on the optimal order size in the presence of random yield, so we refer the reader to it for references on the subject. We do make a note of the model by *Wang and Gerchak* (1996), which is interesting because it incorporates both random yield and random capacity in their analysis. They do not analyze subsidies.

3.2 The model

We will now present the assumptions of our model. In each subsection, we will list the assumptions that were specific for that portion of the analysis.

We start by listing the assumptions that are common to all the versions of our model.

The manufacturer decides on the supply order quantity z_k^t from supplier k. It is assumed that the manufacturer has a reasonable control of the use of the subsidies the supplier will be given. Although the subsidy is quantified in monetary amounts in the formulation, it can be the cost of any action taken by the manufacturer to improve the successful delivery probabilities. An example of this is providing coffee producers with a premium for their coffee if they meet quality and environmental standards. More generally, the subsidy can be the cost of creating "networks and value-adding partnerships", as *Bakos and Brynjolfsson* (1993) points out.

From this amount, the manufacturer will receive a random amount $\beta_k^t z_k^t$, where β_k^t is a random variable. We will assume that the density function of β_k^t is a Bernoulli distribution (which means that no partial amount of supplies will be delivered; only the full amount or nothing). The manufacturing company knows the no-default probability λ_k^t when the manufacturer does *not* subsidize supplier k at all on period t. This probability actually reflects how likely a successful delivery is today when no funding is obtained from the manufacturer. Note that the probabilities for success λ_k^t may change for each supplier and for each period. For instance, it may be more likely for a strike to happen depending on circumstances such as preexisting union agreements, etc.

From now on, we will hold the assumption that the delivered amount is Bernoulli-distributed, except where we specifically relax this assumption. We will also assume that the suppliers are different only in the probability of delivery, but the quality, the cost and price of the supplies are the same for both. These assumptions are valid, for example, when the suppliers are in different geographic locations within a country, but using similar technologies.

The demand is realized at the beginning of each period. If the demand is not met, the manufacturer pays p per unit of shortage (penalty). On the other hand, if there are leftovers in inventory after demand has been satisfied, the manufacturer will pay h per unit (holding cost). Besides the unmet demand penalty, assume that no cost is incurred by the manufacturer for supplies that are not delivered. We do not assume that the manufacturer has an insurance or that she receives compensation from the supplier if the supplies are not delivered. In either of these cases, the problem would have to be modelled using a penalty function that would depend on the shortage amount (delivered amount - current inventory). The results of this analysis would remain valid as long as this penalty function was *convex in the shortage amount.*

Without loss of generality, we do not consider separately the unit cost of supplies and consider only h and p (we have the same cost structure as the work from *Ciarallo et al.* (1994)). This formulation incorporates the ordering cost c through a standard algebraic transformation, assuming that the inventory can be salvaged at the end of the planning horizon at cost. In this case, without loss of generality, we can define $h = h' + c$ and $p = p' - c$ and solve the problem using these redefined variables.

3.2.1 Diversification and subsidy model, immediate subsidy effects, known "debt" amount

As mentioned, we will present the simplifying assumptions of this section, some of which we will relax later. First, we assume that subsidies granted by the manufacturer to the supplier have an *immediate* effect, meaning that today's subsidy increases the probability of successful delivery of the supplies for *today.* We assume that the amount needed to guarantee that a default does not happen is known and deterministic.

At the beginning of period t, the manufacturer records its inventory level x^t and observes the amount of subsidy that guarantees that supplier k does not default at the end of the *current* period, period t. This amount will be denoted as $F_k^t > 0$ and is measured in monetary units.

This amount may be, for instance, the amount that would help to reach a settlement in negotiations with a union. We will call this amount F with the generic name of "debt." Let ω_k^t be the amount of subsidy given by the manufacturer to supplier k ($\omega_k^t \leq F_k^t$). We assume that it costs α to fund either supplier with one monetary unit. In our analysis, if the cost to fund the supplier was an increasing, strictly convex function of the amount of subsidy, the objective function would still be convex, a minimum could be found, and therefore a similar analysis could be carried out. For simplicity, we will consider the funding costs to be linear.

The probability of delivery in period t depends on the amount of subsidy granted.

The more subsidy to the supplier, the more likely it is that the manufacturer will receive the delivery. For simplicity, we will assume that this probability is $\pi_k^t = \frac{\omega_k^t}{F_k^t}(1-\lambda_k^t)+\lambda_k^t$. Note that π_k^t, the after-subsidy probability of default, depends on λ_k^t, and on the amount of subsidy the manufacturer pays ω_k^t. Note that if $\omega_k^t = 0$, i.e., no subsidy is granted, the probability of successful delivery remains as λ_k^t. On the other hand, if full subsidy is granted, the probability of successful delivery is 1.

For the time being, let us assume that the demand, D^t, is deterministic. This is not too strong an assumption in the one-period context, because the manufacturer has an idea of the demand for the current period. We will relax the deterministic assumption in later versions of the model when the manufacturer considers the future.

The state of the system can be characterized by the inventory level in each period x^t. A Markov policy Π is a decision rule that assigns a pair (z_k^t, ω_k^t) to every possible state x^t. In our case, the decision will consist of the order amounts vector and of the subsidies vector: $(\overline{z^t}, \overline{\omega^t})$, where $\overline{z^t} = (z_1^t, z_2^t)$, $\overline{\omega^t} = (\omega_1^t, \omega_2^t)$.

Let ρ be the discount rate for cash flows (i.e., the inverse of the interest rate).

Using all the previous notation, the optimization problem of the manufacturer is the following:

$$\min_{\Pi} E[u^{\Pi}(1, x^1)] \tag{3.2.1}$$

where

$$\begin{aligned} u^{\Pi}(t,x) = \sum_{n=t}^{N} \rho^{n-1} \left[h\left(x + \sum_k z_k^n \beta_k^n - D^n\right)^+ + p\left(x + \sum_k z_k^n \beta_k^n - D^n\right)^+ + \alpha \sum_k \omega_k^n \right] \\ - \rho^N u_T\left[x^{N+1}\right], \end{aligned} \tag{3.2.2}$$

$$x^{t+1} = x^t + \sum_k z_k^t \beta_k^t - D^t. \tag{3.2.3}$$

Let us denote by U the value function of this finite horizon dynamic program. The value function depends on time t and on the available inventory at time t, x. The value function satisfies the following dynamic programming recursion:

$$
\begin{aligned}
U(t,x^t) &= \min_{\Pi} u^{\Pi}(t,x^t) \\
&= \min_{(z_k^t,\omega_k^t)} \left\{ E\left[h[x^t + \sum_k z_k^t \beta_k^t - D^t]^+ + p[D^t - x^t - \sum_k z_k^t \beta_k^t]^+ + \alpha \sum_k \omega_k^t \right] \right. \\
&\left. + \ \rho E U(t+1, x^t + \sum_k z_k^t \beta_k^t - D^t) \right\} \qquad (3.2.4)
\end{aligned}
$$

$U(N+1,x) = -u_T(N,x)$

where $z_k^t \geq 0$; $F_k^t \geq \omega_k^t \geq 0$ for all $k = 1,...,K$. We will assume that $u_T(\cdot) = 0$; that is, the salvage value is 0.

We will assume that both F_k^t and λ_k^t are predetermined and known for period t. In subsequent analysis, we will drop the assumption of a known amount F_k^t to guarantee delivery.

3.2.2 One period, two suppliers

In the model just described, the two-supplier one-period problem, i.e., we will assume that $K = 2$ and $T = 1$. We will omit the superscript $t = 1$ on z, D, ω_k. Recall that Π denotes a policy. From the fact that β_k has a Bernoulli distribution with parameter $\pi(\omega_k)$, the problem becomes:

$$
\begin{aligned}
\min_{\Pi} u^{\Pi}(1,x) = \min_{\Pi} & \left\{ h[x + z_1 + z_2 - D]^+ + p[D - x - z_1 - z_2]^+ \right\} \pi_1(\omega_1)\pi_2(\omega_2) \\
& + \left\{ h[x + z_1 - D]^+ + p[D - x - z_1]^+ \right\} \pi_1(\omega_1)\left[1 - \pi_2(\omega_2)\right] \\
& + \left\{ h[x + z_2 - D]^+ + p[D - x - z_2]^+ \right\} \left[1 - \pi_1(\omega_1)\right] \pi_2(\omega_2) \\
& + \left\{ h[x - D]^+ + p[D - x]^+ \right\} \left[1 - \pi(\omega_1)\right] \left[1 - \pi_2(\omega_2)\right] + \alpha(\omega_1 + \omega_2) \qquad (3.2.5)
\end{aligned}
$$

where $\pi_k(\omega_k) = \frac{\omega_k}{F_k}(1-\lambda_k)+\lambda_k$. For simplicity, where there is no ambiguity, π_k will be used instead of $\pi_k(\omega_k)$.

For fixed ω_k^t, this problem is piecewise-linear on (z_1, z_2). As a consequence, the optimal values of (z_1, z_2) are on the *corner* points. Let us assume for now that $D > x$. The four corner points that are candidates to be the optimal values are the following:

(a) $(z_1, z_2) = (D-x, 0)$,

(b) $(z_1, z_2) = (0, D-x)$,

(c) $(z_1, z_2) = (D-x, D-x)$ and

(d) $(z_1, z_2) = (0, 0)$.

Note that solution (d) is easily discarded because it is suboptimal. The comparison of the objective function (3.2.5) evaluated both in (a) and (d) makes it clear that the expected costs for solution (d) are higher.

Let $\Omega_k \equiv \{\omega_k \geq \lambda_k\}$. Ω_k is the space of the subsidies ω_k for supplier k. The manufacturer has to decide on a pair $(\omega_1, \omega_2) \in \Omega_1 \times \Omega_2$.

First, we will show that under certain conditions, regardless of the amount of subsidy granted, it is not optimal to order from both suppliers (which would imply the creation of backup inventory).

Lemma III.1. *For problem (3.2.5), the following holds:*

- *If $p(1-\lambda_1) < h(\lambda_1)$, regardless of the amount of subsidy granted to supplier 1, ordering from this supplier has a lower cost than ordering from both suppliers.*
- *If $p(1-\lambda_2) < h(\lambda_2)$, regardless of the amount of subsidy granted to supplier 2, ordering from this supplier has a lower cost than ordering from both suppliers.*

Proof. The result is obtained by comparing values of the objective function (3.2.5) at points $(z_1, z_2) = (D-x, 0)$ $(z_1, z_2) = (D-x, D-x)$. ■

In other words, it is possible that the initial parameters are such that it is never optimal to diversify for the manufacturer. This case is simpler: there is no region where diversification is optimal, so that we only need to consider procuring from supplier 1 or procuring from supplier 2; or, in other words, only points (a) and (b) described above should be considered.

However, if two conditions hold ($p(1-\lambda_1) \geq h(\lambda_1)$ and $p(1-\lambda_2) \geq h(\lambda_2)$), there are combinations of subsidies for suppliers 1 and 2 that make it optimal to diversify instead of procuring from an exclusive supplier. This means that in this case, point (c) described above should be considered besides (a) and (b). We will study this case first, because the exclusive supplier case (where either $p(1-\lambda_1) < h(\lambda_1)$ or $p(1-\lambda_2) < h(\lambda_2)$) has a simpler structure and the proofs are very similar.

As we will see in the next subsection, the amount $\widetilde{\omega_i} \equiv \frac{F_i(p(1-\lambda_i)-\lambda_i h)}{(1-\lambda_i)(p+h)}$ defines a threshold: supplier i has to have at least $\widetilde{\omega_i}$ to be reliable enough for the manufacturer. Below this level, this supplier is not reliable enough, and the expected penalty costs for the default are higher than the expected holding costs, and therefore it becomes necessary to diversify. If we look at the numerator of $\widetilde{\omega_i}$, we can see that if $p(1-\lambda_i) \leq h(\lambda_i)$, then $\widetilde{\omega_i} \leq 0$, so for any amount of subsidy ω_i, it is optimal to have i as an exclusive supplier.

3.2.2.1 The case with $p(1-\lambda_1) \geq h(\lambda_1)$ and $p(1-\lambda_2) \geq h(\lambda_2)$: Diversification may or may not be optimal

In order to analyze this case, we will divide $\Omega_1 \times \Omega_2$, the space of the subsidies, into three regions A, B and C. Each region is a subset of (Ω_1, Ω_2), where each of the possible corner points (z_1, z_2)—the order combinations (a), (b) or (c) described above—has a lower cost. In other words, we will determine constraints on (ω_1, ω_2) to define the region in which each possible pair (a), (b) or (c) is optimal. Each of these regions is the intersection of semiplanes (which we will describe below), and therefore, each of them is convex.

Then, within each region (A, B or C), by taking advantage of the simplification obtained by substituting the optimal ordering amounts (z_1, z_2) —(a), (b) or (c), respectively— we will determine which pair of subsidy values (ω_1, ω_2) within each region is the optimal one. The optimal subsidy amounts within each region are possible candidates to be the solution to problem (3.2.5). Finally, we will obtain the conditions where each of these candidate points are the solution to the problem.

We will now address the first step. Before we define the three regions in $\Omega_1 \times \Omega_2$, we will present a Lemma that will help us to obtain them. By substituting (a), (b) or (c) in the objective function (3.2.5) and comparing them, we can determine conditions where each of the combinations (a), (b) and (c) has a lower cost.

Lemma III.2. *For problem (3.2.5) the following holds:*

1. $(z_1, z_2) = (D - x, 0)$ *has a lower penalty cost than* $(z_1, z_2) = (D - x, D - x)$ *if and only if* $\omega_1 > \widetilde{\omega_1} \equiv \frac{F_1(p(1-\lambda_1)-\lambda_1 h)}{(1-\lambda_1)(p+h)}$.

2. $(z_1, z_2) = (0, D - x)$ *has a lower cost than* $(z_1, z_2) = (D - x, D - x)$ *if and only if* $\omega_2 > \widetilde{\omega_2} \equiv \frac{F_2(p(1-\lambda_2)-\lambda_2 h)}{(1-\lambda_2)(p+h)}$.

3. $(z_1, z_2) = (D - x, 0)$ *has a lower cost than* $(z_1, z_2) = (0, D - x)$ *if and only if* $\omega_2 \leq \frac{F_2(\lambda_1-\lambda_2)}{(1-\lambda_2)} + \frac{F_2(1-\lambda_1)}{F_1(1-\lambda_2)}\omega_1$

Proof. The result is obtained by comparing values of the objective function (3.2.5) at points $(z_1, z_2) = (D - x, 0)$ $(z_1, z_2) = (D - x, D - x)$. ■

Note the implications of this Lemma regarding subsidizing an exclusive supplier (without loss of generality, let us consider supplier 1): We assumed that $p(1 - \lambda_1) \geq h(\lambda_1)$. Therefore, $\widetilde{\omega_1}$ is positive. Assume that $\omega_1 < \widetilde{\omega_1}$. For example, assume that $\omega_1 = 0$ (i.e., no subsidy is granted to supplier 1). This means that the order levels $(z_1, z_2) = (D - x, 0)$ have a greater penalty cost than the order levels $(z_1, z_2) = (D - x, D - x)$. In other words, if $\omega_1 = 0$ (or any subsidy level such that $\omega_1 < \widetilde{\omega_1}$), it is suboptimal to order from supplier 1 only.

Now that we have found conditions to compare order alternatives with each other, we can find the regions where each of the possible ordering strategies is optimal (i.e., when each (z_1, z_2) dominates the other two strategies):

(a) $(z_1, z_2) = (D - x, 0)$ is optimal when $\omega_1 > \widetilde{\omega_1}$ and $\omega_2 \leq \frac{F_2(\lambda_1-\lambda_2)}{(1-\lambda_2)} + \frac{F_2(1-\lambda_1)}{F_1(1-\lambda_2)}\omega_1$ (note that we decided to drop the boundary $\omega_1 = \widetilde{\omega_1}$) from region A, which will be included in region C defined below). This is region A depicted in Figure 3.1.

(b) $(z_1, z_2) = (0, D - x)$ is optimal when $\omega_2 > \widetilde{\omega_2}$ and $\omega_2 \geq \frac{F_2(\lambda_1-\lambda_2)}{(1-\lambda_2)} + \frac{F_2(1-\lambda_1)}{F_1(1-\lambda_2)}\omega_1$ (we dropped $\omega_2 > \widetilde{\omega_2}$). This is region B depicted in Figure 3.1.

(c) $(z_1, z_2) = (D - x, D - x)$ is optimal when $\omega_1 \leq \widetilde{\omega_1}$ and $\omega_2 \leq \widetilde{\omega_2}$. This is region C depicted in Figure 3.1.

The feasible space is depicted in Figure 3.1.

Figure 3.1: Partition of the subsidy space $(\Omega_1 \times \Omega_2)$ when diversification may be optimal. In region A, order from supplier 1. Region B, order from supplier 2. Region C, order from both.

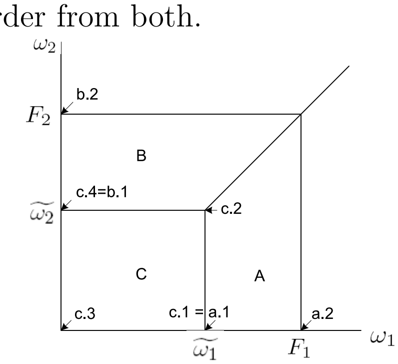

In region A, arithmetic substitution of $z_1 = D - x$ and $z_2 = 0$ into the objective function (3.2.2) will render: $u(z_k, \omega_k) = p[D - x]^+ \left[1 - \frac{\omega_1}{F_1}(1 - \lambda_1) - \lambda_1\right] + \alpha \sum \omega_k$, which is a piecewise-linear function on both ω_1 and ω_2. It will be useful to remember later that the partial derivative with respect to ω_2 is positive. Let us first consider the closure of A, which we denote $Clos(A) \equiv A \bigcup \{\text{border points of A s.t. } \omega_1 = \widetilde{\omega_1}\}$.

The set $Clos(A)$ is a convex set, so we can guarantee the minimum is reached within this set. By the definition of closure, we have that A is *contained* in $Clos(A)$. We will proceed as follows: first, we will find the $\arg\min u(\omega_1, \omega_2)$ over $Clos(A)$; next, we will show that this optimal pair of ω values is contained in A itself. By doing so, we will have accomplished obtaining $\arg\min u(\omega_1, \omega_2)$ over region A.

Let us find the $\arg\min u^{\Pi}$ over $Clos(A)$. Recall that in region A, it is optimal to order from supplier 1. On the border $\omega_1 = \widetilde{\omega_1}$ ordering from supplier 1 costs the same as diversifying. Notice that if we only order from supplier 1, it is not optimal to subsidize supplier 2. Therefore, in order to minimize the objective function with respect to (ω_1, ω_2) in $\overline{A}$, we should set $\omega_2 = 0$ and consider the corner values of ω_1.

Therefore, we have two possible solutions:

(a.1) $(z_1, z_2, \omega_1, \omega_2) = (D - x, 0, \widetilde{\omega_1}, 0)$ or

(a.2) $(z_1, z_2, \omega_1, \omega_2) = (D - x, 0, F_1, 0)$.

Similarly, in region B, from a similar argument, we get that the two possible solutions are:

(b.1) $(z_1, z_2, \omega_1, \omega_2) = (0, D - x, 0, \widetilde{\omega_2})$, or

(b.2) $(z_1, z_2, \omega_1, \omega_2) = (0, D - x, 0, F_2)$.

Finally, in region C, the objective function evaluated on $(z_1 = D - x, z_2 = D - x)$ is non-linear on ω_k. Note that in this case, we have to pay holding costs for the excess of demand $(D - x)$, given that we order this amount from both suppliers. The objective function in this case is:

$$h[D - x]^+ \pi_1(\omega_1)\pi_2(\omega_2) + p[D - x]^+ \left[1 - \pi_1(\omega_1)\right]\left[1 - \pi_2(\omega_2)\right] + \alpha\omega_1 + \alpha\omega_2.$$

In order to minimize it with respect to (ω_1, ω_2) on Region C, we need to be careful. In this case, it is possible to obtain a critical point by computing the partial derivatives, but this critical point is a *saddle point*. We depict the contour plot of the objective function on region C in Figure 3.2 and label each constraint we just defined.

Figure 3.2: Saddle point in Region C. Optimality is at the border

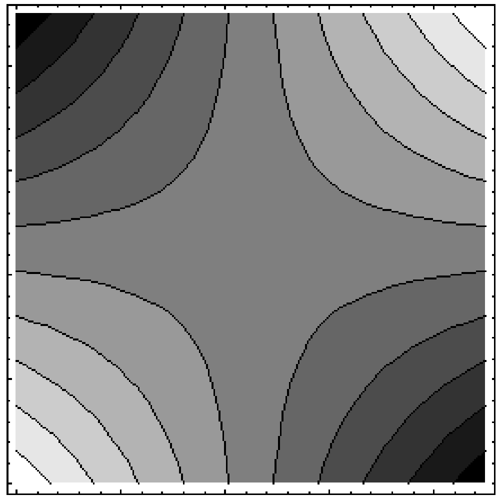

As we can see, the solution for this case is located on the border[1] of the region defined by the following constraints: $\omega_1 \geq 0$, $\omega_2 \geq 0$, $\omega_1 \leq \widetilde{\omega_1}$, $\omega_2 \leq \widetilde{\omega_2}$. The border is illustrated in Figure 3.2. *Segment (i)* corresponds to $\omega_1 = 0$ and $0 \leq \omega_2 \leq \widetilde{\omega_2}$. *Segment (ii)*, corresponds to $\omega_2 = 0$ and $0 \leq \omega_1 \leq \widetilde{\omega_1}$. *Segment (iii)*, defined by the two expressions: $\omega_1 = \widetilde{\omega_1}$ and $0 \leq \omega_2 \leq \widetilde{\omega_2}$. *Segment (iv)* corresponds to $\omega_2 = \widetilde{\omega_2}$ and $0 \leq \omega_1 \leq \widetilde{\omega_1}$. The corner points are the following:

(c.1) $(z_1, z_2, \omega_1, \omega_2) = (D - x, D - x, \widetilde{\omega_1}, 0)$, or

(c.2) $(z_1, z_2, \omega_1, \omega_2) = (D - x, D - x, \widetilde{\omega_1}, \widetilde{\omega_2})$.

(c.3) $(z_1, z_2, \omega_1, \omega_2) = (D - x, D - x, 0, 0)$,

(c.4) $(z_1, z_2, \omega_1, \omega_2) = (D - x, D - x, 0, \widetilde{\omega_2})$.

Let us analyze each of the four segments of the border of Region C individually.

[1] U is continuous on Region C, which is compact; therefore, it reaches its maximum and minimum within this region. However, if the only critical point is a saddle point, the optimum must be reached at the border.

Consider *segment (i)*, corresponding to expression $\omega_1 = 0$ and $0 \leq \omega_2 \leq \widetilde{\omega_2}$. From substitution of $z_1 = D - x$, $z_2 = D - x$, $\omega_1 = 0$, we obtain

$$h[D-x]^+ \times \pi_1(0) \times \left[\frac{\omega_2}{F_2}(1-\lambda_2)+\lambda_2\right]$$
$$+ p[D-x]^+ \times [1-\pi_1(0)] \times \left[1 - \frac{\omega_2}{F_2}(1-\lambda_2) - \lambda_2\right] + \alpha\omega_2,$$

which is a linear function of ω_2, and hence the minimum is attained at one of the corners:

(c.3) $(z_1, z_2, \omega_1, \omega_2) = (D-x, D-x, 0, 0)$,

(c.4) $(z_1, z_2, \omega_1, \omega_2) = (D-x, D-x, 0, \widetilde{\omega_2})$.

By comparing the objective function evaluated at each of these points, we obtain that point (c.3) has a lower cost than (c.4) if and only if

$$h[D-x]^+ \times \pi_1(0) \times \lambda_2 + p[D-x]^+ \times [1-\pi_1(0)] \times (1-\lambda_2)$$
$$\leq h[D-x]^+ \times \pi_1(0) \times (\Delta + \lambda_2) + p[D-x]^+ \times [1-\pi_1(0)] \times (1-\Delta-\lambda_2),$$

where $\Delta \equiv \frac{p(1-\lambda_2)-\lambda_2 h}{p+h}$.

Note that from the definition of $\widetilde{\omega_2}$, point (c.4) yields the same objective value as point (b.1).

Similarly, consider *segment (ii)*, corresponding to $\omega_2 = 0$ and $0 \leq \omega_1 \leq \widetilde{\omega_1}$. From substitution, we obtain a linear function on ω_1, and hence the minimum is attained on one of the corners: (c.3) $(z_1 = D-x, z_2 = D-x, \omega_1 = 0, \omega_2 = 0)$, (c.1) $(z_1 = D-x, z_2 = D-x, \omega_1 = \widetilde{\omega_1}, \omega_2 = 0)$.

Note that from the definition of $\widetilde{\omega_1}$, point (c.1) yields the same objective value as point (a.1).

Consider *segment (iii)*, defined by the two expressions: $\omega_1 = \widetilde{\omega_1}$ and $0 \leq \omega_2 \leq \widetilde{\omega_2}$.

Substitution of $z_1 = D - x$, $z_2 = D - x$, $\omega_1 = \widetilde{\omega_1}$ into the objective function yields

$$\begin{aligned} &h[D-x]^+ \times \pi_1(\widetilde{\omega_1}) \times \left[\frac{\omega_2}{F_2}(1-\lambda_2) + \lambda_2\right] \\ &+ p[D-x]^+ \times [1 - \pi_1(\widetilde{\omega_1})] \times \left[1 - \frac{\omega_2}{F_2}(1-\lambda_2) - \lambda_2\right] + \alpha\widetilde{\omega_1} + \alpha\omega_2. \end{aligned}$$

This is a *linear* function of ω_2, so that the minimum value of the objective function when constrained to this segment is obtained in one of the two corners:

(c.1) $(z_1, z_2, \omega_1, \omega_2) = (D - x, D - x, \widetilde{\omega_1}, 0)$, or

(c.2) $(z_1, z_2, \omega_1, \omega_2) = (D - x, D - x, \widetilde{\omega_1}, \widetilde{\omega_2})$.

In a similar way to what we found above, we find that (c.1) has a lower cost than (c.2) if and only if

$$\begin{aligned} &h[D-x]^+ \times \pi_1(\widetilde{\omega_1}) \times \lambda_2 + p[D-x]^+ \times [1 - \pi_1(\widetilde{\omega_1})] \times (1 - \lambda_2) \\ &\leq h[D-x]^+ \times \pi_1(\widetilde{\omega_1}) \times (\Delta + \lambda_2) + p[D-x]^+ \times [1 - \pi_1(\widetilde{\omega_1})] \times (1 - \Delta - \lambda_2,) \end{aligned}$$

where $\Delta \equiv \frac{p(1-\lambda_2)-\lambda_2 h}{p+h}$ and $\pi_1(\widetilde{\omega_1}) = \frac{(p(1-\lambda_1)-\lambda_1 h)}{(p+h)} + \lambda_1$.

Note that the point (c.1) gives the same objective value as point (a.1) (i.e., $u(z_1 = D - x, z_2 = 0, \omega_1 = \widetilde{\omega_1}, \omega_2 = 0) = u(z_1 = D - x, z_2 = D - x, \omega_1 = \widetilde{\omega_1}, \omega_2 = 0)$).

Similarly, consider *segment (iv),* corresponding to $\omega_2 = \widetilde{\omega_2}$ and $0 \leq \omega_1 \leq \widetilde{\omega_1}$. Substitution yields a linear function on ω_1, and hence the minimum is attained on one of the corners: (c.2) or (c.4).

We have now four points that could possibly achieve the optimum for **case (c)**. We will determine which of these four is the solution for this case.

From substitution we obtain that (c.1) is less costly than (c.2), i.e., that $u(D - x, D - x, \widetilde{\omega_1}, 0) \leq u(D - x, D - x, \widetilde{\omega_1}, \widetilde{\omega_2})$).

Similarly, (c.3) is less costly than (c.1) (i.e., $u(D - x, D - x, 0, 0) \leq u(D - x, D - x, \widetilde{\omega_1}, 0)$).

Also, (c.3) is less costly than point (c.4) too (i.e., $u(D-x, D-x, 0, 0) \leq u(D-x, D-x, 0, \widetilde{\omega_2}))$.

Then, the optimal solution for the subregion defined by case (c), is (c.3) $(z_1, z_2, \omega_1, \omega_2) = (D-x, D-x, 0, 0)$. This is fairly intuitive, because this case refers to the situation when the manufacturer orders from both suppliers, and this double ordering mitigates the risk of one of them defaulting, which in this region has a lower cost than subsidizing.

We go back now to analyze the entire feasible region (composed of the three cases we have described: **region A**, **region B** and **region C**). Note that the point (c.3) gives a lower value than the point (a.1) (given the fact that $u(z_1 = D-x, z_2 = D-x, \omega_1 = 0, \omega_2 = 0) \leq u(z_1 = D-x, z_2 = D-x, \omega_1 = \widetilde{\omega_1}, \omega_2 = 0) = u(z_1 = D-x, z_2 = 0, \omega_1 = \widetilde{\omega_1}, \omega_2 = 0))$. As a consequence, we do not need to consider (a.1) as a possible solution of problem (3.2.5).

From a similar argument, we no longer consider point (b.1) in the search of the optimal subsidy.

From this discussion, the optimal is obtained by comparing the value function given by the three points that were not eliminated by comparison:

(a.2) subsidize and order from supplier 1: $(z_1, z_2, \omega_1, \omega_2) = (D-x, 0, F_1, 0)$,

(b.2) subsidize and order from supplier 2: $(z_1, z_2, \omega_1, \omega_2) = (0, D-x, 0, F_2)$ and

(c.3) order from both suppliers, but subsidize none: $(z_1, z_2, \omega_1, \omega_2) = (D-x, D-x, 0, 0)$.

Observe that if subsidy F_i is given to supplier i, then this supplier becomes perfectly reliable, and as a consequence it is not optimal to both diversify and subsidize. Therefore, we have proved the following proposition.

Proposition III.3. *Assume that* $p(1-\lambda_1) \geq h(\lambda_1)$ *and* $p(1-\lambda_2) \geq h(\lambda_2)$*. The solution to problem (3.2.1) and the corresponding optimum value is given by the point that achieves the lowest value function* U *defined in (3.2.5), from the following three:*

a) $z_1 = D-x, z_2 = 0, \omega_1 = F_1, \omega_2 = 0;\quad u(D-x, 0, F_1, 0) = \alpha F_1,$

b) $z_1 = 0, z_2 = D-x, \omega_1 = 0, \omega_2 = F_2;\quad u(0, D-x, 0, F_2) = \alpha F_2,$

c) $z_1 = D-x, z_2 = D-x, \omega_1 = 0, \omega_2 = 0;$

$u(D-x, D-x, 0, 0) = (\lambda_2)(\lambda_1)h[D-x]^+ + (1-\lambda_2)(1-\lambda_1)p[D-x]^+.$

The cost of choosing supplier i as an exclusive supplier is the cost of providing full subsidy. We could see this as the cost of assuring certainty in the delivery of goods.

The cost of diversification is the cost of excess inventory if *no one* defaults, plus the penalty cost if *both suppliers* default. We could think of this as the cost of the uncertainty regarding the delivery, given that the manufacturer does not provide funding to any of the suppliers.

Thus, when the manufacturer chooses between diversification or subsidizing an exclusive supplier, we are indeed making a selection between certainty versus uncertainty. As a consequence, if we assume that each of the available suppliers can satisfy our total demand, and if the manufacturer prefers to avoid uncertainty, she has a greater incentive to grant subsidies to an exclusive supplier. As we will see, this trade-off can be present under different assumptions, although in this setting the trade-off is extreme in the sense that the manufacturer chooses between full diversification (backing up 100% of the order or subsidizing the full amount F).

3.2.2.2 The case with $p(1-\lambda_1) < h(\lambda_1)$ or $p(1-\lambda_2) < h(\lambda_2)$: Diversification is not optimal: An exclusive supplier must be chosen

We will proceed to analyze this case, where it is not optimal to diversify regardless of the amount of subsidy. The analysis will follow steps very similar to the case with both $p(1-\lambda_1) \geq h(\lambda_1)$ and $p(1-\lambda_2) \geq h(\lambda_2)$ as we described in the previous section, so the details of the analysis will be omitted.

In this case, since diversification is not optimal (as we proved in Lemma III.1), there will be no region C; the space of the subsidies $\Omega_1 \times \Omega_2$ will be divided only in regions A and B, where it is optimal to order $(z_1, z_2) = (D - x, 0)$ or $(z_1, z_2) = (0, D - x)$, respectively.

The following characterizes Regions A and B.

(a) $(z_1, z_2) = (D - x, 0)$ is optimal when $\omega_2 \leq \frac{F_2(\lambda_1-\lambda_2)}{(1-\lambda_2)} + \frac{F_2(1-\lambda_1)}{F_1(1-\lambda_2)}\omega_1$. This characterizes pairs (ω_1, ω_2) that belong to Region A.

(b) $(z_1, z_2) = (0, D - x)$ is optimal when $\omega_2 > \widetilde{\omega_2}$ and $\omega_2 \geq \frac{F_2(\lambda_1-\lambda_2)}{(1-\lambda_2)} + \frac{F_2(1-\lambda_1)}{F_1(1-\lambda_2)}\omega_1$. This characterizes pairs (ω_1, ω_2) that belong to Region B.

The shape of regions A and B changes slightly depending on whether $\lambda_1 \geq \lambda_2$ or $\lambda_2 < \lambda_1$, so without loss of generality, we will focus on the case where $\lambda_1 \geq \lambda_2$ and then point out the minor changes that would be needed to complete a similar analysis in the case where $\lambda_2 < \lambda_1$.

Without loss of generality, assume that $\lambda_1 \geq \lambda_2$. Figure 3.3 shows the space $\Omega_1 \times \Omega_2$ and regions A and B for this case.

Figure 3.3: Partition of the subsidy space $(\Omega_1 \times \Omega_2)$ when diversification is not optimal and $\lambda_1 \geq \lambda_2$. In region A, order from supplier 1. In Region B, order from supplier 2.

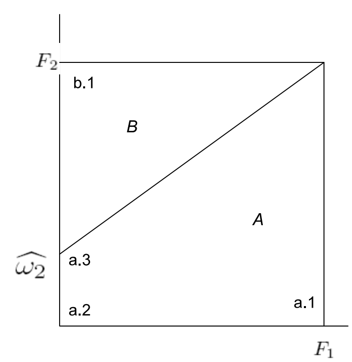

In the same way as in the previous section, it is trivial that the corner points $(D-x,0,F_1,F_2)$ and $(0,D-x,F_1,F_2)$ are suboptimal: Since we are ordering only from one supplier, it is not necessary to subsidize both. In addition, the piecewise-linear structure on (ω_1,ω_2) allows us to find the optimum within region A by focusing on the corner points:

(a.1) $(z_1,z_2,\omega_1,\omega_2)=(D-x,0,F_1,0)$.

(a.2) $(D-x,0,0,0)$.

(a.3) $(D-x,0,0,\widehat{\omega_2})$, where $\widehat{\omega_2}=F_2\frac{\lambda_1-\lambda_2}{1-\lambda_2}$.

Substitution in the objective function will show that point (a.2) has a lower cost than point (a.3), given the fact that $u(D-x,0,0,0)\leq u(D-x,0,\widehat{\omega_1},0)$, with the equality only holding when $\lambda_1=\lambda_2$.

Let us discuss why this is the case. To do so, observe that the amount $\widehat{\omega_2}$ that we subsidize to supplier 2 is the amount that makes the probability of default of both suppliers equal to supplier 1, which is the most reliable one: point (a.3) is equivalent to subsidizing supplier 2 that has a lower λ_2; we force this supplier to be as reliable as supplier 1 is before subsidies (i.e., to have a default probability equal to λ_1).

However, point (a.2) implies the same reliability (the default probability is equal to λ_1), but without any extra spending on subsidies. Therefore, we can focus on points (a.1) and (a.2) for the optimal.

Now, let us focus on region B. We discarded the point $(0,D-x,F_1,F_2)$. The other two corner points are $(0,D-x,0,F_2)$ and $(0,D-x,0,\widehat{\omega_1})$. A simple arithmetic substitution will show that $u(D-x,0,\widehat{\omega_1},0)$ and $u(D-x,0,\widehat{\omega_1},0)$ have the same value, so when we turn to the analysis of the entire region $\Omega_1\times\Omega_2$, we do not need to consider point $(D-x,0,\widehat{\omega_1},0)$. Therefore, the remaining points to consider are the following:

(a.1) $(z_1, z_2, \omega_1, \omega_2) = (D - x, 0, F_1, 0)$,

(a.2) $(D - x, 0, 0, 0)$ and

(b.1) $(0, D - x, 0, F_2)$.

To find the optimum, we need to compare the objective function evaluated in each of them.

In summary, we have proved the following:

Proposition III.4. *Assume that $p(1-\lambda_1) < h(\lambda_1)$ or $p(1-\lambda_2) < h(\lambda_2)$. In addition, without loss of generality, assume that $\lambda_2 \leq \lambda_1$. The solution to problem (3.2.1) and the corresponding optimum value is given by the point from the following three that achieves the lowest value function U defined in (3.2.5):*

a) $z_1 = D - x, z_2 = 0, \omega_1 = F_1, \omega_2 = 0$; $\quad u(D - x, 0, F_1, 0) = \alpha F_1$,

b) $z_1 = 0, z_2 = D - x, \omega_1 = 0, \omega_2 = F_2$; $\quad u(0, D - x, 0, F_2) = \alpha F_2$ *or*

c) $z_1 = D - x, z_2 = 0, \omega_1 = 0, \omega_2 = 0$; $u(D - x, 0, 0, 0) = (h[x - D]^+ + p[D - x]^+)(1 - \lambda_1)$

We have a characterization of the solution of the problem when diversification is not an option. With this, we can examine the solution, and we note that in a similar way to the case we analyzed in section 3.2.2.1 (when diversification was worth considering), the manufacturer is indeed choosing between certainty and uncertainty: certainty in this case is obtained if the manufacturer provides full subsidy. Uncertainty, however, is obtained if the manufacturer simply orders from an *exclusive* supplier without ordering any backup inventory or providing any subsidy.

In this ***one-period model setting***, with the current assumptions, when diversification occurs, the backup inventory amount is very high: the manufacturer orders *twice* the current period demand, even if there is *no future demand.* This does not mean that diversification always requires large amounts of backup inventory. The large order size is the result of the assumptions that demand is deterministic, that the probability distribution f of the percentage successfully delivered follows a Bernoulli distribution (i.e., all or nothing), and that the probability of delivery is *linear* on the subsidy.

These assumptions do not cover every situation, but they are reasonable when the demand can be predetermined by contract and when there exists the possibility of rejecting the entire lot. Although this last assumption occurs in situations such as batch processing of chemicals, rather than focusing on defining the amount of backup inventory in these simplified scenarios, we find it more pertinent to discuss qualitatively when it is more profitable to diversify, when to subsidize and when to create backup inventory to mitigate a firm's risk.

To refine the scope of this work, the issue of backup inventory and its size under different probability distributions will not be addressed here and will be left for future research. Similarly, we will also defer the case when the probability for delivery that is not linear on the subsidy. We do want to emphasize the fact diversification can occur without having to order large amounts of backup inventory. In the ***one-period*** model, when we relax the assumption that the delivery percentage follows a Bernoulli distribution (for example, instead of all or nothing, we allow a positive probability for the delivery of *half* of the ordered amount), diversification does *not* necessarily imply having a large amount of backup inventory to satisfy the demand of the current period. The amount of backup inventory depends on the percentage we expect to be delivered and on the associated probability. Similarly, when we relax the assumption of newsvendor penalty and holding costs (for example, if the penalty costs are a strictly convex function of the shortage and the holding costs are a strictly convex function of the excess inventory), diversification does not imply having a large amount of backup inventory to satisfy the demand. Full backup inventory can be needed at optimality when, besides yield uncertainty, there are supplier disruptions that can last for several periods, as Schmitt and Snyder (2008) show. They show that a high ratio $p/(p+h)$ requires full demand backup inventory. They also discuss examples where low probability of recovery from disruptions make the optimal backup inventory amount equal to 100% of the demand.

As we explained, we are interested in obtaining *qualitative* insights about the incentives to *subsidize* vs. incentives to *diversify* between both suppliers and the circumstances in which either or both strategies are combined with *backup inventory* (the size of which we do not address at this time and leave for future research). Therefore, we will focus our analysis on the simpler case where the delivery amount follows a Bernoulli distribution and where costs are of the linear, newsvendor type. Our intention is to understand the motivations for a manufacturer to choose either strategy (diversification vs. subsidizing).

Until now, we have discussed that the probability distribution of the actual delivered amount for the current period is one of the factors for diversification, *even if future periods are not considered.* Adding future periods and their inherent uncertainty gives the manufacturer even more reasons to diversify. When there is uncertainty in future demand, there are cases where it is optimal to diversify the order between different suppliers, while *at the same time* it is better to order more than today's demand, in order to have backup inventory *for the future.* In order to analyze this scenario, we drop the assumption of the model for a *single period.* We will study the case with two periods.

3.2.3 Two periods, one supplier

So far, we have analyzed the case when the manufacturer considers only immediate costs and returns and assumes that the demand is deterministic. However, the problem is not so simple in practice, given future costs and earnings. First, we analyze the consequences of considering the future when the manufacturer has only one supplier.

Let us consider the two-period, one-supplier case ($K = 1, T = 2$) for problem (3.2.1). It will be convenient to drop the subindex k that distinguishes each supplier.

For simplicity, we will assume that both λ^t and F^t are constant over t. Without loss of generality, we will assume that $\alpha = 1$. To avoid trivial solutions, we will assume that $D^2 + D^1 \geq x^1$ (meaning that we do not have inventory that exceeds the demand for both periods). Recall that Π is a decision rule that assigns a pair (z_k^t, ω_k^t) to every possible state (x^t, D^t).

The time horizon for this problem is finite; thus, it can be solved by backward induction. The optimization for the last period (period 2) is similar to the one-period problem we studied in the previous section, but the model is simpler because it has only one supplier. Therefore, the solution to the period 2 problem is the order/subsidy strategy that has the least cost between the following two:

1) $(z^2 = D^2 - x^2, \omega^2 = F^2)$, or

2) $(z^2 = D^2 - x^2, \omega^2 = 0)$

and the value function is $U(2, x) = \min\{F^2, p[D^2 - x^2]^+(1 - \lambda^2)\}$, where each element between the brackets is the cost of strategies 1) and 2), respectively.

If we substitute $U(2, x)$ in the objective function (3.2.5), we obtain the following:

$$\begin{aligned} U(1, x) = \min_{\Pi} \sum & \left\{h[x + z^1 - D^1]^+ + p[D^1 - x - z^1]^+\right\} \pi^1(\omega^1) \\ & + \left\{h[x - D^1]^+ + p[D^1 - x]^+\right\} (1 - \pi^1(\omega^1)) + \alpha\omega^1 \\ & + \min\{F^2, p[D^2 - (x + z^1 - D^1)]^+(1 - \lambda^2)\}\pi^1(\omega^1) \\ & + \min\{F^2, p[D^2 - (x - D^1)]^+(1 - \lambda^2)\}(1 - \pi^1(\omega^1)). \end{aligned} \tag{3.2.6}$$

Note that this is a piecewise-linear function on z^1, so we only need to consider the corner points: $z^1 = D^2 + D^1 - x^1$, and $z^1 = D^1 - x^1$.

In order to obtain the solution to this problem, we will first present a result that will help us.

Lemma III.5. *If the manufacturer does not order anything, the optimal subsidy is* $\omega^1 = 0$. *If the manufacturer orders a positive amount (either* $z^1 = D^1 - x^1$ *or* $D^2 + D^1 - x^1$*, given that the problem is piecewise-linear on* z^1*), the optimal subsidy amount is as follows:*

1. *If* $p[D^1 - x^1]^+(1-\lambda)/F^1 > 1 + h(D^2)(1-\lambda)/F^1$*, then* $\omega^1 = F^1$.

2. *If* $p[D^1 - x^1]^+(1-\lambda)/F^1 \leq 1$*, then* $\omega^1 = 0$.

3. *If* $1 < p[D^1 - x^1]^+(1-\lambda)/F^1 \leq 1 + h(D^2)(1-\lambda)/F^1$.

   - *For* $z^1 = D^2 + D^1 - x^1$*, then* $\omega^1 = 0$
   - *For* $z^1 = D^1 - x^1$*, then* $\omega^1 = F^1$

Proof. Note that the problem is linear on ω^1. The result is obtained by substituting the order size $z^1 = 0$, $z^1 = D^1 - x^1$ or $z^1 = D^2 + D^1 - x^1$ in the objective function u in the minimization problem 3.2.6, and by then computing the derivative $\partial u/\partial \omega^1$. If the derivative is positive, it is optimal to provide full subsidy. Otherwise, it is optimal to not provide any subsidy at all. ∎

Proposition III.6. *The solution to problem (3.2.6) is the following:*

1. $z^1 = D^1 - x^1$ *and* $\omega^1 = F^1$ *if*

   (a) $p[D^1 - x^1]^+(1-\lambda) > F^1 + hD^2(1-\lambda)$ *and*

   $hD^2 > \min\{F^2, pD^2(1-\lambda)\}$,

   OR

   (b) $F^1 < p[D^1 - x^1]^+(1-\lambda) < F^1 + hD^2(1-\lambda)$ *and*

   $\min\{F^2, pD^2(1-\lambda)\} < hD^2\lambda + p[D^1 - x^1]^+(1-\lambda)$.

2. $z^1 = D^1 - x^1$ *and* $\omega^1 = 0$ *if*

   $p[D^1 - x^1]^+(1-\lambda) < F^1$ *and*

   $\min\{F^2, pD^2(1-\lambda)\}\lambda + \min\{F^2, p[D^2 + D^1 - x^1]^+(1-\lambda)\}(1-\lambda) < hD^2\lambda.$

3. $z^1 = D^2 + D^1 - x^1$ *and* $\omega^1 = 0$ *if*

   (a) $p[D^1 - x^1]^+(1-\lambda) < F^1$ *and*

   $hD^2\lambda < \min\{F^2, pD^2(1-\lambda)\}\lambda + \min\{F^2, p[D^2 + D^1 - x^1]^+(1-\lambda)\}(1-\lambda),$

   OR

   (b) $F^1 < p[D^1 - x^1]^+(1-\lambda) < F^1 + hD^2(1-\lambda)$ *and*

   $hD^2\lambda + p[D^1 - x^1]^+(1-\lambda) < \min\{F^2, pD^2(1-\lambda)\}.$

4. $z^1 = D^2 + D^1 - x^1$ *and* $\omega^1 = F^1$ *if*

   $p[D^1 - x^1]^+(1-\lambda) > F^1 + hD^2(1-\lambda)$ *and*

   $\min\{F^2, pD^2(1-\lambda)\} > hD^2.$

Proof. In the same fashion as the one-period two-supplier model, we start by taking the first derivative of the objective function with respect to ω^1. We obtain the solution by arithmetic comparison, while keeping in mind that $\partial B/\partial\omega^1$ is positive and by paying attention to which of the terms in $\min\{F^2, p[D^2 - (x^1 + z^1 - D^1)]^+(1-\lambda^2)\}$ actually achieves the minimum. As an example, we will derive the result for cases 1a) and 4) of the proposition. The remaining cases are derived in a similar manner. From Lemma III.5, we have that if $p[D^1 - x^1]^+(1-\lambda) > F^1 + hD^2(1-\lambda)$, the optimal subsidy is $\omega_1 = F_1$. In order to determine the optimal amount, we substitute $z_1 = D_1 - x_1$ and $z_1 = D_2 + D_1 - x_1$, and $\omega_1 = F_1$ in the objective function (3.2.5), from which we obtain the following:
If $hD^2 > \min\{F^2, pD^2(1-\lambda)\}$, then it is optimal to order $z_1 = D_1 - x_1$. Otherwise, it is optimal to order $z_1 = D_2 + D_1 - x_1$. Putting these conditions together gives case 1a) and case 4) of the proposition. The remaining cases can be proved in a similar

manner. ■

In order to get insights into the optimal strategy, we will now take a look at Figure 3.4, which is a graph of the penalty costs for both periods vs. the holding costs.

We will assume for this graph that the subsidies required in the second period do not exceed the penalty of not meeting the demand for that period, i.e., $F^2 < p[D^2]^+(1-\lambda)$. In practical terms, this means that it is more expensive to quit by not doing anything to satisfy the demand in the last period than to guarantee that the demand will be covered by subsidizing.

Figure 3.4: Penalty vs. Holding costs—Each region defines the strategy(ies): Ordering for one or both periods and/or subsidizing. (Graphical representation of Proposition III.6)

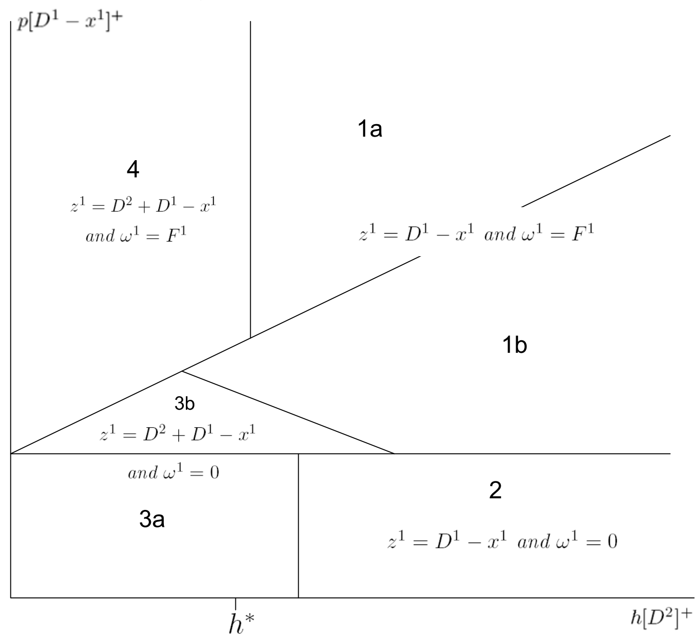

Consider the case when the holding costs are such as the cost marked with h^* in Figure 3.4. Suppose that the penalty cost is extremely low (Region 3a). In this case, the manufacturer wants to have some diversification of the risk concerning the supplies she needs for the next period: she can order them now, pay the holding cost, and if the supplier defaults, she can have one more opportunity to order them again. As the penalty cost increases, the manufacturer transits into Region 3b, keeping the order size and subsidy constant. If the penalty cost continues to rise, the manufacturer will want to actually fund the supplier instead of risking to have to pay the penalty, and she will want to order only for the current period, because she will want to save the holding cost for ordering for the second period too (Region 1b). Finally, if the penalty cost is very high (region 4), the manufacturer will opt for the most protection she can get against the penalty: full funding in period 1, and ordering for both periods now and saving the future funding costs.

As we can see, even if there is only one supplier, the time horizon gives the manufacturer some flexibility. Depending on the required subsidy amount, the holding costs and the shortage penalties, the supplier has the option to subsidize, to have backup inventory or to wait and see if subsidy will be required to satisfy the accumulated demand in period 2.

Note, however, that this flexibility depends on the fact that the subsidy is assumed to have an immediate effect on the current default probability. In practice, unfortunately, it is not always the case that a subsidy can instantly eliminate the reasons for potential default. In the following sections, we will analyze the case when the subsidy does not produce an immediate effect and helps to avoid default with a delay, i.e., in the *second* period. We will also relax the assumption of deterministic demand in later sections.

3.3 Diversification and subsidy model, subsidy with delayed effects, random assets

In the previous section, we assumed that the effects of subsidizing a supplier was immediate, meaning that any subsidy granted in period t would change the default probability for period t. However, in reality, it may take some time for the subsidy to change the default probability of a supplier. In order to consider the consequences of a delay in the effects of the probability of a production default, we will change some of the assumptions of the model. We will no longer consider the deterministic amount of debt F (which represented the amount to guarantee that no default would happen in future settings) in this and future sections.

Instead of focusing on the amount of debt, we will consider the supplier's level of *assets* $c^t > 0$ at the end of the previous period (before any random effects from the current period happen). Similar to the previous model, for the manufacturer an attempt to increase the level of either supplier assets by a unit costs α. In this setting, we will assume that each supplier's assets are affected by a random shock ε_k^t in each period. The random shock distribution f_k^t is stationary, i.e., $f_k^t = f_k$, which means that the probability distribution does not change over time.

In this context, the manufacturer once again decides on the order quantity z_k^t, depending on the current inventory x^t, and will receive $\beta_k^t z_k^t$. The percentage delivered β_k^t is Bernoulli-distributed with parameter $\pi_k(c_k^t)$ (probability of successful delivery, i.e., $\pi_k(c_k^t) = P[\beta_k^t = 1]$). We assume that for all c, we have $0 < \pi_k(c) < 1$, and that π_k is strictly increasing and strictly concave (if we assume that π is twice differentiable, these assumptions translate as $\pi_k'(c) > 0$ and $\pi_k''(c) < 0$). The strictly increasing assumption means that the higher the level of assets in period $t-1$, the lower the probability of the supplier's default in period t.

To derive the closed-form solution of the problem, we will assume that π_k is twice differentiable. This closed-form solution will contain thresholds and critical points that depend on the inverse and the derivatives of a function of π_k. The closed-form solution is useful, because it allows us to see how the thresholds move when different parameters vary and therefore allow us to explore different scenarios. It is important to note that although the arguments of π changed from the last section and that we no longer give a specific functional form for it, π continues to refer to the probability of successful delivery.

The manufacturer also decides on the amount of subsidy to be given to each supplier. Let us denote by θ_k^t the **intended *level of assets*** that the manufacturer desires supplier k to have in the next period; that is to say, the actual amount of subsidy delivered by the manufacturer is $(\theta_k^t - c_k^t)$. We will call θ_k^t *the subsidy-up-to amount*, which highlights the fact that θ_k^t is similar in nature to the order-up-to amount that is often used in inventory problems. We chose to use the subsidy-up-to level to simplify the notation in the two-period problem, but we should note that it is equivalent to finding the actual amount of subsidy $(\theta_k^t - c_k^t)$.

The supplier assets c_k^{t+1} at the beginning of the next period (period $t+1$) depend not only on the subsidy received today, period t, but also on the random shock of the period: $c_k^{t+1} = \theta_k^t + \varepsilon_k^t$. This means that the subsidy provided in period t will *not* affect the probability of default in the *current* period, but it will affect the default probability in the *following* period, $t+1$. In the previous section, we solved the problems in terms of ω_k, the **net** subsidy given to the supplier. Here we solve the problem in terms of θ_k, the *subsidy-up-to amount* (net subsidy$= \theta_k^t - c_k^t$). The change in notation is necessary since we are no longer using the *net* subsidy ω as we did in previous sections.

3.3.1 Deterministic demand, one supplier, two periods, newsvendor costs

For the time being, we assume that the demand, D^t, is deterministic and that it is backlogged if necessary. Given the fact that there is only one supplier, we drop the subindex k for the time being.

The manufacturer wishes to minimize the costs coming from procurement and subsidies granted to the supplier.

The state of the system can be characterized by the inventory level in each period and by the level of assets of the period, i.e., (x^t, c^t). A Markov policy Π is a decision rule that assigns a pair (z_k^t, θ_k^t) to every possible state (x^t, c^t).

The manufacturer's problem can be stated as follows:

$$\min_{\Pi} E[v^{\Pi}(1, x^1, c^1)] \tag{3.3.1}$$

where

$$\begin{aligned} v^{t,\Pi}(t, x^t, c^t) = \sum_{n=t}^{N} \rho^{n-1} \left[h[x^n + z^n\beta^n - D^n]^+ + p[D^n - x^n - z^n\beta^n]^+ + \alpha(\theta^n - c^n) \right] \\ - \rho^N v_T \left[x^{N+1}, c^{N+1} \right], \end{aligned} \tag{3.3.2}$$

$$\beta^t \sim Bernoulli(\ \pi^t(c^t)\ ), \qquad \varepsilon^t \sim f.$$

and the transition rules are:

$$\begin{aligned} c^{t+1} &= \theta^t + \varepsilon^t, \\ x^{t+1} &= x^t + z^t\beta^t - D^t. \end{aligned} \tag{3.3.3}$$

We will assume a salvage value of 0; that is, $v_T(\cdot) = 0$.

Denote by V the value function. The value function satisfies the following dynamic programming recursion:

$$\begin{aligned} V(t, x^t, c^t) = \min_{(z_k^t, \theta_k^t)} E \Bigg[h[x^t + \sum_k z_k^t\beta_k^t - D^t]^+ + p[D^t - x^t - \sum_k z_k^t\beta_k^t]^+ + \sum_k \alpha(\theta_k^t - c_k^t) \\ + \rho EV(t+1, x + \sum_k z_k^t\beta_k^t - D^t, \theta^t + \varepsilon^t). \Bigg] \end{aligned} \tag{3.3.4}$$

$$V(N+1, x, c) = 0$$

Then,

$$V(t, x, c) = \min_{\Pi} v(t, x, c). \tag{3.3.5}$$

Let us consider the solution to the two-period problem; i.e., we will assume that $N = 2$.

Considering that β^t has a Bernoulli distribution, and rewriting it in terms of $y^t = x^t + z^t$ (order-up-to level before demand) and of $c^{t+1} = \theta^t + \varepsilon^t$ (next period's assets), the problem becomes:

$$\begin{aligned} V(t, x^t, c^t) = \min_{\Pi} & \left\{ h[y^t - D^t]^+ + p[D^t - y^t]^+ \right\} \pi(c^t) \\ & + \left\{ h[x^t - D^t]^+ + p[D^t - x^t]^+ \right\} (1 - \pi(c^t)) + \alpha(\theta^t - c^t) \\ & + \rho \, E_{\varepsilon^t} V(t+1, y^t - D^t, c^{t+1}) \pi(c^t) \\ & + \rho \, E_{\varepsilon^t} V(t+1, x^t - D^t, c^{t+1})(1 - \pi(c^t)). \end{aligned} \tag{3.3.6}$$

In a similar fashion to what we have done in previous sections, this problem can be solved by backwards induction, so we first present the last period's solution.

Proposition III.7. *For period 2, the solution to the problem is to not subsidize and to order up to D^2. In other words, for period 2, the optimal solution is $z^2 = [D^2 - x^2]^+$, $\theta^2 = c^2$. Therefore, the value function becomes*

$$V(2, x^2, c^2) = \left[1 - \pi(c^2)\right] p[D^2 - x^2]^+ + h[x^2 - D^2]^+.$$

Proof. First, we note that the objective function for the last period is

$$
\begin{aligned}
V(2,x^2,c^2) &= \min_{(z^2,\theta^2)\in\mathcal{A}(t)} \left\{h[x^2+z^2-D^2]^+ + p[D^2-x^2-z^2]^+\right\}\pi(c^2) \\
&\quad + \left\{h[x^2-D^2]^+ + p[D^2-x^2]^+\right\}(1-\pi(c^2)) + \alpha(\theta^2-c^2) + 0 \\
&= \min_{(y^2,\theta^2)\in\mathcal{A}(t)} \left\{h[y^2-D^2]^+ + p[D^2-y^2]^+\right\}\pi(c^2) \\
&\quad + \left\{h[x^2-D^2]^+ + p[D^2-x^2]^+\right\}(1-\pi(c^2)) + \alpha(\theta^2-c^2).
\end{aligned}
\tag{3.3.7}
$$

As we can see, this equation is linear on θ^2, and since α is positive, the optimal action is to reduce θ^2 as much as possible. Also, we can see that the optimal order-up-to amount y is D^2, given the piecewise-linear structure on y. This implies that $z^2 = [D^2 - x^2]^+$. Substitution of z in V will result in the following:

- $x^2 \leq D^2$ (and hence $z^2 = D^2 - x^2$)

$$
\begin{aligned}
V(2,x^2,c^2) &= (p[0]^+ + h[0]^+)\pi(c^2) + (p[D^2-x^2]^+ + h[x^2-D^2]^+)(1-\pi(c^2)) \\
&= p[D^2-x^2]^+\left[1-\pi(c^2)\right].
\end{aligned}
$$

- $x^2 \geq D^2$ (and hence $z^2 = 0$)

$$
\begin{aligned}
V(2,x^2,c^2) &= (p[0]^+ + h[x^2-D^2]^+)\pi(c^2) + (p[0]^+ + h[x^2-D^2]^+)\left[1-\pi(c^2)\right] \\
&= h[x^2-D^2]^+.
\end{aligned}
$$

The term that multiplies h is the same in both cases. Also, the term that multiplies p is positive *only* when $x^2 \leq D^2$. As a consequence, we can write a simplified version of V composing both cases: $V(2,x^2,c^2) = [1-\pi(c_2)]\, p[D^2-x^2]^+ + h[x^2-D^2]^+$. ∎

Note the implications of this simplified expression $V(2,x^2,c^2) = (1-\pi(c^2))p[D^2 - x^2]^+ + h[x^2-D^2]^+$. V has a portion of holding costs and a portion of penalty costs. Assume that h and p are of similar magnitude (meaning that the ratio h/p is not too small, like in the order of $1-\pi(c^2)$). Observe that the factor $1-\pi(c^2) < 1$, which is the probability of default, is only multiplying p and not h. For a manufacturer to choose a supplier, it is reasonable to assume that this probability of default is relatively low, and it will be lower when subsidizing. The probability of default (which in practice is not a large amount very often), $1-\pi(c^2) < 1$, is only multiplying p; we can see that the total holding costs will be a large part of V, and the total penalty costs will be a small part of $V(2, x^2, c^2)$. As a consequence, when we substitute $V(2, x^2, c^2)$ into the two-period model, we observe that the holding costs have a much greater impact in the cost-to-go, due to the fact that we will choose θ_1 such that we make $(1-\pi(c^2))$ *as low as possible*, where $c^2 = \theta_1 + \varepsilon^1$.

Now that we have rewritten the value function for the last period, we can rewrite the original two-period problem in a way that will allow us to analyze the optimal decision for the first period. Substitution of the cost-to-go on period 2 yields:

$$
\begin{aligned}
V(1, x^1, c^1) = & \min_{(z,\theta)\in\mathcal{A}(t)} v^{\Pi}(1, x^1, c^1) \\
= & \min_{(z,\theta)\in\mathcal{A}(t)} \left(\left\{h[y^1 - D^1]^+ + p[D^1 - y^1]^+\right\}\pi(c^1)\right. \\
& + \left\{h[x^1 - D^1]^+ + p[D^1 - x^1]^+\right\}\left[1-\pi(c^1)\right] \\
& + \alpha(\theta^1 - c^1) + \rho\left\{E_{\varepsilon^1}\left[1 - \pi(\theta^1 + \varepsilon^1)\right] p[D^2 + D^1 - y^1]^+\right. \\
& + h[y^1 - D^2 - D^1]^+\}\pi(c^1) + \rho\left\{E_{\varepsilon^1}\left[1 - \pi(\theta^1 + \varepsilon^1)\right] p[D^2 + D^1 - x^1]^+\right. \\
& \left. + h[x^1 - D^2 - D^1]^+\}(1 - \pi(c^1)\right).
\end{aligned}
$$

From this expression, we note that if we fix the value of θ^1, V is convex on y^1. Therefore, for each fixed value of θ^1 and c^2, there is an optimal policy for z^1 of *order-up-to* type.

Before proceeding with the analysis, we will define the notation and establish some properties that will be useful later, as they will simplify our calculations. Define function g as follows:

$$
g(\theta^1) \equiv E_{\varepsilon^1}\left[1 - \pi(\theta^1 + \varepsilon^1)\right].
$$

Function g is the probability of default when the level of assets is θ^1 (recall that the probability of default depends both on the subsidy *and* on the random shock that affects the assets).

We will also establish the following lemma.

Lemma III.8. *g is monotonically decreasing (and hence invertible) and convex. In addition, if π is twice differentiable, g is twice differentiable too.*

Proof. $\pi(\cdot)$ is increasing and concave. Therefore, $1-\pi(\cdot)$ is decreasing and convex. From the convexity of π, $1-\pi(\lambda\theta_a^1+\lambda\epsilon^1+(1-\lambda)\theta_b^1+(1-\lambda)\epsilon^1) \leq \lambda[1-\pi(\theta_a^1+\epsilon^1)] + (1-\lambda)[1-\pi(\theta_b^1+\epsilon^1]$. Taking expectation on both sides renders the convexity of g. A similar argument holds for the monotonicity of g. The twice differentiable property of g follows from the definition and the properties of the expected value operator if π is twice differentiable. ∎

In other words, the expected probability of default decreases with the amount of subsidy granted, and it is convex on that same variable.

Now that we have defined function g and established some of its properties, we return to the problem of finding the solution for the problem for period 1. Note the piecewise-linear structure of the objective function (3.3.1) on y. As a consequence, the order amount, z, should be a corner point. This means that the order amount in the optimal solution is one of the following: *(a)* $z^1 = [D^1 - x^1]^+$ (order enough to reach an inventory level of D^1), *(b)* $z^1 = [D^1 + D^2 - x^1]^+$ (order enough to reach an inventory level of $D^1 + D^2$), *(c)* $z^1 = 0$ or *(d)* $z^1 = \infty$. Note that we can immediately discard option (d). The reason is that the positive holding costs would make the value function go to infinity when $z^1 = \infty$. Next, we show that, to achieve optimality, we must at least order enough to cover the demand in the first period. Given the fact that the problem is piecewise-linear on y^1 (the order-up-to level), we only need to focus on the corner points. In this case, it suffices to show that the corner point $z = 0$

is suboptimal.

Proposition III.9. *A policy that orders $z^1 = [D^1 - x^1]^+$ performs better than a policy that orders $z^1 = 0$ for every initial inventory x^1. In other words, the former has a lower or equal cost.*

Proof. We need to prove that ordering $[D^1 - x^1]^+$ outperforms ordering 0 for every possible initial inventory x^1. Note that for any fixed θ^1, the policies give a different value of the objective function v only when $x^1 < D^1$. Let Π be a policy such that the order is $z^1_\Pi = [D^1 - x^1]^+$ and the subsidy is θ^1. Let Π' be a policy such that the order is $z^1_{\Pi'} = 0$ and the subsidy is θ^1 as well. Hence, the result will be proved if the inequality $v^\Pi(1, x^1, c) \le v^{\Pi'}(1, x^1, c)$ holds for $x^1 < D^1$. Substituting and using the fact that $D^1 > x^1$:

$$\begin{aligned} &v^\Pi(1, x^1, c) - v^{\Pi'}(1, x^1, c) \\ &= \rho\ g(\theta^1)p(D^2)\pi(c^1) - \left[p(D^1 - x^1)\pi(c^1) - \rho\ g(\theta^1)p(D^2 + D^1 - x^1)\pi(c^1)\right]. \end{aligned}$$

Canceling terms and rearranging, we get that $v^\Pi(1, x^1, c) - v^{\Pi'}(1, x^1, c) \le 0$ holds if $1 + \rho\ g(\theta^1) \ge 0$. It can easily be verified that this inequality holds for all θ by doing the arithmetical computation for g and observing that we assumed $0 < \pi(\cdot) < 1$. Hence, we have proved that the objective function achieves a lower value when ordering $z^1 = [D^1 - x^1]^+$. ∎

We have simplified the problem, where the two possible options are ordering either *(a)* $z^1 = [D^1 - x^1]^+$ or *(b)* $z^1 = [D^1 + D^2 - x^1]^+$. Which of these two possible options is less costly? The answer depends on the amount of subsidy θ^1 that we need to grant the supplier in each of those two cases. We will now state a theorem that determines for which region of θ^1 it is better to order up to D^1 and vice versa. In order to do so, let us define

$$\widehat{\theta^1} \equiv g^{-1}\left(\frac{h}{\rho p}\right). \tag{3.3.8}$$

Recall that $g(\theta^1) = E_{\varepsilon^1}[1 - \pi(\theta^1 + \varepsilon^1)]$. Note that g depends on the distribution of ϵ^1.

In the next proposition, we will show that $\widehat{\theta^1}$ defines a threshold that determines how much we should order: for today only or for today and tomorrow. Besides, as we will see later, $\widehat{\theta^1}$ determines a threshold that defines different cases (regions) for the solution of the problem.

Proposition III.10. *Assume we fix the subsidy-up-to value θ^1.*

a) If $\theta^1 \geq \widehat{\theta^1}$, it is better to order $z^1 = [D^1 - x^1]^+$.

b) If $\theta^1 \leq \widehat{\theta^1}$, it is better to order $z^1 = [D^1 + D^2 - x^1]^+$.

Proof. In a similar fashion to what we did for the previous proposition, the result is obtained from arithmetic substitution of the order amounts $z^1 = [D^1 - x^1]^+$ and $z^1 = [D^1 + D^2 - x^1]^+$ in the objective function in each of the regions $[\widehat{\theta^1}, \infty)$ (for part a) and $(-\infty, \widehat{\theta^1}]$ (for part b). Then, we compare them.

Let us start with part a) of the result. In order to be able to easily compare the objective function values for each of the policies, it is easier to analyze different cases defined by the level of the initial inventory x^1.

Let us prove part a) for the case when $x^1 < D^1$. Let us prove that $z^1 = [D^1 - x^1]^+$ has a lower cost than $z^1 = [D^1 + D^2 - x^1]^+$, i.e., that $v^{\Pi}(1, x, c) < v^{\Pi'}(1, x, c)$, where policy Π is such that the order is $z^1 = [D^1 - x^1]^+$, and policy Π' is such that $z^1 = [D^1 + D^2 - x^1]^+$. After substitution, and using the fact that $x^1 < D^1$, we get that

$$v^{\Pi}(1, x, c) - v^{\Pi'}(1, x, c) = \rho\, g(\theta^1)p(D^2)\pi(c^1) - \left[h(D^2)\pi(c^1)\right].$$

Solving for θ, we find that to guarantee that $v^{\Pi}(1, x, c) - v^{\Pi'}(1, x, c) < 0$, the following condition is both necessary and sufficient:

$$\theta > \widehat{\theta^1}. \tag{3.3.9}$$

Similarly, one can prove that the same condition guarantees $v^{\Pi}(1,x,c) - v^{\Pi'}(1,x,c), < 0$, for the case $D^1 < x^1 < D^1 + D^2$. This fact, together with the fact that the feasible region is $\theta^t \in [c^t, \infty)$, concludes the proof for part a of the result.

Part b) is proven using an analogous two-part argument. ∎

From what we have seen, the optimal order amount varies according to the level of the subsidy: after a certain threshold $\widehat{\theta^1} = g^{-1}\left(\frac{h}{\rho p}\right)$, the manufacturer will order $[D^1 - x^1]^+$. The intuition here is clear: once the manufacturer has given enough subsidy to the supplier, there is enough confidence that the supply will be delivered next period, and therefore it is not necessary to take the extra precaution of ordering for the future today. In other words, $\widehat{\theta^1}$ is actually the threshold that defines a *strong* supplier, in the sense that we can trust that the supplier will deliver the supplies, and therefore, there is no need to add backup inventory for the future.

The previous proposition gives insights on the solution, and it will be helpful in providing a complete solution to problem (3.3.1). We still need to find which level of subsidy would be optimal in each region: $\theta \leq \widehat{\theta^1}$ or $\theta \geq \widehat{\theta^1}$.

In other words, to complete the solution of the optimization problem, we can divide the feasible region in two disjointed portions (as defined by $\widehat{\theta^1}$) and solve the following two subproblems:

Subproblem a): Determine the optimal θ^1 within the region $\theta^1 \leq \widehat{\theta^1}$, the region where it is optimal to order $z^1 = [D^1 + D^2 - x^1]^+$, and

Subproblem b): Determine the optimal amount θ^1 within the region $\theta^1 \geq \widehat{\theta^1}$, the region where it is optimal to order $z^1 = [D^1 - x^1]^+$.

Each of these problems is solved by computing the first- and second-order conditions for the subsidy.

The following proposition deals with the local solution to the *unconstrained* version of *subproblem a)*. This solution to the unconstrained problem may not coincide with the solution of the *constrained* version (the optimal may not be in the region $\theta^1 \le \widehat{\theta^1}$). However, we will use this solution as a tool to deal with problem (3.3.1) by considering the location of $\widehat{\theta^1}$ and its consequences.

Proposition III.11. *Suppose that we order $z^1 = [D^1 + D^2 - x^1]^+$. If $D^1 + D^2 \ge x^1$, the optimal amount of subsidy in this case is*

$$\theta^{1,D1+D2} \equiv (g')^{-1}\left[\frac{-\alpha}{\rho p(D^2 + D^1 - x^1)\left[1 - \pi(c^1)\right]}\right]. \tag{3.3.10}$$

Otherwise, the optimal amount of subsidy is $\theta = 0$.

Proof. First, we should note that when we order $z^1 = [D^1+D^2-x^1]^+$, the objective function 3.3.2 becomes

$$\begin{aligned} v(1, x^1, c^1) = h(D^2)\ I_{[D^1+D^2>x^1]}\pi(c^1) + (p[D^1 - x^1]^+ \\ +h[x^1 - D^1]^+ I_{[D^1<x^1<D^1+D^2]})(1 - \pi(c^1)) + \alpha(\theta^1 - c^1) \\ +\rho h[x^1 - D^1 - D^2]^+ + \rho\ g(\theta^1) p[D^1 + D^2 - x^1]^+(1 - \pi(c^1)). \end{aligned} \tag{3.3.11}$$

We will compute the first derivative and obtain the critical point. Later, we will prove that V is convex on θ^1, and therefore the critical point will attain the minimum. The first derivative is as follows:

$$\frac{\partial v}{\partial \theta} = \alpha + \rho\ g'\left(\theta^1\right) p[D^2 + D^1 - x^1]^+(1 - \pi(c^1)).$$

If $D^2 + D^1 > x^1$, we can solve for θ^1 in the first-order condition equation $\partial v/\partial\theta^1 = 0$. We obtain that the critical point is $\theta^{1,D1+D2} \equiv (g')^{-1}\left[\frac{-\alpha}{\rho p[D^2+D^1-x^1]^+(1-\pi(c^1))}\right]$. It remains to be shown that v is a convex function on θ^1. This can be accomplished by obtaining the second derivative with respect to θ^1:

$$\frac{\partial^2 v}{\partial \theta^2} = \rho\, g''\left(\theta^1\right) p[D^2 + D^1 - x^1]^+(1 - \pi(c^1)).$$

From the previous proposition about g, we know that $g'' \geq 0$, which proves the first part of the result.

To conclude the proof, observe that if $D^2 + D^1 < x^1$, the derivative above simplifies to

$$\partial v / \partial \theta = \alpha,$$

and therefore, given the fact that $\alpha > 0$, the optimal amount of subsidy is 0. ■

We have a similar result, which solves the *unconstrained* version of *subproblem b)* in a similar fashion to the proposition above, and hence we omit the details in the proof. Once again, this is the unconstrained version of the subproblem, and therefore the solution may not be in the feasible region $\theta \geq \widehat{\theta^1}$, but, eventually, we will use this proposition *and the location of* $\widehat{\theta^1}$ to obtain the solution of problem (3.3.1).

Proposition III.12. *Suppose that we order* $z^1 = [D^1 - x^1]^+$*. If* $D^1 + D^2 \geq x^1$*, the optimal amount of subsidy is*

$$\theta^{1,D1} \equiv (g')^{-1}\left[\frac{-\alpha}{\rho p(D^2 + (1 - \pi(c^1))(D^1 - x^1))}\right]. \tag{3.3.12}$$

If $D^1 + D^2 \leq x^1$*, the optimal amount of subsidy is* $\theta = 0$*.*

Proof. In a similar fashion to the previous result, substituting the ordered amount $z^1 = [D^1 - x^1]^+$, and using the fact that $[D^1 + D^2 - x^1]^+ = D^2\ I_{[x^1 < D^1 + D^2]} + (D^1 -$

$x^1)\ I_{[x^1<D^1+D^2]}$, the objective function becomes:

$$\begin{aligned} v(1,x^1,c^1) &= h(x^1-D^1)+p[D^1-x^1]^+\ (1-\pi(c^1))+\alpha(\theta^1-c^1) \\ &\quad +\rho\ g(\theta^1)p\ D^2 I_{[x^1<D^1+D^2]}+\rho\pi(c^1)\ h[x^1-D^1-D^2]^+ \\ &\quad +\rho\ g(\theta^1)\ p\ (D^1-x^1)(1-\pi(c^1))I_{[x^1<D^1+D^2]} \\ &\quad +\rho\ h[x^1-D^1-D^2]^+(1-\pi(c^1)). \end{aligned}$$

Solving for θ^1 in the first-order condition $\partial v/\partial\theta^1 = 0$, and verifying the second derivative to check the convexity, proves the result when $D^1+D^2 \geq x^1$.

The second part of the result, when $D^1+D^2 \leq x^1$, follows trivially. This is because the derivative above simplifies to $\partial v/\partial\theta = \alpha$. ∎

We have obtained the optimal solutions for both cases (when we order either $[D^2+D^1-x^1]^+$ or $[D^1-x^1]^+$). An arithmetic calculation, and the fact that g is decreasing, show that $\theta^{1,D1+D2} \leq \theta^{1,D1}$. The intuition behind this inequality is straightforward: If we are not willing to provide much subsidy to decrease the default probability for tomorrow, we should take provisions and build some inventory today to compensate for the risk. Alternatively, we can think that when ordering for the future today, we are not as concerned with the possibility of default. Therefore, subsidizing today to guarantee delivery tomorrow is not as attractive.

Note as well that while there is an order relationship between $\theta^{1,D1+D2}$ and $\theta^{1,D1}$, there is no obvious order relationship between these two quantities and $\widehat{\theta^1}$, the threshold that defines whether it is better to order for the current period only, or for the present and the future. We have now defined two regions for the subsidy and the ordering amount, so by comparing the optimal solutions for each region, we can obtain the solution to (3.3.1). In other words, to obtain the optimal solution, we will need to use the solutions to both *subproblems a)* and *b)*. The solution to (3.3.1) actually depends on the location of $\widehat{\theta^1}$. As we can see in Figures 3.5, 3.6 and 3.7, there are three general cases. In each of them, the objective function graphs of the two subproblems intersect at the point $\widehat{\theta^1}$. In each of them, $\theta^{1,D1+D2}$ is the arg min of the subproblem from Proposition III.11, while $\theta^{1,D1+D2}$ is the arg min of the subproblem from Proposition III.12.

Figure 3.5: Case 1: The threshold $\widehat{\theta^1}$ is to the left of the optimal subsidies for order sizes $[D^1 - x]^+$ and $[D^1 + D^2 - x]^+$

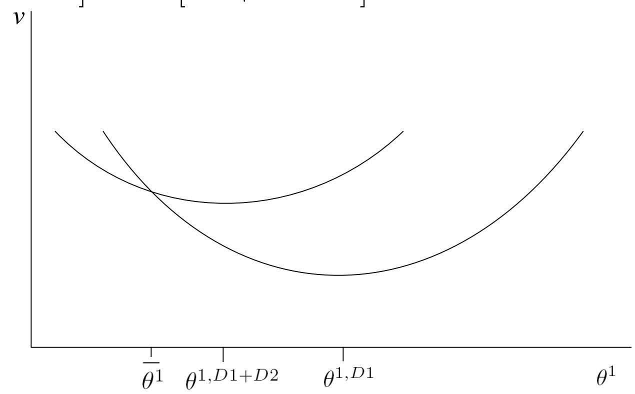

Figure 3.6: Case 2: The threshold $\widehat{\theta^1}$ is in the middle of the optimal subsidies for order sizes $[D^1 - x]^+$ and $[D^1 + D^2 - x]^+$

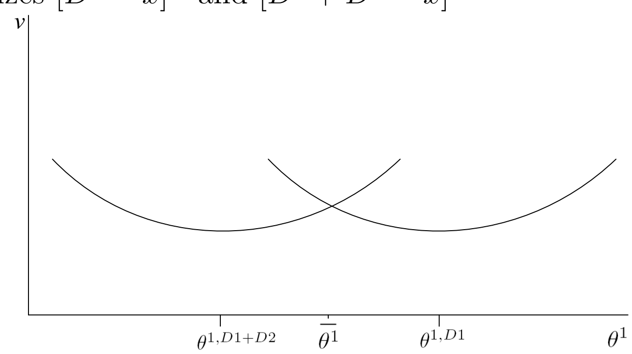

Note that given the fact that $\theta^{1,D1+D2} \leq \theta^{1,D1}$, the three cases we depict in Figures 3.5, 3.6, 3.7 are all the possible ones. We find the solution by analyzing each of them.

Case 1) refers to the case when the point at the borders A and B, $\widehat{\theta^1}$, is to the left of $\theta^{1,D1+D2}$.

Case 2) refers to the case when the border point $\widehat{\theta^1}$ is between $\theta^{1,D1+D2}$ and $\theta^{1,D1}$.

Figure 3.7: Case 3: The threshold $\widehat{\theta^1}$ is to the right of the optimal subsidies for order sizes $[D^1 - x]^+$ and $[D^1 + D^2 - x]^+$

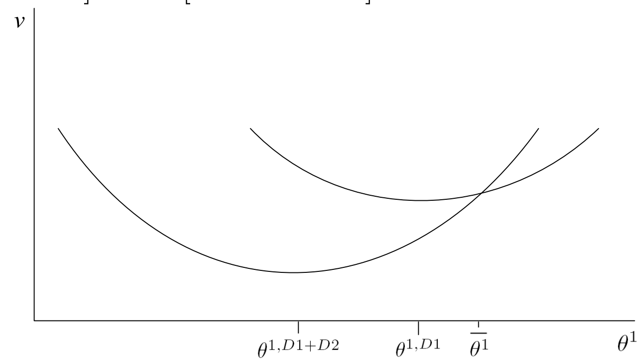

Case 3) refers to the case when $\widehat{\theta^1}$ is to the right of $\theta^{1,D1}$. We will need to compare the value functions of the two problems (III.12) and (III.11), and having these characterizations will be necessary in order to do so.

Let us define a *strategy* as a combination of subsidy and level after supply delivery: (y^1, θ^1).

These cases will be useful to characterize the solutions. We refer the reader to Proposition A.1 in Appendix A for the statement and proof of the solution.

Let us discuss the intuition behind the solution to each of the cases.

Case 1): $\widehat{\theta^1}$ is bounded above by $\theta^{1,D1+D2}$. We know that $\widehat{\theta^1} = g^{-1}(\frac{h}{\rho p})$ and that g^{-1} is decreasing. Therefore, for this case h, is large and ρp is small. In this case, we order less and subsidize more, given the fact that the holding costs are higher.

Case 2): $\widehat{\theta^1} = g^{-1}(\frac{h}{\rho p})$ is between $\theta^{1,D1+D2}$ and $\theta^{1,D1+D2}$, which in practical terms means that the ratio $\frac{h}{\rho p}$ is neither too large nor too small. In this case we have to compare both alternatives: ordering for both periods with a low subsidy and ordering for period 1 only with a higher subsidy.

Case 3): $\widehat{\theta^1}$ is bounded below by $\theta^{1,D1}$. In this case, we know that h is small and ρp is large. Therefore, there is more incentive to order more and subsidize less (case 3a), or, if the current assets are high enough, we will not need more subsidy to guarantee future delivery, nor need to order for the future.

3.3.1.1 Sensitivity analysis

In this section, we analyze how changes in some parameter values affect the optimal solution to problem (3.3.1) and discuss the economic implications of these changes.

To do the sensitivity analysis, we should study how $\theta^{1,D1+D2}$, $\theta^{1,D1}$ and $\widehat{\theta^1}$ are affected when the exogenous variables change. For clarity, we reproduce expressions (3.3.8), (3.3.10) and (3.3.12):

$$\widehat{\theta^1} = g^{-1}\left[\frac{h}{\rho p}\right].$$

$$\theta^{1,D1+D2} \equiv (g')^{-1}\left[\frac{-\alpha}{\rho p(D^2 + D^1 - x^1)(1 - \pi(c^1))}\right].$$

$$\theta^{1,D1} \equiv (g')^{-1}\left[\frac{-\alpha}{\rho p(D^2 + (1 - \pi(c^1))(D^1 - x^1))}\right].$$

Note that, given the fact that g is decreasing and convex, g' is increasing. This fact is useful in deriving all the results below.

Changes in demand for period 1

Assume that the demand in period 1 increases. As we mentioned before, the amounts of subsidy $\theta^{1,D1+D2}$ and $\theta^{1,D1}$ may be at either side of $\widehat{\theta^1}$. From the expressions above, we note that in this case $\theta^{1,D1+D2}$ and $\theta^{1,D1}$ increase, but $\widehat{\theta^1}$ does not change, so $\theta^{1,D1+D2}$ and $\theta^{1,D1}$ may actually shift from being to the left of $\widehat{\theta^1}$ to the right.

Assume for now that optimization problem l can be characterized at the beginning by Case 3. Suppose that the demand increases. Given the fact that the order-up-to amount depends on D^1, at the beginning the order-up-to will increase. This is not the end of the story, though. If the subsidies $\theta^{1,D1+D2}$ and $\theta^{1,D1}$ increase, it may (or may not) happen that the problem can no longer be characterized by Case 3 and will be characterized by Case 2 instead. If the increase in demand is very large, it may even be possible that the problem characterization changes from Case 3 to Case 1. A change of this nature would imply that the optimum amount to order will change too. Changing from Case 3 to Case 1, say, implies that the order-up-to amount will jump down after initially increasing.

From this discussion, and from Proposition A.1, when D^1 increases, the solution to problem (3.3.1) will always change by increasing the subsidy awarded to the supplier. See Figure 3.8.

Figure 3.8: Sensitivity of the Subsidy θ^1 vs. Demand D^1

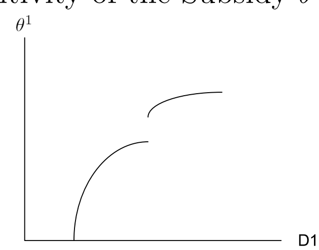

Note that the order-up-to amount y^1 will not necessarily change monotonically; say that the problem can be characterized by Case 3. We know that the manufacturer has to order to cover the demand for period 1 (plus the demand in period 2). In order to meet the new demand for period 1, the order-up-to amount must increase. However, as we just discussed, the subsidy level also has to increase as a backup measure (if the order made today is not delivered, the subsidy will increase the probability of successful delivery *tomorrow*). If the increase in the subsidy is big enough, the level of assets of the supplier may be large enough to become very dependable for future deliveries (i.e., the chances of default tomorrow will be low). In that case, the manufacturer no longer needs to order for both periods to assure that there will be enough supplies to satisfy future demand. As a consequence, at some point the characterization changes from Case 3 to Case 1, and the order-up-to size changes from $D^1 + D^2$ to D^1. Therefore, the changes in the order-up-to size are not monotonic.

An illustration of the changes is in Figure 3.9.

Figure 3.9: Sensitivity of Order-up-to Quantity y^1 vs. Demand D^1

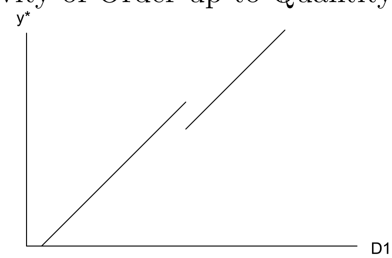

The changes in the order-up-to amount are inherited by the actual order size, $|y^1 - x^1|$, so the ordered amount is not monotonic. See Figure 3.10.

Figure 3.10: Sensitivity of Order Quantity $y^1 - x^1$ vs. Demand D^1

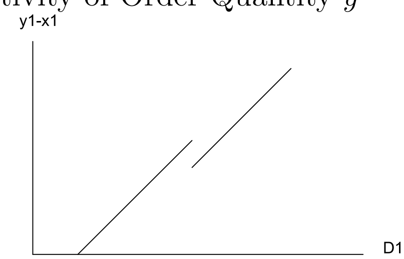

The changes we discussed above were based on the possibility of changing from Case 3 to Case 2. Similarly, it is possible that before the demand increase, the problem could be characterized by Case 2, but after the change, the problem characterization changes to Case 1, and the directions of the changes would be the same as discussed above.

Note that it could never happen that as a result of a demand increase the problem characterization changes from Case 2 to Case 3, or from Case 1 to Cases 2 or 3 (recall the direction of the changes in $\theta^{1,D1+D2}$ and $\theta^{1,D1}$), so that the manufacturer will never switch from ordering for period 1 to ordering for both periods as a result of an increase in demand. The increase in the order size will be of the same magnitude of the increase in the demand for period 1.

Demand in period 2

Let us consider now what happens when the demand in period 2 increases. From expressions (3.3.8), (3.3.10) and (3.3.12), we get that, in a similar way to what happened with increases in demand in the first period, $\theta^{1,D1+D2}$ and $\theta^{1,D1}$ increase, but $\widehat{\theta}^1$ does not change, and therefore the direction of the changes in the subsidy will be the same as discussed in the case of increases of demand for period 1. See Figure 3.11.

Figure 3.11: Sensitivity of Subsidy θ^1 vs. Demand D^2

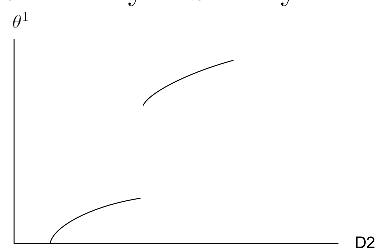

Changes in the order-up-to amount have a different nature, though. In the case of changes in demand in period 1, the order-up-to amount would always adjust to reflect the increased demand in period 1. However, for changes in demand in period 2, the order-up-to amount will adjust *only* if the problem characterization is given by Case 3 or in Case 2aii. If the problem is characterized by Case 1, the subsidy will change, but the order-up-to will remain constant (in Case 1 the order size covers only the first period). When increasing D^2 enough, we reach a threshold to the right of the graph (when the problem is finally characterized by Case 1) where the order-up-to is constant. At this point, it becomes optimal to order only for period 1, rather than for both. Note that the actual order size $y^1 - D^1$ inherits this pattern. The changes are illustrated in Figures 3.12 and 3.13.

We should note that the order size is more sensitive to changes in demand today than to changes in demand tomorrow, the reason being that ordering today is not the only source for the supplies we need tomorrow.

Figure 3.12: Sensitivity of Order-up-to Quantity y^1 vs. Demand D^2

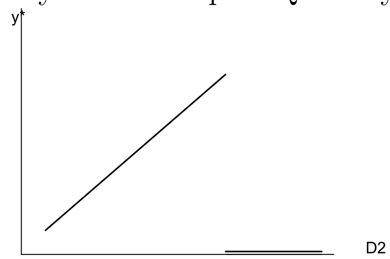

Figure 3.13: Sensitivity of Order Quantity $y^1 - x^1$ vs. Demand D^2

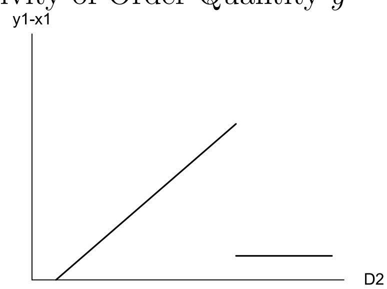

Changes in supplier assets

Let us analyze now what happens when the initial asset level of the supplier c^1 increases. From expressions (3.3.8), (3.3.10) and (3.3.12), we see that $\theta^{1,D1+D2}$ and $\theta^{1,D1}$ decrease, but $\widehat{\theta^1}$ is unchanged.

Before proceeding, note that Case 3 has two subcases, Subcase i and Subcase ii A.1. Subcase i corresponds to lower levels of c^1 ($c^1 \leq \widehat{\theta^1}$), and subcase ii corresponds to higher levels ($c^1 > \widehat{\theta^1}$). The shifts in $\theta^{1,D1+D2}$ and $\theta^{1,D1}$ imply that the characterization may change, for example, from Case 1 to Case 3a and then to Case 3b. In that situation, at the beginning, the manufacturer only orders supplies for the first period. Note from (3.3.12) how the optimal subsidy is inversely proportional to the expected supplies to be received; as a result, when c^1 increases a little, the amount of subsidy awarded will decrease. However, when the decrease is considerable and the supplier becomes less reliable, it will be necessary to increase the order-up-to amount to cover the future demand too. Given the fact that the order-up-to amount increased, the actual ordered amount increases accordingly. If c^1 continues growing after this, the optimal subsidy-up-to level θ^1 will have to grow once the constraint $\theta^1 \geq c^1$ becomes binding. Also, if c^1 gets high enough, the supplier will be reliable enough even before subsidies, so the order-up-to amount will decrease again to $D^1 - x^1$. The changes in the optimal subsidy, in the order-up-to and in the actual ordered amount, are illustrated in the graphs in Figures (3.14), (3.15) and (3.16).

Figure 3.14: Sensitivity of Subsidy θ^1 vs. Assets c^1

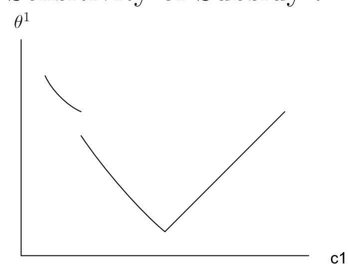

Figure 3.15: Sensitivity of Order-up-to Quantity y^1 vs. Assets c^1

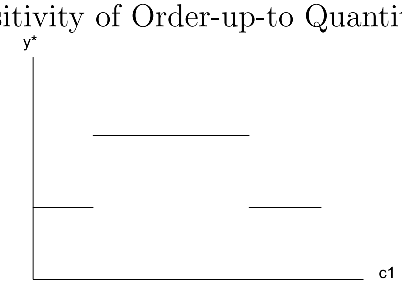

Figure 3.16: Sensitivity of Order Quantity $y^1 - x^1$ vs. Assets c^1

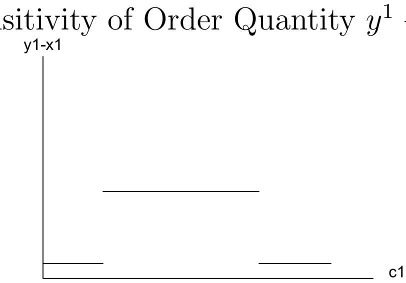

Changes in the cost of funding the supplier (α)

When the cost of funding the supplier α increases, $\theta^{1,D1+D2}$ and $\theta^{1,D1}$ decrease, but $\widehat{\theta^1}$ is unchanged. In such a situation, for instance, we can switch from Case 1

Figure 3.17: Sensitivity of Subsidy Quantity θ^1 vs. Cost of Funding α

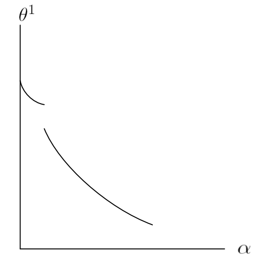

to Case 3a and then to Case 3b. In that situation, we switch from ordering-up-to D^1 only to ordering-up-to $D^1 + D^2$. At the same time, at the beginning, the subsidy decreases slowly when α increases, and it jumps farther down when we change the order-up-to amount. See Figures 3.17, 3.18 and 3.19.

Figure 3.18: Sensitivity of Order-up-to Quantity y^1 vs. Cost of Funding α

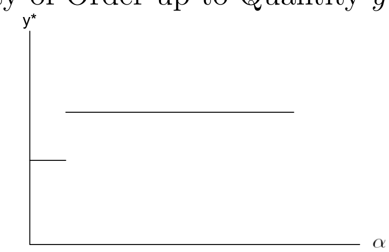

Figure 3.19: Sensitivity of Order Quantity $y^1 - x^1$ vs. Cost of Funding α

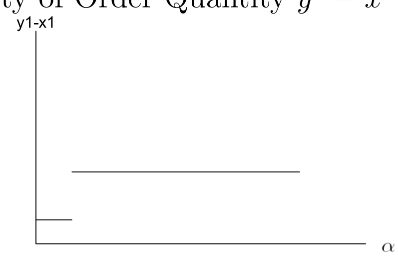

Changes in initial inventory

In this case, $\theta^{1,D1+D2}$ and $\theta^{1,D1}$ decrease, but $\widehat{\theta^1}$ does not change, so $\theta^{1,D1+D2}$ and $\theta^{1,D1}$ may shift from being to the right of $\widehat{\theta^1}$ to the left.

Assume that before the demand increase, the parameters are such that the optimization problem l can be characterized by Case 1. If $\theta^{1,D1+D2}$ and $\theta^{1,D1}$ increase, it may happen that the problem can no longer be characterized by Case 1 and can be characterized instead by Cases 2 or 3. The change of the case that characterizes the problem implies that the optimum amount to order will change too. Changing from Case 1 to Case 3, say, implies that the order-up-to amount will increase from ordering for the current period to ordering for both periods. On the other hand, the actual ordered amount depends both on the order-up-to level and on x^1, and therefore the ordered amount is not monotonic: it decreases as x^1 increases. Then, when the order-up-to level increases from D^1 to $D^1 + D^2$, the actual ordered amount jumps up too, and then it starts to decrease again. See Figures 3.20 and 3.21.

Figure 3.20: Sensitivity of Order-up-to Quantity y^1 vs. Initial Inventory x^1

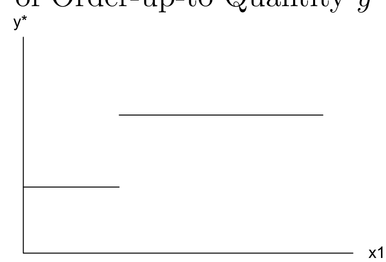

Figure 3.21: Sensitivity of Order Quantity $y^1 - x^1$ vs. Initial Inventory x^1

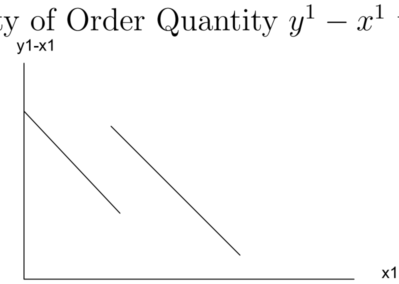

Finally, given the fact that both $\theta^{1,D1+D2}$ and $\theta^{1,D1}$ shift to the left, we can conclude that when x^1 increases, the solution to problem (3.3.1) will change by decreasing the subsidy awarded to the supplier. See Figure 3.22.

Figure 3.22: Sensitivity of Subsidy θ^1 vs. Initial Inventory x^1

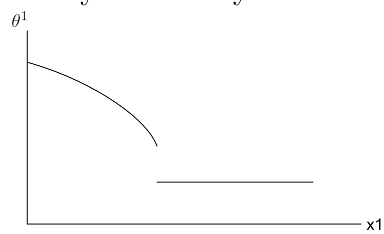

Intuitively, the order-up-to level will increase with the initial inventory: If the inventory level is low to start with, it is necessary to order at least for period 1, as it would be expensive to order from both. However, if the initial inventory is a little higher, we can afford to order for both periods. In addition, if the initial inventory is higher, the need for a subsidy to guarantee delivery is less.

Changes in the holding cost

If h increases, $\theta^{1,D1+D2}$ and $\theta^{1,D1}$ remain unchanged, while $\widehat{\theta}^1$ decreases. In this case, the characterization of problem (3.3.1) can change from Case 3 to Case 2, from Case 3 to Case 1 or from Case 2 to Case 1, but not the other way around. Therefore, the optimum solution will change by increasing the subsidy awarded to the supplier and possibly decreasing the order-up-to amount. The changes are illustrated in the following graphs. See Figures 3.23, 3.24 and 3.25.

Figure 3.23: Sensitivity of Subsidy θ^1 vs. Holding Cost h

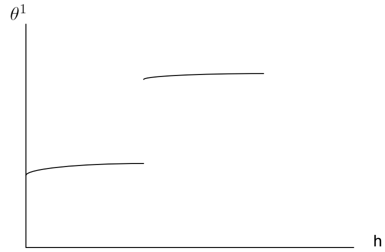

Figure 3.24: Sensitivity of Order-up-to Quantity y^1 vs. Holding Cost h

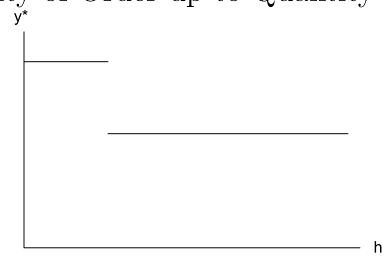

Figure 3.25: Sensitivity of Order Quantity $y^1 - x^1$ vs. Holding Cost h

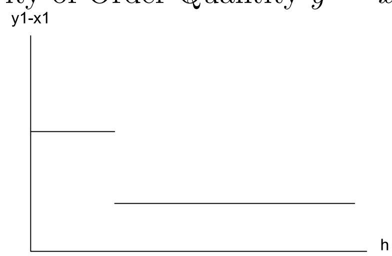

Intuitively, what happens is that the threshold $\widehat{\theta}^1$ is lower as a reaction to the increase in the holding cost. Given the fact that it is more expensive to hold inventory, lower levels of subsidy will allow the ordering of today's demand only.

Changes in the discount rate or in the penalty

If either the rate of discount ρ or the penalty p decreases, $\theta^{1,D1+D2}$, $\theta^{1,D1}$ and $\widehat{\theta}^1$ increase. Note that in this case, $\widehat{\theta}^1$ increases as well, and as a result we cannot be sure of how the characterization changes (and it is possible that, say, it changes from Case 2 to Case 1 or to Case 3). As we pointed out before, a change in the characterization (Case 1, Case 2 or Case 3) of the problem implies changes in the optimum amount to order. In the current situation, we do not know if Case 2 changes to 1 or to 3, say, so we cannot conclude if the ordering amount will increase or decrease. We are sure, however, that if the optimal ordering amount does not change, the subsidy will increase (since $\theta^{1,D1+D2}$ and $\theta^{1,D1}$ increase when ρ increases). The reason that we cannot determine the direction of the changes is that it depends on the current levels of D^1, D^2, α, c^1 and x^1 **and** on the shape of the function π, which in turn determines the convexity of g.

3.3.2 Stochastic demand (period 2), one supplier

Let us analyze the case when the demand in period 2 is stochastic. We will first derive results that come from the analysis of the algebraic formulation. We will see that when we relax the assumption of the demand being deterministic, the algebraic treatment we did on the deterministic case cannot be carried out in the same way, given the fact that the objective function is no longer piecewise-linear on y^1. We will provide the system of equations that implicitly characterize the solution.

A quick reminder of the notation follows: x^t is the level of inventory at the beginning of period t. The manufacturer observes the supplier's level of assets $c^t > 0$ at the beginning of the current period. It costs α to raise one unit to subsidize a supplier. The supplier's assets are affected by a random shock ε^t with a probability distribution f. The manufacturer orders z^t and will receive $\beta^t z^t$ at the end of the period, where β^t is a random variable with a Bernoulli distribution with parameter $\pi(c^t)$. The supplier's assets at the beginning of the next period are $c^{t+1} = \theta^t + \varepsilon^t$. The demand is realized at the beginning of each period. In period 1, the demand for period 2 is unknown to the manufacturer. Before the beginning of period 1 of the planning horizon, the demand levels for both periods are independently distributed.

Assume that the demand for period 2, $D^2 > 0$, is realized at the beginning of period 2, and therefore it is unknown at the beginning of period 1, when we model it as a stochastic variable, distributed according to a probability mass function ϕ (or probability density function, if it is a continuous random variable). Demand is backlogged. The manufacturer pays p per unit of shortage and h per unit for storage for unmet and excess demand, respectively.

The state of the system can be characterized by the inventory level in each period and by the level of assets of the period, i.e., (x^t, c^t). A Markov policy Π is a decision rule that assigns a pair (z_k^t, θ_k^t) to every possible state (x^t, c^t).

The optimization problem with the stochastic demand is the following:

$$\min_{\Pi} E[v^{\Pi}(1, x^1, c^1)] \tag{3.3.13}$$

where

$$\begin{aligned} v(t, x^t, c^t) = &\sum_{t=1}^{2} \rho^{t-1} \left[h[x^t + z^t\beta^t - D^t]^+ + p[D^t - x^t - z^t\beta^t]^+ + \alpha(\theta^t - c^t) \right] \\ &- \rho^2 v_T \left[x^3, c^3 \right], \end{aligned}$$

$$\beta^t \sim Bernoulli(\ \pi^t(c^t)\ ),$$

$$\varepsilon^t \sim f,$$

$$D^t \sim \phi,$$

and the transition rules are:

$$c^{t+1} = \theta^t + \varepsilon^t$$

$$x^{t+1} = x^t + z^t\beta^t - D^t \tag{3.3.14}$$

We will assume that $v_3(\cdot) = 0$

Denote by V the value function. The value function satisfies the following dynamic programming recursion:

$$\begin{aligned} V(t, x^t, c^t) = &\min_{(z,\theta)\in\mathcal{A}(t)} \left[h[x^t + z\beta^t - D^t]^+ + p[D^t - x^t - z\beta^t]^+ + \alpha(\theta - c^t) \right] + \\ &\rho^T E_{\beta^t, \varepsilon^t, D^{t+1}} V(t+1, x + z\beta^t - D^t, c^{t+1}) \end{aligned}$$

$$V(3,x,c) = -v_3(x,c)$$

where optimization region $\mathcal{A}(t)$ is given by

$$z \geq 0; \tag{3.3.15a}$$

$$\theta \geq c \tag{3.3.15b}$$

We will assume that $v_3(x,c) = 0$, i.e., that the salvage value is 0. Considering that β^t has a Bernoulli distribution, and rewriting in terms of $y^t = x^t + z^t$ (order-up-to level) and of $c^{t+1} = \theta^t + \varepsilon^t$ (future capacity), the problem can be rewritten as follows:

$$\begin{aligned} V(t,x^t,c^t) = \min_{(y,\theta)\in\mathcal{A}(t)} & \Big[\big\{h[y-D^t]^+ + p[D^t-y]^+\big\}\,\pi(c^t) \\ & + \big\{h[x^t-D^t]^+ + p[D^t-x^t]^+\big\}\,(1-\pi(c^t)) + \alpha(\theta-c^t) \\ & + \rho\ E_{\varepsilon^t,D^t}V(t+1,y-D^t,c^{t+1})\pi(c^t) \\ & +\rho\ E_{\varepsilon^t,D^t}V(t+1,x^t-D^t,c^{t+1},D^{t+1})(1-\pi(c^t))\Big] \end{aligned} \tag{3.3.16}$$

This expression looks very similar to the case with deterministic demand, and the first question that comes into mind is how close this solution is to the deterministic solution, and if there is a rule of thumb to guess about the solution to the stochastic case based on the solution to the deterministic case of the following type: "the optimal order-up-to will be higher/lower in the stochastic case than in the deterministic case."

By solving several instances of the stochastic problem with the expected demand for period 2 equal to the deterministic demand, we created a plot (see Figures 3.26 and 3.27) that permits the comparison of the optimal order-up-to for different variance levels with the optimal order-up-to in the deterministic case. Similarly, we can compare the optimal subsidy in the stochastic case for different variance levels with the optimal subsidy in the deterministic case. What we found in the plots is that the order-up-to level in the stochastic case can be higher or lower depending on the variance. It is interesting that in some examples we looked at, the subsidy was higher for the stochastic case.

Figure 3.26: Optimal order-up-to level, stochastic and deterministic cases

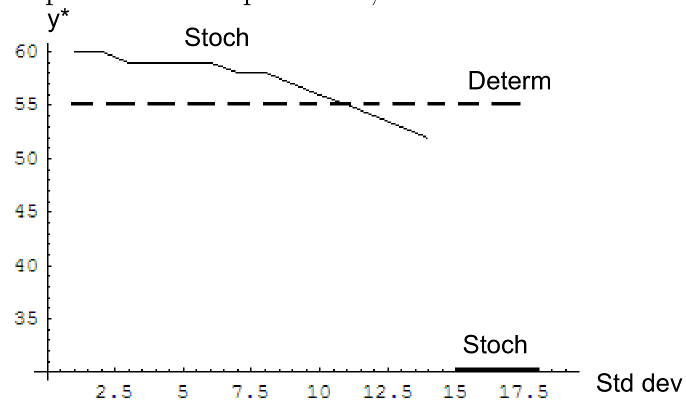

Figure 3.27: Optimal subsidy, stochastic and deterministic cases

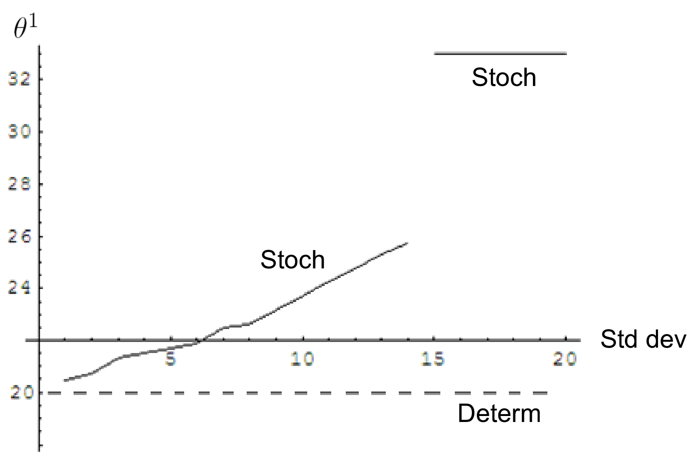

In those examples, we also found that the larger the variance, the smaller the optimal order. This seemingly counterintuitive fact can be explained by looking at the optimal subsidy. If the parameters vary enough to make subsidizing cheaper, it will be optimal to subsidize more and order less. The combination of parameters determines when this is the case. We will perform a sensitivity analysis later that will help us to understand when these conclusions hold.

The solution of the manufacturer's problem is harder to find when the demand is stochastic, and there is no obvious way to derive the solution from the one we found in the case with the deterministic demand. Fortunately, the structure of the stochastic problem allows us to infer some properties of its solution.

Using the recursive expression (3.3.16), we will first calculate the value function for the last period, which will be substituted later as the cost-to-go in the objective function for period 1.

Proposition III.13. *The optimal strategy for period 2, the last period in this context, is the same one as in Proposition III.7:* $z^2 = [D^2 - x^2]^+$, $\theta^2 = c^2$. *The value function is*

$$\begin{aligned} &V(2, x^2, c^2) \\ &= \left[1 - \pi(c^2)\right] p[D^2 + x^2]^+ + h[x^2 - D^2]^+. \end{aligned} \tag{3.3.17}$$

Proof. Note that the problem for the last period is *the same* as the case we dealt with before (deterministic demand). The reason is that $v_3(\cdot) = 0$. In addition, the demand D^2 is already known in period 2. Therefore, the optimal strategy suggested by Proposition III.7 carries over with no modifications. ∎

The expression above for the value function of period 2 does not have a stochastic component in it because, in the second period, both the random shock ε^1 and the demand D^2 are already known. Note that in period 1, this expression is stochastic. Therefore, we have to substitute the expected value (III.13) in the value function for period 1 (3.3.16). We now provide two different expressions for the expected value of (III.13). Note that both contain only costs from default: if the supplier successfully delivers the amount ordered in the second period to cover for the demand of the same period (which is known), there are no holding costs or penalties.

$$
\begin{aligned}
&E_{\varepsilon^1,D^2}V(2,x^2,c^2) \\
&= E_{\varepsilon^1}\left[1-\pi(\theta^1+\varepsilon^1)\right]\int_0^{\infty} p[D^2-x^2]^+ + h[x^2-D^2]^+ d\ \Phi(D^2) \qquad (3.3.18)\\
&= E_{\varepsilon^1}\left[1-\pi(\theta^1+\varepsilon^1)\right]\int_{\max\{0,x^2\}}^{\infty} p(D^2-x^2) d\ \Phi(D^2)\\
&+ \int_0^{\max\{0,x^2\}} h(x^2-D^2) d\ \Phi(D^2). \qquad (3.3.19)
\end{aligned}
$$

Let us now substitute expressions (3.3.18) and (3.3.19) in the value function V for period 1, expression (3.3.16). We get:

$$
\begin{aligned}
V(1,x^1,c^1) = \min_{(y,\theta)\in\mathcal{A}(t)} &\left\{h[y-D^1]^+ + p[D^1-y]^+\right\}\pi(c^1)\\
&+\left\{h[x^1-D^1]^+ + p[D^1-x^1]^+\right\}(1-\pi(c^1)) + \alpha(\theta-c^1)\\
&+\rho\ \{\int_0^{\infty} E_{\varepsilon^1}\left[1-\pi(\theta+\varepsilon^1)\right] p[D^2+D^1-y]^+\\
&+h[y-D^2-D^1]^+ d\ \Phi(D^2)\ \}\pi(c^1)\\
&+\rho\ \{\int_0^{\infty} E_{\varepsilon^1}\left[1-\pi(\theta+\varepsilon^1)\right] p[D^2+D^1-x^1]^+\\
&+h[x^1-D^2-D^1]^+ d\ \Phi(D^2)\ \}(1-\pi(c^1)) \qquad (3.3.20)
\end{aligned}
$$

and hence

$$
\begin{aligned}
V(1,x^1,c^1) = \min_{(y,\theta)\in\mathcal{A}(t)} & \left\{h[y-D^1]^+ + p[D^1-y]^+\right\}\pi(c^1) \\
& + \left\{h[x^1-D^1]^+ + p[D^1-x^1]^+\right\}(1-\pi(c^1)) + \alpha(\theta-c^1) \\
& + \rho\ \{\ E_{\varepsilon^1}\left[1-\pi(\theta+\varepsilon^1)\right] \int\limits_{\max\{0,y-D^1\}}^{\infty} p(D^2+D^1-y)d\ \Phi(D^2) \\
& + \int\limits_{0}^{\max\{0,y-D^1\}} h(y-D^2-D^1)d\ \Phi(D^2)\ \}\ \pi(c^1) \\
& + \rho\ \{\ E_{\varepsilon^1}\left[1-\pi(\theta+\varepsilon^1)\right] \int\limits_{\max\{0,x^1-D^1\}}^{\infty} p(D^2+D^1-x^1)d\ \Phi(D^2) \\
& + \int\limits_{0}^{\max\{0,x^1-D^1\}} h(x^1-D^2-D^1)d\ \Phi(D^2)\ \}\ (1-\pi(c^1)).
\end{aligned} \tag{3.3.21}
$$

Now that we have an explicit expression of the objective function, we will first prove that it is never optimal to order less than we need to cover at least the demand for the first period. Recall that the order-up-to level is defined as $y^t \equiv x^t + z^t$.

Lemma III.14. *Given any fixed level of subsidy θ, consider an order-up-to level y^1 lower than the demand for period 1 (i.e., $y^1 < D^1$). The strategy (D^1, θ) is less costly than the strategy (y^1, θ) (i.e., the strategy (y^1, θ) is suboptimal).*

Proof. Let Π be a policy such that the order and subsidy are $(y^1 - x^1, \theta^t)$ for some $y^1 < D^1$ and let Π' be a policy such that the order and subsidy are $(D^1 - x^1, \theta^t)$. Arithmetical comparison of $v^{\Pi}(1,x,c) < v^{\Pi'}(1,x,c)$ renders the result. ■

From the last proposition, we know that it is suboptimal to order so little that not even the demand for period 1 is covered. If we assume that the demand for period 2 is bounded, the fact that ordering too much is suboptimal it is also very intuitive. At the very extreme, we will order enough for both periods, but no more than that.

Lemma III.15. *Assume that the demand D^2 is bounded, i.e., that $0 \leq D^2 \leq \overline{D^2}$. Given any fixed level of subsidy θ, it is suboptimal to order $y^1 > D^1 + \overline{D^2}$.*

Proof. The proof is done in the same way as Lemma III.14. ∎

So far, we have proved that it is not optimal to order less than D^1 or to order more than $D^1 + \overline{D^2}$. Therefore, we know that optimality is achieved between D^1 and $D^1 + \overline{D^2}$.

Let us assume now that $\pi(c) = 1$ for some $c < \infty$.

From the extreme value theorem in calculus, there exists a pair (y, θ) that minimizes expression (3.3.21). We cannot obtain a closed-form solution, since expression (3.3.21) has cross terms, and therefore its critical points are saddle points. However, we can draw some results about the optimal costs when the parameters change.

3.3.3 Sensitivity analysis

Looking at equation (3.3.20), we observe that the stochastic demand relaxation preserves a similar structure of the objective function v for the deterministic demand case. However, we no longer have a closed-form solution for the optimal actions.

While it is not possible to obtain the same type of results about the order size and subsidy amount, we can still obtain some results regarding the direction of the changes. For this purpose, we recall the envelope theorem, commonly used in Microeconomics. We present the theorem and refer the reader to *Turkington* (2007) for more details.

Proposition III.16. ***The envelope theorem.***— *Consider an arbitrary maximization (or minimization) problem where the objective function $f(\overline{x}, \overline{r})$ depends on some parameters $\overline{r}$:*

$$f^*(\overline{r}) = \max_{\overline{x}} f(\overline{x}, \overline{r})$$

The function $f^*(\overline{r})$ is the problem's optimal-value function: it gives the maximized (or minimized) value of the objective function $f(\overline{x},\overline{r})$ as a function of its parameters r.

Let $\overline{x}^*(\overline{r})$ be the (arg max) value of $\overline{x}$, expressed in terms of the parameters, that solves the optimization problem, so that $f^*(\overline{r}) = f(\overline{x}^*(\overline{r}),\overline{r})$. The optimal-value function $f^*(\overline{r})$ derivatives with respect to parameters (which described how $f^*(\overline{r})$ will change when the parameters $\overline{r}$ change) can be obtained as follows:

$$\frac{\partial f^*(\overline{r})}{\partial r_i} = \left.\frac{\partial f(\overline{x},\overline{r})}{\partial r_i}\right|_{\overline{x}=\overline{x}^*(\overline{r})}$$

That is, the derivative of $f^*(\overline{r})$ with respect to r_i is given by the partial derivative of $f(\overline{x},\overline{r})$ with respect to r_i, holding $\overline{x}$ fixed, and then evaluating at the optimal choice.

From this, we can obtain some results about the optimal cost. These results concern the changes of the optimal cost when the parameters change. First, we present the derivatives.

Lemma III.17. *The following are the derivatives of the optimal value of the objective function in 3.3.21 with respect to different parameters:*

$$\begin{aligned}\frac{dv}{dc^1} &= \left\{h[y^{1*} - D^1]^+ + p[D^1 - y^{1*}]^+\right\}\pi'(c^1)\\ &+ \left\{h[x^1 - D^1]^+ + p[D^1 - x^1]^+\right\}(-\pi'(c^1)) - \alpha\\ &+ \rho\textstyle\int_0^\infty E_{\varepsilon^1}[-\pi'(\theta^{1*} + \varepsilon^1)]p[D^1 + D^2 - y^{1*}]^+ d\Phi(D^2) \quad \pi(c^1)\\ &+ \rho\textstyle\int_0^\infty E_{\varepsilon^1}[1 - \pi(\theta^{1*} + \varepsilon^1)]p[D^1 + D^2 - y^{1*}]^+ + h[y^{1*} - D^1 - D^2]^+ d\Phi(D^2) \quad \pi'(c^1)\\ &+ \rho\textstyle\int_0^\infty E_{\varepsilon^1}[-\pi'(\theta^{1*} + \varepsilon^1)]p[D^1 + D^2 - x^1]^+ d\Phi(D^2) \quad (1 - \pi(c^1))\\ &+ \rho\textstyle\int_0^\infty E_{\varepsilon^1}[1 - \pi(\theta^{1*} + \varepsilon^1)]p[D^1 + D^2 - x^1]^+ + h[x^1 - D^1 - D^2]^+ d\Phi(D^2)(-\pi'(c^1))\end{aligned} \tag{3.3.22}$$

$$\frac{dv}{d\alpha} = \theta^{1*} - c^1$$

$$\frac{dv}{dh} = [y^{1*} - D^1]^+ \pi(c^1) + [y^{1*} - D^1]^+ (1 - \pi(c^1))$$
$$+ \rho \int_0^\infty [y^{1*} - D^1 - D^2]^+ d\Phi(D^2) \quad \pi(c^1) \tag{3.3.23}$$
$$+ \rho \int_0^\infty [x^1 - D^1 - D^2]^+ d\Phi(D^2)(1 - \pi(c^1))$$

$$\frac{dv}{dp} = [y^{1*} - D^1]^+ \pi(c^1) + [x^1 - D^1]^+ (1 - \pi(c^1))$$
$$+ \rho \int_0^\infty E_{\varepsilon^1}[-\pi'(\theta^{1*} + \varepsilon^1)] \quad [D^1 + D^2 - y^{1*}]^+ d\Phi(D^2) \quad \pi(c^1) \tag{3.3.24}$$
$$+ \rho \int_0^\infty E_{\varepsilon^1}[-\pi'(\theta^{1*} + \varepsilon^1)] \quad [D^1 + D^2 - x^1]^+ d\Phi(D^2) \quad (1 - \pi(c^1))$$

$$\frac{dv}{dD^1} = \left\{-hI_{[y^{1*} - D^1 > 0]} + p_{[D^1 - y^{1*} > 0]}\right\} \pi(c^1)$$
$$+ \left\{-hI_{[x^1 - D^1 > 0]} + p_{[D^1 - x^1 > 0]}\right\} (1 - \pi(c^1))$$
$$+ \rho \int_0^\infty E_{\varepsilon^1}[-\pi'(\theta^{1*} + \varepsilon^1)] p I_{[D^1 + D^2 - y^{1*} > 0]^+} - hI_{[y^{1*} - D^1 - D^2 > 0]} d\Phi(D^2) \quad \pi(c^1)$$
$$+ \rho \int_0^\infty E_{\varepsilon^1}[-\pi'(\theta^{1*} + \varepsilon^1)] p I_{[D^1 + D^2 - x^1 > 0]^+} - hI_{[x^1 - D^1 - D^2 > 0]} d\Phi(D^2) \quad (1 - \pi(c^1)) \tag{3.3.25}$$

$$\frac{dv}{dx^1} = \left\{hI_{[y^{1*} - D^1 > 0]} - p_{[D^1 - y^{1*} > 0]}\right\} \pi(c^1)$$
$$+ \left\{hI_{[x^1 - D^1 > 0]} - p_{[D^1 - x^1 > 0]}\right\} (1 - \pi(c^1))$$
$$+ \rho \int_0^\infty E_{\varepsilon^1}[-\pi'(\theta^{1*} + \varepsilon^1)](-p) I_{[D^1 + D^2 - y^{1*} > 0]^+} + hI_{[y^{1*} - D^1 - D^2 > 0]} d\Phi(D^2) \quad \pi(c^1)$$
$$+ \rho \int_0^\infty E_{\varepsilon^1}[-\pi'(\theta^{1*} + \varepsilon^1)](-p) I_{[D^1 + D^2 - x^1 > 0]^+} + hI_{[x^1 - D^1 - D^2 > 0]} d\Phi(D^2) \quad (1 - \pi(c^1)) \tag{3.3.26}$$

Proof. The results follow from Leibniz's rule and from the envelope theorem. ■

From this result, even if we do not know how the actual order and subsidy change, we do know the magnitude of the changes in the optimal cost, and for most cases, even the direction of the changes. From the derivatives we computed above, we can see that whenever α, h or p increase, the optimal cost will increase too. On the other hand, if c^1, D^1 and x^1 change, it is not clear whether the optimal cost will increase or decrease. The reason for this is that, depending on the parameters, when the initial assets are higher, this may result in higher *expected* holding costs, since it is more likely that the manufacturer will have to pay for inventory. Similarly, an increase in demand (or decrease in initial inventory) may result in lower *expected* holding costs.

We have described some of the implications of changing a parameter. We saw that there's no definite implication of an increase in the demand for period 1.

Figure 3.28: Sensitivity of the Subsidy θ^1 vs. Initial Inventory x^1 for a Bernoulli-distributed D^2

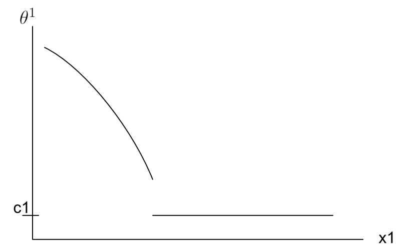

Figure 3.29: Sensitivity of the Subsidy θ^1 vs. initial inventory x^1 for a discrete uniformly distributed D^2

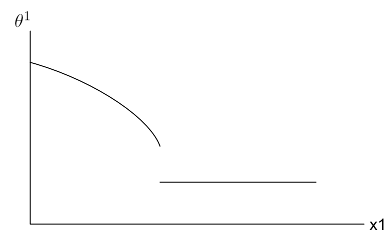

In order to check if similar conclusions follow from the sensitivity analysis for the decision variables when the demand in the second period is stochastic, we worked out some examples. One of them has a Bernoulli-distributed demand, and the other has a uniformly distributed demand. The Bernoulli distribution can be thought of as an extreme case of a bimodal distribution. In the examples, we found that the conclusions from the sensitivity analysis we did in previous sections for the case of deterministic demand in period 2 follow smoothly when the demand is stochastic and follows a *uniform distribution*, but the analysis does not apply in a straightforward manner to the Bernoulli-distributed demand case.

Since the conclusions for the sensitivity analysis of the case with Bernoulli-distributed demand are different, we will take a moment to give several details in the analysis for changes in the initial inventory for this case. We will also give the main conclusions for changes in the demand in period 1.

- Initial inventory

Figure 3.30: Sensitivity of Order-up-to Quantity y vs. Initial Inventory x^1 for a Bernoulli-distributed D^2

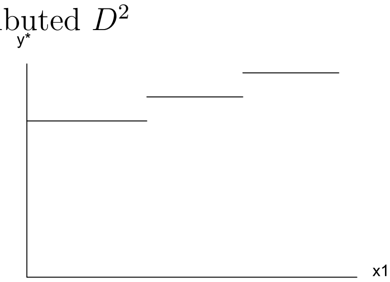

Figure 3.31: Sensitivity of the Order-up-to Quantity y vs. Initial Inventory x^1 for a discrete uniformly distributed D^2

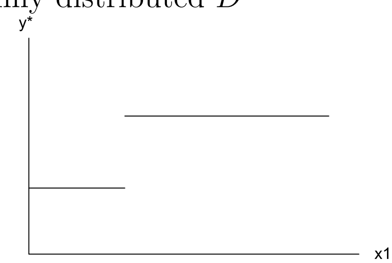

At the beginning, increasing the initial inventory x^1 causes the order size to decrease in the amount that the inventory increased, while the order-up-to level stays constant. For example, if it was optimal to order only for period 1 ($z^1 = D^1 - x^1$), the order size would decrease.

Eventually, the order-up-to amount jumps up to cover some of the demand of

D^2. The order size jumps up with the order-up-to level. After this jump, we see again that the order size decreases as x^1 increases. When the demand has a Bernoulli distribution, the order amount jumps up twice: the first time to increase the order to cover the low amount of demand D^L in period 2 (i.e., $z^1 = D^1 + D^L - x^1$), the second time to cover the high amount of demand D^H in period 2 (i.e., $z^1 = D^1 + D^H - x^1$).

Figure 3.32: Sensitivity of the Order z vs. Initial Inventory x^1 for a Bernoulli-distributed D^2

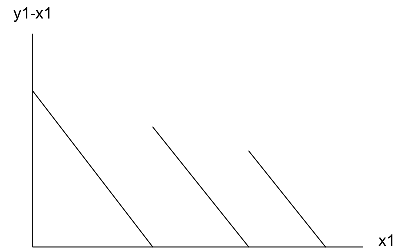

Figure 3.33: Sensitivity of the Order z vs. Initial Inventory x^1 for a discrete uniformly distributed D^2

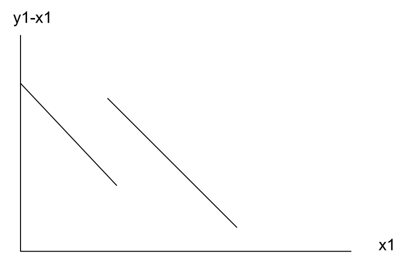

The more initial inventory we have, the less we need to subsidize for the future, given the fact that a higher x^1 will give incentives to order today for the two periods. The subsidy jumps down every time that the order-up-to jumps up, unless the condition $\theta^1 = c^1$ is holding. As we mentioned, when the stochastic demand has a uniform distribution, the changes are similar to what we found in the deterministic case. As a reference, we include both graphs in the diagrams. See Figures 3.28, 3.29, 3.30, 3.31, 3.32 and 3.33.

- Demand in period 1

Figure 3.34: Subsidy vs D^1 for a Bernoulli distributed D^2

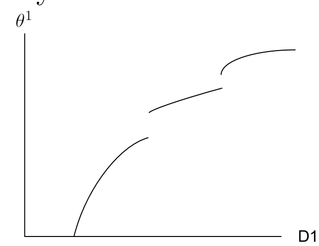

Figure 3.35: Order-up-to vs D^1 for a Bernoulli distributed D^2

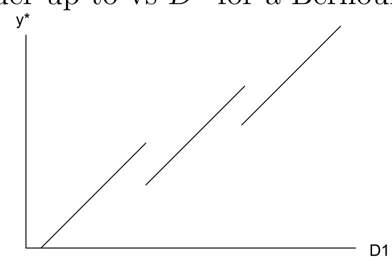

As we mentioned, when the demand distribution is *uniformly-distributed*, the conclusions are similar to the deterministic case. For brevity, we refer the reader to the figures in the deterministic demand section.

We will focus in the case when the demand is *Bernoulli-distributed*, given that in that scenario the situation is a little different. In this case, it is possible to have more than one non-monotonic jump. In the example we present, at the beginning the order-up-to amount covers for the higher level of demand in period 2. When D^1 increases more, it is better to order only the minimum amount for period 2. Finally, if D^1 is very high, it is optimal to order only for the first period.

As the demand in the current period increases, and as the order-up-to amount decreases, it is necessary to provide the supplier with a greater amount of subsidy to guarantee delivery in the future, and therefore the subsidy jumps up when the order-up-to jumps down. The order level inherits the same structure as the order-up-to level; and therefore we will just show the diagram for the

Figure 3.36: Subsidy vs ED^2 for a Bernoulli distributed D^2

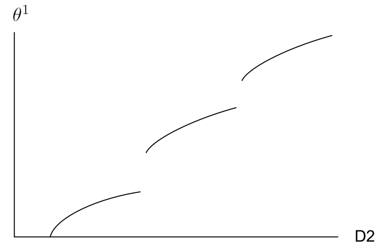

Figure 3.37: Order-up-to vs ED^2 for a Bernoulli distributed D^2

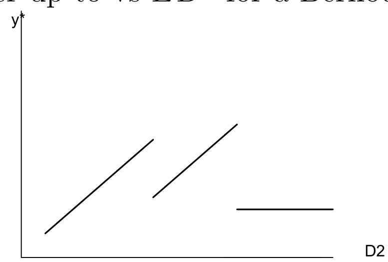

order-up-to.

See figures 3.34 and 3.35.

- Expected value of demand in period 2.

Conclusions very similar to the ones we described for demand in period 1 follow for the demand in the second period. The order-up-to level will jump down if $E\ D^2$ grows considerably, as it will be better to not order for both periods, but only for the first one. When this happens, the subsidy will jump up. The order-up-to level will jump down more than once (and, the subsidy level will jump up more than once unless the constraint $\theta^1 = c^1$ has been reached). Note that when once we reach the point when it is optimal to order only for period 1, the order-up-to level does not change any more, so the graph becomes flat to the right. The order level inherits the same structure as the order-up-to level; and therefore we will just show the diagram for the order-up-to.

See figures 3.36 and 3.37.

3.4 Conclusions

We first studied the one-period model, where there was a finite amount that would make the probability of successful delivery equal to 1 and where the subsidy would have an immediate effect. In this case, we first determined conditions that defined whether diversification was or was not an alternative to be considered when compared to subsidizing a supplier.

In this context, the amount $\widetilde{\omega}_i \equiv \frac{F_i(p(1-\lambda_i)-\lambda_i h)}{(1-\lambda_i)(p+h)}$ defines a threshold that determines if supplier i will be reliable enough for the manufacturer. Below this level, this supplier is not reliable enough and the expected penalty costs for the default are higher than the expected holding costs, and therefore it becomes necessary to diversify. If $p(1-\lambda_i) < h(\lambda_i)$, then every level of subsidy satisfies $w > \widetilde{w}$. This means that supplier i is strong enough to have a reasonable probability of successful delivery, so having this supplier exclusively has a lower cost than diversifying.

Therefore, the conditions $p(1-\lambda_1) < h(\lambda_1)$ and $p(1-\lambda_2) < h(\lambda_2)$ are needed to determine whether diversification may or may not be optimal.

In the case when $p(1-\lambda_1) \geq h(\lambda_1)$ and $p(1-\lambda_2) \geq h(\lambda_2)$, by choosing if having an exclusive supplier or diversifying, the manufacturer is choosing between certainty and uncertainty: the cost of choosing supplier i as an exclusive supplier is the cost of providing full subsidy. The cost of diversification is the cost of excess inventory if no one defaults, plus the penalty cost if both suppliers default. As a consequence, if we assume that each of the available suppliers can satisfy our total demand of supplies, and if the manufacturer prefers to avoid uncertainty, she has a greater incentive to grant subsidies to an exclusive supplier.

In the case when $p(1-\lambda_1) < h(\lambda_1)$ or $p(1-\lambda_2) < h(\lambda_2)$, the manufacturer is choosing between certainty and uncertainty. Certainty in this case is obtained if the manufacturer provides full subsidy. Uncertainty, however, is obtained if the manufacturer simply orders from an *exclusive* supplier without ordering any backup inventory or providing any subsidy.

In this one-period setting, under some circumstances, such as when diversification occurs, we find that under this model, the backup inventory amount is high: the manufacturer orders *twice* the current period demand, even if there is no future demand.

We later considered the case with two periods, while all the other assumptions remained unchanged. We obtained the solution and analyzed it. We found that depending on the initial parameters, the manufacturer's optimal actions would reflect the flexibility that the introduction of time as a dimension would add to his choice of possible actions. In the example we examined, the penalty cost is initially extremely low. In this situation, the manufacturer wants to have some diversification of the risk concerning the supplies she needs for the next period: she can order them now, pay the holding cost, and if the supplier defaults, she can have one more opportunity to order them again. If the penalty cost increases a lot, the manufacturer will want to actually fund the supplier instead of risking to have to pay the penalty, and she will want to order only for the current period, given the fact that she will want to save the holding cost for ordering for the second period too. Finally, if the penalty cost is very high (region 4), the manufacturer will opt for the most protection she can get against the penalty: full funding in period 1 and ordering for both periods now and saving the future funding costs. In summary, the time horizon gives the manufacturer some flexibility. Depending on the required subsidy amount, the holding costs, and the shortage penalties, the supplier has the option to subsidize, to have backup inventory or to wait and see if subsidy will be required to satisfy the accumulated demand in period 2.

We then switched to study the case when we had a two-period setting, but the subsidy would not be achieved in the same period in which it was granted. In addition, we relaxed the assumption that a finite amount F existed that would guarantee that no default would happen. We assumed that the demand in the second period was known and deterministic. In this setting, we found a threshold $\widehat{\theta^1} = g^{-1}(\frac{h}{\rho p})$. If the subsidy is lower than this amount, it is optimal to order $[D^1 - x^1]^+$. Intuitively, once the manufacturer has given enough subsidy to the supplier, there is enough confidence that the supply will be delivered next period, and therefore it is not necessary to take the extra precaution of ordering for the future today. In other words, $\widehat{\theta^1}$ actually is the threshold that defines a *strong* supplier, in the sense that we can trust that the supplier will deliver the supplies, and therefore there is no need to add backup inventory for the future.

This partition allowed us to solve the problem optimally, obtaining closed-form solutions. With these closed-form solutions, we performed a sensitivity analysis.

We found that although changes in the optimal subsidy may be monotonic, this is not necessarily the case for the order size. For example, for the case when the demand for either period increases, the subsidy increases and the order size may increase for a while, until the supplier becomes so reliable that it is no longer necessary to order for both periods, and the order size jumps down. A second example is when the current inventory amount increases. The subsidy amount decreases, but this forces the order-up-to amount to go up, by ordering only from the current period to ordering for both.

Relaxing the assumption of deterministic demand in the second period does not prevent us from giving a set of equations that characterize the solution. However, a closed-form solution can no longer be found so that a similar sensitivity analysis be performed as in the deterministic demand case. There is no *rule of thumb* to determine if the subsidy or the order will be greater or lower than their counterpart in the deterministic case. In the stochastic case, the conclusions of the sensitivity analysis for the deterministic demand may or may not hold. However, despite this limitation, the envelope theorem enables us to determine in some cases the direction of the changes in the costs (the objective function), which is good news given the limitations we face in this case. Using the envelope theorem, and computing the corresponding derivatives, we obtain that whenever α, h or p increase, the optimal cost will increase too.

CHAPTER IV

Room Assignment Optimization

4.1 Introduction

The tourism industry is an important part of the economy in many countries throughout the world. *Vogel* (2001) points out that three hundred million people work in the tourism industry worldwide, and approximately 1.5 trillion US dollars of direct and indirect revenue, out of a total world economic output of around 40 trillion US dollars, is generated by this industry. Vogel also mentions that "in the United States, tourism is estimated to account for approximately 5% of gross domestic product and to be the third largest retail industry after automobile dealers and food stores." The statistics from *Lum and Moyer* (2001) at the Bureau of Economic Activity show that almost 1% of gross domestic product from the United States is actually related to the accommodation segment, which is one fifth of the tourism industry revenue.

Hotel operations are complex and encompass a great variety of tasks. We refer the reader to *Malhotra* (1997) for a comprehensive description of all of them. We now will discuss one of these tasks, which is the process of assigning a room to each requested booking.

The room assignment task consists of finding an available room that matches the fare type and attribute requests to as many existing bookings as possible. Attribute requests can be varied; for example, the type of bedding, some specific view from the room, wheelchair accessibility, floor number, requests for connecting rooms for pairs of bookings from customers traveling together, and so on. However, it is not always possible to find a room for every booking, so we would like to assign rooms to as many bookings as we can. For this purpose, we would like to assign rooms in a way that creates as few "holes" as possible (a hole is a small block of empty room-nights between two reservations).

In summary, the room assignment problem can be stated as follows:

Maximize the number of room-nights assigned to existing bookings, with the following constraints:

- the assigned room is of the right type,
- the assigned room is available for the entire length-of-stay and
- the assigned room has the requested attributes, which can include connecting rooms.

With respect to the room assignment task, *ConRunner, a Why and How-To reference for Convention organizers*[1] states that "Assigning particular rooms is a nice touch, and it's often the only way to guarantee things like requests for connecting rooms. But it's a lot of work. Some hotel reservation systems can assign room numbers to reservations as they come in from a list of rooms in each block." It is true that some hotel management tools do assign room numbers to bookings *as these bookings are being made*. Many of these tools are designed primarily for accounting and automation of some hotel operations (such as daily reporting of checkouts). Some of these applications assign rooms automatically when the booking is made. An example of this is the hotel application from *IT Edge Softwares*[2].

Other tools allow the front desk to assign rooms manually to bookings at any time they decide (when the booking comes in, when the guest arrives, or at any time in

[1] *ConRunner.net* (2009)
[2] *VCL Components* (2009)

between). Examples of this are Hotel and Property Management software from *Business Software Solutions* (*Business Software Solutions, Inc.* (2009)), Ezee FrontDesk Hotel software from *Technosys*[3] and the software from *Resort Data Processing Inc.*[4].

Some hotels obtain some or all of their bookings online via their own websites (rather than from Internet travel agencies such as Expedia, Travelocity, or Orbitz). Still others use software that automatically assigns rooms as bookings are made. Examples of such software tools are the Hotel Booking software from *Best Software Inc.*[5] and the Hotel Reservation System by *Soft Acid*[6].

Although it is true that some hotels assign a room number as soon as the booking is made, many hotels assign room numbers at a *later* stage. *Bitran and Gilbert* (1996) give an overview of the planning procedure of a hotel, which they divide into several phases. The last phase is *daily inventory planning*, which, as they point out, "involves the daily allocation of rooms to customers as they arrive at the hotel." A reason that many hotels assign rooms late in the booking cycle is that many of them practice *overbooking*. Overbooking is one of the simplest revenue management techniques and the first one that was thoroughly researched and implemented. We refer the reader to the work from *McGill and van Ryzin* (1999) for more details. It consists of taking reservations beyond the capacity for accommodation with the purpose of "increasing capacity utilization in a reservation-based system where there are significant cancellations" ((*Talluri and van Ryzin*, 2004, p.129)). Early room assignment does not work well with this technique. Therefore, many hotels prefer to assign rooms on the day of arrival or perhaps a few days in advance: for example, they may assign room numbers in advance for arrivals happening *today* and in the next n days (where n is, for example, 7 days). They would not assign rooms for all the bookings they currently have for the coming months.

[3] *eZee Technosys Pvt. Ltd.* (2009)
[4] *Resort Data Processing, Inc.* (2009)
[5] *Best Software, Inc.* (2009)
[6] *SoftAcid* (2009)

To illustrate this and why many hotels prefer to delay the room number assignment, we present an example. Assume that we are assigning rooms for the period of June 24 to June 30. Consider a fictitious facility where there are only two rooms available, both of the same type. Assume that three bookings are made on May 15, with check-in and check-out dates as follows:

Booking #: (Check-In /Check-Out)

Booking A: June 24th/25th.

Booking B: June 25th/30th.

Booking C: June 25th/26th.

The room number assignment is done as the bookings are received. Therefore, on May 15, bookings A and B are assigned to room 1, and booking C is assigned to room 2. Assume that on June 1 there is a new booking request. The booking request is as follows:

Booking D: June 24th/26th.

Assume that the rate of cancellation is high enough for the management to practice overbooking without quality of service concerns. There is no room available for booking D, but the booking is still accepted.

Assume that on June 5, booking B is canceled.

When the actual stay dates arrive, if there were no changes in the room number assignment performed on May 15, there is no place for booking D according to the schedule created on May 15, so the customer is sent to another hotel (the customer is "walked") and offered a free night at a later time to compensate for the inconvenience.

Let us now change the assumption regarding the timing of the room number assignment. *Assume that the room number assignment is done each day for the next* *seven days*. After booking 2 is canceled on June 5, only bookings 1, 2 and 4 need a room. On June 18, the room number assignment is performed for the next seven days, so bookings 1, 2 and 4 are given a room number if possible. On June 18, because of the cancellation, bookings 1 and 3 can be assigned to room 1, and booking 4 can be assigned to room 2. In other words, all the bookings were successfully assigned a room. As we can see, when overbooking occurs, it can be better to delay the room number assignment.

In the example we described above with two rooms, four bookings and one cancellation, the optimal room assignment is obvious and easy to find, since it is possible *to inspect the entire inventory* to find a room with the requested length-of-stay *by trying all the possible combinations*. However, this is not feasible when there are thousands of bookings to assign to hundreds of rooms. It is even harder when customers have attribute requests as described (bedding, view, floor, wheelchair-accessible, connecting rooms, etc.), which must be considered in addition to the length-of-stay requirement. The connecting rooms attribute presents an additional complication to the problem when compared to other attribute requests. The reason is that a connecting rooms request involves two separate bookings and two rooms, while other attributes involve only one room.

To our knowledge, there is no literature related to tailoring scheduling heuristics to the hotel industry, as we will see in the literature review in the next section. A possible reason for the dearth of literature is that the task is usually performed by front desk staff, who usually review the list of rooms available for each type and assign each guest the first one they find that meets their requests. The front desk staff may also have some program to assist with the task, such as the commercial software we described above. Most of the programs described assign a room when a booking is made, very likely using a version of the procedure we just described (assign the booking to the first room in a list of rooms that matches availability and the requested length-of-stay, repeat the operation for the second booking and so on, until we reach the end of the booking list). We will call this procedure the lexicographic algorithm. As we showed in the example with four bookings, with this type of approach there is a possibility of unintentionally blocking a booking.

Another possible reason for the lack of research on room assignment is that for many places, hotel demand is highly seasonal. Spring, fall and winter are usually low seasons, with only some holidays having high demand. At these times and places, it is not difficult for many hotels to provide lodging to every visitor. However, in the summer, some popular destinations have many visitors, and their hotels are close to 100% occupancy. For those months, having a good tool that allows efficient room assignment is a good way to promote customer satisfaction. In addition, for destinations like New York City, the average occupation for an entire year is usually 80% on average, as we can see in *NYC & Company, Inc.* (2009). Efficient room assignment is more important for hotels in such locales.

Besides the overbooking technique we already discussed, there are other revenue management techniques that are commonly applied in the hotel industry. Generally speaking, these strategies allow the consideration of future demand, including walk-ins, in order to decide how much to charge for a room, and whether to accept an incoming booking request. This is generally done with an allocated amount of inventory to each type of customer, or with bid prices. These techniques have proven to be valuable to increase profits. ((*Cross*, 1998, p. 3),); for example, reports that revenue management helps Marriott Hotel to gain US$100 million in additional annual revenue. For more details on these techniques, we refer the reader to the following revenue management surveys: *Weatherford and Bodily* (1992) and *Bitran and Caldentey* (2003), and the book from *Talluri and van Ryzin* (2004).

We would like to point out that for the revenue management techniques, the modeling usually focuses on the expected number or percentage of bookings, on cancellations and no-shows, and on expected values. In general, assigning room numbers to bookings is *not* an explicit part of this type of model, as we can see from the surveys listed above. The focus of these models is generally the *amount* of inventory to allocate to each fare type on a specific night rather than on the assignment of a specific room to each booking. Unfortunately, given the stochastic nature of arrivals, cancellations and no-shows, it occasionally happens that more customers of a specific type do show up (even when the *total* amount of bookings accepted does not exceed total hotel capacity). It is in the best interest of the hotel to assign room numbers so that the number of bookings without an assigned room is minimized, if not eliminated. In this way, hotels can provide better customer service. Some hotels give customers in this situation an "upgrade," which blocks a room that could otherwise be sold. Some others may have to let the customer walk away and lose the customer's trust and goodwill. By efficiently assigning rooms, hotels can reduce the amount of times they have to pay a penalty when they cannot match the specified room in a booking request, be it in the form of an upgrade or in the form of goodwill loss.

Even if a hotel does not use revenue management techniques, an efficient allocation of rooms allows more walk-ins to be accommodated. We will focus on the task of maximizing the number of assigned room-nights to *existing* bookings that will arrive in the next n days, where n is the number of days in advance when most cancellations will have happened. We do not incorporate the modeling of walk-ins. Such modeling is not needed to assign as many rooms as needed to existing bookings. We assume that the decision to accept an incoming booking has already been made as part of the revenue management policies of the hotel.

While the upgrading process (giving a better type of room to a customer) could potentially be incorporated as part of the optimization problem, there are two reasons why we choose to not to do so. We consider it to be appropriate for the front desk staff to determine which customers to give an upgrade to rather than an automated algorithm because the front desk staff can perceive easily when there is dissatisfaction, and, thus, the need to compensate in some way the customer for any inconvenience. Secondly, not every hotel guarantees attributes when their customer requires a booking, and of those who do, not each of them practices upgrading when the attribute request cannot be honored. For this reason, our work will be focused on assigning rooms, trying to satisfy the room type and attribute requests as possible, and will not include upgrading. In practice, this should not be a major issue for the front desk staff: if we assign the rooms as efficiently as possible to the bookings, there will only be a few bookings that could not be assigned, and, therefore, there will not be many bookings that will need an upgrade.

Our purpose in this paper is to develop a heuristic algorithm that is specific to the lodging industry by taking advantage of the characteristics that are specific to hotels. This can assist the front desk staff by automating the process in an efficient way, i.e., increasing the number of bookings that are successfully assigned. By doing so, customer satisfaction is increased, more walk-ins can be accommodated, and therefore more income can be generated, decreasing the amount of penalties and loss of goodwill.

This paper is organized as follows: In section 2, we present a literature review. In section 3, we present the mathematical definition of the problem. In section 4, we present the heuristic algorithm that we propose to solve this problem. In section 5, we present the results of comparing this algorithm to a lexicographic method. Finally, in section 6, we present our conclusions and suggest future research directions.

4.2 Literature Review

The assignment of rooms to bookings should take into account the room type, desired attributes, connecting room requests, and current available time for each room. Attributes of rooms may include view, floor, bedding, etc., while available times depend on when customers currently in house check out of the hotel or on scheduled maintenance times. Connection requests are more complicated, since they involve matching the attributes and type for two rooms instead of one.

The problem of assigning rooms to a set of reservations is an example of a constrained satisfaction problem. This can be understood as the problem of finding an assignment of values to a set of variables such that a set of constraints is satisfied. These problems are called *the constraint satisfaction problem (CSP).* For introductory material on CSP, see Chapter 5 in the work of *Russell and Norvig* (1995). For a survey of algorithms for CSP, we refer the reader to *Kumar* (1992).

It is important to notice that in this context, it is assumed that the problem at hand is to find a feasible assignment such that *every* reservation is honored. When a feasible assignment for every reservation is not possible, the problem becomes one of assigning a room to as many bookings as possible. In the context of general scheduling, *Schiex et al.* (1995) suggest including some metrics of the seriousness of a violation of a constraint in the problem formulation. They call this approach a Valued Constraint Satisfaction problem.

This approach is generic and not tailored to the hotel industry. There are many settings where examples of constraint satisfaction problems are present. Some examples are the problems known as the university lecture scheduling or the job shop in manufacturing. They present some similarities to the room assignment problem.

A quick review of the work done for these problems shows that the solution is usually found using heuristics. For the university lecture scheduling problem, see, for instance *Abramson* (1991) and *Elmohamed et al.* (1998), for annealing algorithms to solve this problem. See *Luan and Yao* (1996), *Erben and Keppler* (1996), and *Ueda et al.* (2001) for genetic algorithms in this context. For a survey of techniques, see the work from *Burke et al.* (1997) and *Burke and Petrovic* (2002).

In the case of the job shop scheduling problem, see for instance the work of

Coffman and Bruno (1976), *Vaessens et al.* (1996), *Jain and Meeran* (1999) (1999), *Blazewicz et al.* (1996) and *Sadeh and Fox* (1996), *Sanlaville and Schmidt* (1998). *Davis* (1985) and *Bierwirth and Mattfeld* (1999) discuss the application of genetic algorithms to this problem. *van Laarhoven et al.* (1992) discuss the theory and applications of the simulated annealing techniques to the same problem. *Dell'Amico and Trubian* (1993) discuss how to apply tabu search in this context.

Many job shop problems are of the type known as *variable interval scheduling* or *variable job scheduling.* This means that each job has a flexible start and end time, or in other words, a flexible window of execution. A lot of the job shop literature is on this type of model, which is NP-complete (*Garey et al.* (1976)). Other job shop models, which are more relevant to the type of problem we are studying, belong to the category known as *fixed interval scheduling, fixed job scheduling*, or *k-track assignment* or simply *interval scheduling,* which has the characteristic that each job must be performed starting at a specific time and ending at a specific time. The room assignment problem is a fixed interval scheduling problem, given that every booking must be assigned to a room for the entire length-of-stay, starting exactly on the requested arrival date.

A survey for fixed intervals is provided by *Kolen et al.* (2007). As we will see, many interval scheduling problems are computationally expensive, and the room assignment problem is no exception. In order to understand the complexity in reference to the work that has been done in this area, let us temporarily drop the connection requests from the room assignment problem.

Kolen et al. provide a classification for interval scheduling problems. One of the categories of interest to us is the Interval Scheduling with Machine Availabilities (ISMA) problem. In this problem, several tasks have to be assigned to processors for processing. The processors are identical, except for the availability times. Kolen et al. point out that ISMA was proved to be NP-complete by *Papadimitriou* (1982).

Notice that a simplified version of the room assignment problem (where there are no room types) is an ISMA-type problem. In this version, rooms can be thought of as processors, while the bookings are the jobs to be processed.

Kolen et al. also define the hierarchical interval scheduling (HIST) problem with T types. In this setting, all processors are available for the same period of time, but they differ because each job can be processed by some processors: Jobs and processors have a hierarchical type (i.e., types are ordered). Jobs of type t can only be processed by a processor of type t or lower. Kolen et al. prove that a HIST problem with T types (where the number of types is $T = 2$ or greater) is NP-complete. A room assignment problem with T room types without unavailable blocks in the middle of the x day period, but with the possibility of upgrading a customer, can be considered a HIST problem with T types.

It is important to clarify that in practice, upgrades do not exceed one category (it is not in the best interest of the business to give the highest-category room to customers who paid the lowest possible fare). Also, as we pointed out, it is appropriate to let front desk staff or managers make the decision as to who gets upgraded, given that they perceive better the need to do so when customer satisfaction is at stake. Therefore, the HIST problem is not the best way to model the room assignment problem in practice. Our purpose of reviewing it here is to make the reader aware of the complexity of the room assignment problem.

We have seen that the room assignment problem falls into two of the categories defined by Kolen et al. for the interval scheduling problems, and these categories are problems that are NP-complete.

We found an algorithm in the literature (*Brucker and Nordmann* (1194)) that attempts to solve problems in which each processor's track (available window) can process only a given set of jobs. They call this the generalized k-track problem. This is similar to the room assignment problem, where only a given set of bookings can be scheduled for each room, depending on the room availability and attributes; however, this algorithm is computationally expensive. Brucker and Nordmann do point out that, in practice, this algorithm is ***not*** *useful for* $k > 5$. In the room assignment context, this algorithm is $O(n^k)$, where n is the number of bookings and k is the *number of rooms.* [7] For this reason, their algorithm is not even useful for very small hotels, except for the rare ones that do have $k = 5$ rooms at most.

Recall that we dropped the assumption that the bookings contain connecting room requests. As we have seen so far, the nature of this simpler version of the room assignment problem is complex. Connection requests can be considered within the scope of adjacent resource scheduling (ARS), which, as *Duin and van Sluis* (2006) point out, has not been widely studied. The room assignment problem with connection requests can be thought of as an ARS problem: In a connection request, two jobs (bookings) should be processed by adjacent processors (connected rooms).

ARS is closely related to interval scheduling and to multiprocessor processing (when a single job requires at least two processors). The difference between multiprocessor scheduling and ARS is that in multiprocessor scheduling, it is necessary that a pre-specified number of processors performs a job, while in ARS, it is necessary that such processors are selected from a collection of adjacent processors rather than from the whole set.

Duin and van Sluis (2006) study the ARS problem when all processors are available for the same period of time (which is not true in the room assignment problem; each room is available for a different period of time). They study different versions of the problem, one of which they call rectangular ARS. In this version, the need for resources remains constant during the entire interval in which the job is processed

[7] This is true assuming that each room is available for one continuous period of time rather than having some blocked time (for maintenance, for example) *in the middle* of the n days creating two separate blocks. If the current available time for a room is not continuous and is two separate blocks instead, k is the number of *available blocks*. If all rooms are available for a continuous period of time, the number of open slots matches the number of rooms.

(this is the case for the room assignment problem: the same room is maintained for the entire length-of-stay of a booking). They prove that the rectangular ARS is NP-complete.

As a summary of what we have found in the literature, several traits of the room assignment problem are present in related problems that are NP-complete:

- Each room is available for a different period of time.
- each booking can be assigned only to a subset of rooms, depending on its type and attributes.
- The fact that some pairs of bookings request connecting rooms.

To our knowledge, a problem that exhibits the three mentioned features together has not been studied so far.

As we mentioned in the introduction, we are interested in using the room inventory in the most efficient way so that we honor as many requests as possible, and that we are able to accommodate more future walk-in customers. Consistent with what happens in the industry, we do assume that appropriate prices have been offered to each type of customer who made a booking. In other words, we assume that traditional revenue management pricing strategies have already been performed when we perform the room assignment process.

4.3 Mathematic Formulation

We now present the mathematical formulation of the problem we are interested in solving.

A booking i has the following characteristics: arrival date, length-of-stay, requested room type, and requested attributes (if any).

A room j has the following characteristics: initially available time, type of room, attributes of the room.

We will call a room j *compatible* with booking i if the attributes of the room match the following:

1. Room type
2. Requested attributes
3. Availability during length-of-stay.

Notice that rooms may be unavailable for specific nights before the assignment process begins, because of maintenance, the common practice of group-blocking for conventions, weddings, etc., or by guests who are already in house at the time we are assigning rooms for incoming bookings, and whose room is blocked until their departure.

We will define the notation that we will use in the problem.

Let us assume that the assignment will be done for the arrivals happening during the next n days.

Let $x_{ij} = 1$ if i-th booking is assigned to the j-th room; $x_{ij} = 0$ otherwise.

$a_{jk} = 1$ if room j is available at night k; $\delta_{ik} = 0$ otherwise.

R: the set of all rooms.

Denote the set of compatible rooms with a booking i, as the set $R_i \subseteq R$. Notice that the set of the rooms that are not compatible with booking i is $R \setminus R_i$.

R_c: the set of rooms that have a connecting room.

Define h: R_c -> R_c , $h(x) \equiv y$ if and only if room x is connected with room y.

B: the set of all bookings starting during the next n days. This is the set of bookings that we want to assign rather than waiting until the customer is on site.

We will call a booking i' *overlapping* with booking i if there is a time conflict between both bookings that would prevent them from being assigned to the same room.

$B_i \subseteq B$: the set of *overlapping bookings with booking i.*

B_c: the set of bookings that request a connecting room.

Define k: B_c -> B_c , $k(x) \equiv y$ if and only if booking x has a request of a connecting room for booking y. We will say that bookings x and y are connected.

s_i: the length-of-stay of booking i.

n_j: the number of nights available for room j.

We would like to maximize the number of bookings that get a room number assigned, subject to the constraint that the assignment is feasible. The mathematical formulation for the Room Assignment Optimization problem (RAO) can be formulated as:

$$\max \sum_{i \epsilon B} \sum_{j \epsilon R} x_{ij} \tag{4.3.1}$$

subject to:

$\sum_{j \epsilon R} x_{ij} \leq 1$ for all $i \epsilon B$

The booking is assigned to one room at most.

$\sum_{j \epsilon R \setminus R_i} x_{ij} = 0$ for all $i \epsilon B$

The booking is not assigned to a room that is not compatible.

$\sum_{i=1}^{n} \sum_{i \epsilon B} x_{ij} \leq 1$ for all $j \epsilon R$

Each room-night is assigned to at most one booking.

$\sum_{i \epsilon B} s_i x_{ij} \leq n_j$ for all $j \epsilon R$

room-nights assigned are less than the available room-nights for each room.

$x_{ij} + \sum_{i' \epsilon F_i} x_{i'j} \leq 1$ for all $i \epsilon B$

Overlapping bookings are not assigned to the same room as booking i.

$x_{ij} = x_{k(i),h(j)}$ for all $i \epsilon B_c$, all $j \epsilon R_c$

A pair of connected bookings is assigned to a pair of connected rooms.

$x_{ij} = 0, 1$

This problem has a very large size. Given the characteristics of the problem (NP-complete), as one can expect, it is not possible to obtain in a short time an exact solution of this Mixed Integer Problem, even for a small hotel with 150 rooms and a small set of reservations (approximately 280). We made attempts to reduce the size of the problem and to make more tractable the actual Integer Program for commercial software (such as C++ or SAS O.R.). Even with the size reductions we attained, an exact solution could not be found for a small test hotel (150 rooms). Therefore, the best option we found was to create a heuristic method to obtain the solution. The heuristic does not present problems in dealing with an actual size hotel. We describe this methodology in the next section.

4.4 Heuristic method

In this section, we suggest an algorithm that creates a heuristic solution based on sorted lists. In order to obtain a good heuristic solution, it is better to assign the *most scarce* resources *first*. Which resource is more scarce depends on each specific hotel, and the management can help define which resources should be assigned first. An example of this is the view attribute. For hotels close to the seaside, this attribute may be highly demanded. Another example is the connecting rooms requests. Connection requests are more demanded in hotels oriented to families or small groups traveling together. Handling the assignment of connection requests can be hard, given the fact that the inventory of connecting rooms is limited. Therefore, when sorting the attribute list, we would list *connections* first.

We created an algorithm to obtain a heuristic solution to this problem, respecting attribute requests and avoiding overlappings.

We describe on a high level the algorithm before presenting the pseudocode.

The first part of the algorithm initializes all the variables and creates an ordered list of feasible rooms for each booking, except for the connecting room request attribute, which we check in the main loop.

In the main loop, there is a variable κ. This variable is used to run a few iterations of the actual routine that finds a room assignment schedule. Inspired by genetic algorithms, the suggested algorithm runs a routine to propose an assignment of room numbers to bookings and takes note of the bookings that did not get a room assigned; then it starts over the routine to look for a different room assignment schedule *starting with the bookings, that did not get a room assigned in the previous iteration,* and so on, until κ iterations have been performed. Only the assignment that assigns the most bookings is saved. This step is similar to *mutation* in genetic algorithms, which is used in order to "prevent the algorithm to be trapped in a local minimum" (Sivanandam and Deepa, 2007). However, our algorithm would not be classified as a typical genetic algorithm, given that there is no selection or reproduction procedures.

The routine that proposes a possible room assignment schedule takes each booking i and performs the following steps:

- Clear $minR$ and smallest hole variables.
- Check if booking i has not been assigned. If it has, go to next booking. Otherwise, do the following:
- Check each room j in the set R_i. If it is still available for booking i, go to next steps. If not, skip this room and check the next one.
  - Check if booking i has a connected booking. If it does, denote it i_2 and do the following:
    * Check if room j has a connecting room. If it does not, skip room j and examine the next room in the list. If it does, denote the connecting room j_2:
    * Check if j_2 is in the compatible set R_{i2} and if it is still available for booking i_2 length-of-stay. If not, stop checking room j any further

(and room j_2) and check the next one.

  - Check if room j creates a *smaller* hole than the smallest one so far. If it does, save it in $minR$. If it creates no hole (no empty room-nights), do not check any other rooms for booking i.

- If there were rooms available for i ($minR$ not NULL), assign $minR$ to booking i and update room $minR$ availability.

- If booking i has a connected booking i_2, assign the room that connects to $minR$ to booking i_2. Update the room availability.

We now present the pseudocode that performs the steps just outlined:

```
Sort:
B= b_1, b_2, ..., b_n    Sorted set of bookings (sorted by type, connection request,
dayIn, Number of nights, other attributes)
R= r_1, r_2, ..., r_m    Sorted set of rooms (sorted by type, connection request,
dayIn, empty nights, other attributes)
Initialize:
    for all i in B
        Set(b_i->assigned)=0
        Set(b_i ->room)=NULL
        Set(b_i ->CurrentBestRoom)=NULL

    for all j in R
        Clear r_j->bookings
        r_j ->AvailableOn=r_j ->OriginalAvailableOn

    for all i in B
        Create ordered set R_i, where a room r_j is in R_i
        if and only if compatible with booking b_i, i.e.:
                r_j->type = b_i->type
                && b_i->attr in r_j->attr
                && b_i-> dayIn > r_j->availableOn
                && b_i-> dayOut > r_j->blockedFrom
                Order by: type, connection request, dayIn,
                number of nights, other attributes;
                // (connecting room requests will be examined later)
        maxNbBookingsAssigned=0;
```

Main loop:

```
for (κ=0 to MaxIteration; κ++)
    nbBookingsassigned = 0;

    for all i in B
        // if the booking has been assigned skip to the next booking.
        if (b_i ->assigned > 0) continue;
        // initialize minimum hole length for this booking to large number
        minH = 9999;
        // initialize room that minimizes hole for this booking
        minR = NULL;
        for all j in set R_i
            if b_i->dayIn earlier than r_j->availableOn,
then skip this room; // arrival when not available;

            if (b_i2 != NULL) // booking b_i has a connected booking b_i2

                // room r_j has no connecting room, skip to next room
                if (r_j2 == NULL) continue;
                // if room connected to r_j not available,
                // then skip to next room
                if (b_i2->dayIn < r_j2->availableOn) continue;

            // find length of the hole created by b^i in the room r_j
            diff:= b_i->dayIn - r_j->availableOn
            // if smaller than minimum length hole so far:
            if (diff <= minH)

                // save the new minimum length hole
                minH = diff;
                // update room with minimum length hole
                minR = r_j;
                // if booking in room does not create a hole,
                // stop reviewing rooms
                if (minH == 0) break;

         // End for j
        // if no room available for booking, nothing to do,
        // skip to next booking
        if (minR == NULL) continue;
        // Otherwise, assign booking b_i to room
        // with minimum length hole (minR)
        Add b_i to minR->bookings;
        // assign room with minimum length hole minR to booking b_i
        Set(b_i->assignedRoom)=minR;
```

```
            // change room minR availability
            Set(minR->availableOn)=b_i->dayOut;
            // flag booking b_i as assigned
            Set(b_i->assigned)=1;
            // counting number of assigned bookings
            nbBookingsAssigned++;
            // repeat steps above for connecting rooms if necessary
            // if b_i does not have a connected booking, no further action
            if (b_i2 == NULL) continue;
            // Otherwise, determine the room that connects with minR
            // repeat steps above for connected booking and connected room
            minR2= minR->ConnectedRoom;
            Add b_i2 to minR2->bookings;
            Set(b_j2-> room)=minR2;
            Set(minR2-> availableOn)=b_i2->dayOut;
            Set(b_j2-> assigned)=1;
            nbBookingsAssigned++;
        // End for i
        // ======================================
        // If assignment obtained from Current Iteration κ is better,
        // save it in"CurrentBestRoom".
        if (nbBookingsAssigned > maxNbBookingsAssigned)
            maxNbBookingsAssigned = nbBookingsAssigned;
            for all i in B
            b_i->CurrentBestRoom = b_i->room
             // End for i

        // Unassigned bookings go first in the next iteration κ
        Order bookings by: type, dayIn, assigned, connection request,
Number of nights, other attributes

        // Clear data to run another iteration
        for all i in B
            Set(b_i->assigned)=0
            Set(b_i ->room)=NULL


        for all j in R
            Clear r_j->bookings
        Set(r_j->availableOn)=0;

    // End for κ
```

4.5 Computational Results

In order to test this algorithm, we used a production environment data set from a hotel with 1477 rooms, and a set of bookings from the peak season (Easter 2009). The total number of booking requests is 2815. The peak season was chosen because this is when it is harder to give a room to each existing booking; it is easier to assign a room to each booking. The results are summarized in the table.

Table 4.1: An Optimal Solution of Room Assignment to an Actual Hotel

Measurement	Lexicographic	RAO
Number of bookings successfully assigned	2518	2773
Percentage of bookings successfully assigned	89.4%	98.5%
Number of rooms used	1477	1418
Number of holes (intervals of empty nights)	212	69
Average number of empty nights per hole	2	1

From these results, we can observe that our heuristic makes a considerable difference in the number of bookings served when compared with the lexicographic approach. We obtained results that made a difference of approximately 10% (250) more assigned bookings using our heuristic. If we assume that the hotel room costs around $60 per night, the difference between RAO and the lexicographic method is about $15,300 in a 10-day period. Assuming that there is one full month of peak season in the hotel, this is more than $45,900 for that month. These savings come only from the difference in room-nights assigned.

Our heuristic also helps to greatly reduce the number of empty room-nights between bookings, which are room-nights that are less likely to be sold. Notice the potential difference in income caused by those empty room-nights: With the lexicographic approach, we obtain $212 \times 2 = 424$ empty room-nights between bookings. On the other hand, with RAO, we have only 69 empty room-nights between bookings. Assuming that all of these nights can be sold (which is a reasonable assumption during peak time), if room cost is around $60 per night, using the lexicographic method, the inventory that is potentially blocked costs about \$25,400. On the other hand, with the Room Assignment algorithm, this amount is considerably lower: \$4,100. The difference is around \$21,000. These amounts correspond to a period of 10 days. For a peak-season month, this amount would be about \$63,000.

The total income difference from both the total number of assigned rooms and the additional income from selling nights that would be hole-nights with the lexicographic method is about \$100,000 per peak-season month.

Notice that these results are obtained without any computational time problems. Our heuristic only takes a few minutes to run. It is, therefore, a tool that can be used routinely. This tool is useful for any hotel, but it is even more attractive for destinations where occupancy is generally high. New York City, for example, has an average occupancy of 80% year round, according to the statistics provided by NYCgo.com.

4.6 Concluding Remarks

As we saw in the introduction, this problem is NP-complete, and therefore finding an exact solution is computationally expensive. Therefore, we recommended a heuristic algorithm in order to obtain an approximate solution. In order to assign as many rooms as possible, it is better to assign the most scarce resource first. Connecting rooms are generally scarcer than others, and therefore our heuristic gives priority to these requests.

In our tests, we successfully assigned 9% more bookings using our heuristic than with the lexicographic method. We reduced the number of room-nights between bookings by 18%. The planning horizon was for 10 days during the peak season. If if it is possible to sell the room-nights that were part of a “hole” (which is usually the case during peak season), the savings obtained are approximately \$100,000 per month. As a reference, if we assume a 95% occupancy (typical of peak time) for a

1477-room hotel, the income for one day is around $79,750. This means that not using any optimization method during peak season is similar to closing up the hotel for one day.

This is an important difference, especially when occupancy levels are high, which often happens in tourist locales or in towns and cities frequently visited by business travelers.

Finally, we hope that this model motivates further research on the little studied area of room allocation in the lodging industry. The ability to solve a problem of this nature can help businesses in the lodging industry to improve their revenue, given the increment in room utilization (and therefore a greater number of booking requests honored), the decrease in the need for unpaid upgrading, and improved customer satisfaction.

CHAPTER V

Conclusions

In the first essay, joint work with Mark Lewis and David Kaufman, we introduced a new method for load balancing in the case for highly variable service distributions. The method introduced is robust to changes in the parameter settings even in the case where it is not adjusted to optimize the implementation. The most reasonable alternative to our heuristic appears to be a non-idling heuristic. In this case, the question is simply, is the consistency and savings worth the difficulty of implementing our heuristic. In many cases we believe more than 8.5% savings is worth the time to implement our heuristic.

At the same time, we showed that the use of Markov decision processes can mitigate the challenges of a general service time distribution. We believe that the ideas described here can lead to insights for other queueing models. The example of admission controlled $M/G/1$ has already been alluded. Exactly the same intuition holds for service rate control in a $G/M/1$. Of course, these are just the building blocks for more sophisticated models. We note that an extension of the current work is to consider a larger network of queues, and we conjecture that the *two pairing* heuristics described in and Down and Lewis *Down and Lewis* (2006)) would be useful. We leave this for future research.

In the second essay, we studied the problem solved by a manufacturer who faces supplier disruptions. In order to understand the interactions between three strategies (subsidizing the supplier, supplier diversification, and the creation of back-up inventory) we used a simple model with inventory storage costs and shortage penalties. we studied first the one-period model, where there was a finite amount that would make the probability of successful delivery equal to 1, and where the subsidy would have an immediate effect. In this case, we first determined conditions which defined whether diversification was or was not an alternative to be considered when compared to subsidizing a supplier.

In this context, the amount $\widetilde{\omega}_i \equiv \frac{F_i(p(1-\lambda_i)-\lambda_i h)}{(1-\lambda_i)(p+h)}$ defines a threshold which determines if supplier i will be reliable enough for the manufacturer. Below this level, this supplier is not reliable enough and the expected penalty costs for the default are higher than the expected holding costs, and therefore it becomes necessary to diversify. If $p(1-\lambda_i) < h(\lambda_i)$ then every level of subsidy satisfies $w > \widetilde{w}$. This means that supplier i is strong enough to have a reasonable probability of successful delivery, so having this supplier exclusively has a lower cost than diversifying.

Therefore, the conditions $p(1-\lambda_1) < h(\lambda_1)$ and $p(1-\lambda_2) < h(\lambda_2)$ are needed to determine whether diversification may or may not be optimal.

In the case when $p(1-\lambda_1) \geq h(\lambda_1)$ and $p(1-\lambda_2) \geq h(\lambda_2)$, by choosing if having an exclusive supplier or diversifying, the manufacturer is choosing between certainty and uncertainty: the cost of choosing supplier i as an exclusive supplier is the cost of providing full subsidy. The cost of diversification is the cost of excess inventory if no one defaults, plus the penalty cost if both suppliers default. As a consequence, if we assume that each of the available suppliers can satisfy our total demand of supplies, and if the manufacturer prefers to avoid uncertainty, she has a greater incentive to grant subsidies to an exclusive supplier.

In the case when $p(1-\lambda_1) < h(\lambda_1)$ or $p(1-\lambda_2) < h(\lambda_2)$, the manufacturer is choosing between certainty and uncertanty. Certainty in this case is obtained if the manufacturer provides full subsidy. Uncertainty, however, is obtained if the manufacturer simply orders from an *exclusive* supplier without ordering any backup inventory or providing any subsidy.

In this one-period setting, under some circunstances, such as when diversification occurs, we find that under this model, the backup inventory amount is high: the manufacturer orders *twice* the current period demand, even if there is no future demand.

We later considered the case with two periods, while all the other assumptions remained unchanged. We obtained the solution and analyzed it. Depending on the initial parameters, the manufacturer's optimal actions would reflect the flexibility that the introduction of time as a dimension would add to his choice of possible actions. In the example we examined, the penalty cost is initially extremely low. In this situation, the manufacturer wants to have some diversification of the risk concerning the supplies she needs for the next period: she can order them now, pay the holding cost, and if the supplier defaults, she can have one more opportunity to order them again. If the penalty cost increases a lot, the manufacturer will want to actually fund the supplier instead of risking to have to pay the penalty, and she will want to order only for the current period, given the fact that she will want to save the holding cost for ordering for the second period too. Finally, if the penalty cost is very high (region 4), the manufacturer will opt for the most protection she can get against the penalty: full funding in period 1, and ordering for both periods now and saving the future funding costs. In summary, the time horizon gives the manufacturer some flexibility. Depending on the required subsidy amount, the holding costs, and the shortage penalties, the supplier has the option to subsidize, to have backup inventory or to wait and see if subsidy will be required to satisfy the accumulated demand on period 2.

We then switched to study the case when we had a two-period setting, but the subsidy would not be achieved in the same period that it was granted. Additionally, we relaxed the assumption that a finite amount F existed that would guarantee that no default would happen. We assumed that the demand in the second period was known and deterministic. In this setting, there is a threshold $\widehat{\theta^1} = g^{-1}(\frac{h}{\rho p})$. If the subsidy is lower than this amount, it is optimal to order $[D^1 - x^1]^+$. Intuitively, once the manufacturer has given enough subsidy to the supplier, there is enough confidence that the supply will be delivered next period, and therefore it is not necessary to take the extra precaution of ordering for the future today. In other words, $\widehat{\theta^1}$ actually is the threshold that defines a *strong* supplier, in the sense that we can trust that the supplier will deliver the supplies, and therefore, there is no need to add backup inventory for the future.

This partition allowed to solve the problem optimally, obtaining closed form solutions. With these closed form solutions, we performed a sensitivity analysis.

We found that although changes in the optimal subsidy may be monotonic, this is not necessarily the case for the order size. For example, for the case when the demand for either period increases, the subsidy increases, and the order size may increase for a while, until the supplier becomes so reliable that it is no longer necessary to order for both periods, and the order size jumps down. A second example is when the current inventory amount increases. The subsidy amount decreases, but this forces the order-up-to amount to go up, by ordering only from the current period to ordering for both.

Relaxing the assumption of deterministic demand in the second period does not prevent us from giving a set of equations that characterize the solution. However, a closed form solution can no longer be found so that a similar sensitivity analysis be performed as in the deterministic demand case. There is no *rule of thumb* to determine if the subsidy or the order will be greater or lower than their counterpart in the deterministic case. In the stochastic case, the conclusions of the sensitivity analysis for the deterministic demand may or may not hold. However, despite this limitation, the envelope theorem enables us to determine in some cases the direction of the changes in the costs (the objective function), which is good news given the limitations we face in this case. Using the envelope theorem, and computing the corresponding derivatives, we obtain that whenever α, h or p increase, the optimal cost will increase too.

In the third essay, which is joint work with Yihua Li, we studied a hotel room assignment problem, which is generally performed by the front desk staff on the arrival day using a lexicographic approach. We found features of the room assignment problem that makes it NP complete, and therefore, finding an exact solution is computationally expensive. Therefore, we recommended a heuristic algorithm in order to obtain an approximate solution. In order to assign as many bookings as possible, it is better to assign the most scarce resource first. Connecting rooms are generally more scarce than other resources, and therefore our heuristic gives priority to these requests.

In our tests, we successfully assigned 9% more bookings more using our heuristic than with the lexicographic method. We reduced the number of room nights between bookings by 18%. The planning horizon was for 10 days in peak season. If if it is possible to sell the room-nights that were part of a "hole" (which is usually the case during peak season), the savings obtained are of approximately $100,000 per month. As a reference, if we assume a 95% occupancy (typical of peak time), for a 1477 room hotel, the income for one day is around $79,750. This means that not using any optimization method during peak season is similar to closing up the hotel for one day.

This is an important difference, especially when occupancy levels are high, which often happens in tourist locales or in towns and cities frequently visited by business travelers.

The ability to solve a problem of this nature can help businesses in the lodging industry to improve their revenue, given the increment in room utilization (and therefore a greater number of booking requests honored), the decrease in the need of unpaid upgrading, and the *increment* in customer satisfaction.

APPENDIX

APPENDIX

Solution to problem 3.3.1

Proposition A.1. *The optimal strategy to problem (3.3.1), and the corresponding optimum value, is as follows:*

1) If $\widehat{\theta^1} \leq \theta^{1,D1+D2}$ *, then* $z^{1*} = [D^1 - x^1]^+$*,* $\theta^{1*} = \max\{\theta^{1,D1}, c^1\}$ *and*

$$\begin{aligned} v^*(1, x^1, c^1) = &\left[p[D^1 - x^1]^+ + h[x^1 - D^1]^+\right]\left[1 - \pi(c^1)\right] + \alpha(\theta^{1*} - c^1) \\ &+ \rho\left\{g(\theta^{1*})p(D^2)\pi(c^1) + \left[g(\theta^{1*})p[D^2 + D^1 - x^1]^+\right.\right. \\ &\left.\left.+h[x^1 - D^1 - D^2]^+\right]\left[1 - \pi(c^1)\right]\right\}. \end{aligned}$$

Let $\Pi^{1,\theta 1}$ *be the policy associated with ordering* $z^1 = D^2 + D^1 - x^1$ *and a subsidy of* $\theta^{1,D1}$*. Let* $\Pi^{2,\max 2}$ *be the policy associated with ordering* $z^1 = D^2 + D^1 - x^1$ *and a subsidy of* $\max\{\theta^{1,D1+D2}, c^1\}$*.*

2ai)

If $\theta^{1,D1+D2} \leq \widehat{\theta^1} \leq \theta^{1,D1}$*,* $c^1 \leq \widehat{\theta^1}$ *and* $v^{\Pi^{1,\theta 1}}(1, x^1, c^1) \leq v^{\Pi^{2,\max 2}}(1, x^1, c^1)$*, then*

$z^{1*} = [D^1 - x^1]^+$, $\theta^{1*} = \theta^{1,D1}$ *and*

$$v^*(1, x^1, c^1) = \left[p[D^1 - x^1]^+ + h[x^1 - D^1]^+\right]\left[1 - \pi(c^1)\right] + \alpha(\theta^{1*} - c^1) + \rho\left\{g(\theta^{1*})p(D^2)\pi(c^1) + \left[g(\theta^{1*})p[D^2 + D^1 - x^1]^+ + h[x^1 - D^1 - D^2]^+\right]\left[1 - \pi(c^1)\right]\right\}.$$

2aii) If $\theta^{1,D1+D2} \leq \widehat{\theta^1} \leq \theta^{1,D1}$, $c^1 \leq \widehat{\theta^1}$ *and* $v^{\Pi^{2,\max 2}}(1, x^1, c^1) \leq v^{\Pi^{1,\theta 1}}(1, x^1, c^1)$, *then* $z^{1*} = [D^2 + D^1 - x^1]^+$, $\theta^{1*} = \max\{\theta^{1,D1+D2}, c^1\}$ *and*

$$v^*(1, x^1, c^1) = \left[h(D^2)I_{[D^2+D^1>x^1]} + h(x^1 - D^1)^+ I_{[D^2+D^1<x^1]}\right]\pi(c^1) + \left[p[D^1 - x^1]^+ + h[x^1 - D^1]^+\right]\left[1 - \pi(c^1)\right] + \alpha(\theta^{1*} - c^1) + \rho\left\{\left[g(\theta^{1*})p[D^2 + D^1 - x^1]^+ + h[x^1 - D^1 - D^2]^+\right]\left[1 - \pi(c^1)\right]\right\}.$$

2b) If $\theta^{1,D1+D2} \leq \widehat{\theta^1} \leq \theta^{1,D1}$, $c^1 > \widehat{\theta^1}$, *then* $z^{1*} = [D^1 - x^1]^+$, $\theta^{1*} = c^1$ *and*

$$v^*(1, x^1, c^1) = \left[p[D^1 - x^1]^+ + h[x^1 - D^1]^+\right]\left[1 - \pi(c^1)\right] + \rho\left\{g(\theta^{1*})p(D^2)\pi(c^1) + \left[g(\theta^{1*})p[D^2 + D^1 - x^1]^+ + h[x^1 - D^1 - D^2]^+\right]\left[1 - \pi(c^1)\right]\right\}.$$

3a) If $\widehat{\theta^1} \geq \theta^{1,D1}$ *and* $c^1 \leq \widehat{\theta^1}$ *then* $z^{1*} = [D^2 + D^1 - x^1]^+$, $\theta^{1*} = \max\{\theta^{1,D1+D2}, c^1\}$, *and*

$$v^*(1, x^1, c^1) = \left(h(D^2)I_{[D^2+D^1>x^1]} + h(x^1 - D^1)^+ I_{[D^2+D^1<x^1]}\right)\pi(c^1) + \left[p[D^1 - x^1]^+ + h[x^1 - D^1]^+\right]\left[1 - \pi(c^1)\right] + \alpha(\theta^{1*} - c^1) + \rho\left\{\left[g(\theta^{1*})p[D^2 + D^1 - x^1]^+ + h[x^1 - D^1 - D^2]^+\right]\left[1 - \pi(c^1)\right]\right\}.$$

3b) If $\widehat{\theta^1} \geq \theta^{1,D1}$ *and* $c^1 > \widehat{\theta^1}$ *then* $z^{1*} = [D^1 - x^1]^+$, $\theta^{1*} = c^1$, *and*

$$\begin{aligned} v^*(1, x^1, c^1) = & \left[p[D^1 - x^1]^+ + h[x^1 - D^1]^+\right]\left[1 - \pi(c^1)\right] \\ & + \rho\left\{g(\theta^{1*})p(D^2)\pi(c^1) + \left[g(\theta^{1*})p[D^2 + D^1 - x^1]^+\right.\right. \\ & \left.\left. + h[x^1 - D^1 - D^2]^+\right]\left[1 - \pi(c^1)\right]\right\}. \end{aligned}$$

Proof. Recall *subproblems a)* and *b)*, which we have solved in Propositions III.11 and III.12. These are *unconstrained* problems. However, as pointed out, when imposing the constraint $\theta^1 \leq \widehat{\theta^1} = g^{-1}(\frac{h}{\rho p})$ to *subproblem a)* and $\theta^1 \geq \widehat{\theta^1}$ to *subproblem b)*, the solution to each of them may change (now the optimal subsidy for each problem may be at $\widehat{\theta^1} = g^{-1}(\frac{h}{\rho p})$, the border point). We will call each of these feasible sets "the feasible regions for *subproblems a)* and *b)*,"respectively. Note, as well, that the constraints for the subsidy level that define the feasible regions for each subproblem are not the only constraints we are facing. From the formulation of problem (3.3.1) we see that $\theta^1 \geq c^1$ is another feasibility constraint. Therefore, we need to take it into consideration.

We divided the solution in three cases, the ones that we depicted in Figures 3.5, 3.6, 3.7.

Case 1) refers to the case when the point at the border s A and B, $\widehat{\theta^1}$, is to the left of $\theta^{1,D1+D2}$. As illustrated in the diagram, the optimal solution for (*subproblem b)* ($z^1 = [D^1 - x^1]^+$ and $\theta^{1*} = \max\{\theta^{1,D1}, c^1\}$) clearly outperforms the optimal for *subproblem a)*.

Case 2) refers to the case when the border point $\widehat{\theta^1}$ is between $\theta^{1,D1+D2}$ and $\theta^{1,D1}$. Subcase 2a) refers to the case when both subregions intersect the feasible region $\theta^1 \geq c^1$, and therefore, both minima must be compared (therefore, we have parts (i) and (ii), depending which of the two minima outperforms the other). Subcase 2b) refers to the case when $c^1 > \widehat{\theta^1}$, where only the subregion for *subproblem b)* intersects the region $\theta^1 \geq c^1$. Therefore, we only need to consider the solution to *subproblem b).*

Case 3) refers to the case when $\widehat{\theta^1}$ is to the right of $\theta^{1,D1}$. Before considering the constraint $\theta^1 \geq c^1$, we see that the optimal solution for *subproblem a)* ($z^1 = [D^2 + D^1 - x^1]^+$ and $\theta^{1*} = \max\{\theta^{1,D1+D2}, c^1\}$) outperforms the optimal for *subproblem b).* Subcases 3a) and 3b) are considered separately just to distinguish when the optimal for unconstrained *subproblem a)* is actually attainable (subcase (a.1)) or not (subcase (a.2)) given the original constraint $\theta^1 \geq c^1$. In this last case, the best we can do is to not subsidize, given the fact that the initial level of assets c^1 is above $\theta^{1,D1+D2}$. ∎

BIBLIOGRAPHY

BIBLIOGRAPHY

Abramson, D. (1991), Constructing school timetables using simulated annealing: sequential and parallel algorithms, *Management Science*, *37*(1), 98–113, doi:http://dx.doi.org/10.1287/mnsc.37.1.98.

Babich, V. (2006), Vulnerable options in supply chains: Effects of supplier competition, *Naval Research Logistics 53 (2006): 656-673.*, *53*, 656/673.

Babich, V. (2008), Independence of Capacity Ordering and Financial Subsidies to Risky Suppliers, *SSRN eLibrary*.

Bakos, J. Y., and E. Brynjolfsson (1993), Information technology, incentives, and the optimal number of suppliers, *J. Manage. Inf. Syst.*, *10*(2), 37–53.

Beard, C., and V. Frost (2001), Prioritized resource allocation for stressed networks, *IEEE-ACM Transactions on networking*, *9*(5), 618–633.

Best Software, Inc. (2009), Automated hotel website, `http://www.bestsoftinc.com/products.html`.

Bierwirth, C., and D. C. Mattfeld (1999), Production scheduling and rescheduling with genetic algorithms, *Evolutionary Computation*, *7*(1), 1–17.

Bitran, G., and R. Caldentey (2003), An overview of pricing models for revenue management, *Manufacturing and Service Operations Management*, *5*(3), 203–229.

Bitran, G., and S. Gilbert (1996), Managing hotel reservations with uncertain arrivals, *Operations Research*, *44*(1), 35–49.

Blazewicz, J., W. Domschke, and E. Pesch (1996), The job shop scheduling problem: Conventional and new solution techniques, *European Journal of Operational Research*, *93*(1), 1–33.

Brucker, P., and L. Nordmann (1194), The k-track assignment problem, *Computing*, *52*(2), 97–122.

Burke, E., and S. Petrovic (2002), Recent research directions in automated timetabling, *European Journal of Operational Research*, *140*(2), 266–280.

Burke, E., K. Jackson, J. Kingston, and R. Weare (1997), Automated timetabling: The state of the art, *The Computer Journal*.

Business Software Solutions, Inc. (2009), Hotel and property management software, http://www.businessoftware.com/hotelscreens.asp.

Ciarallo, F., R. Akella, and T. Morton (1994), A periodic review, production planning model with uncertain capacity and uncertain demand-optimality of extended myopic policies, *Management Science*, pp. 320–332.

Coffman, E. G., and J. Bruno (1976), *Computer and job-shop scheduling theory*, Wiley, New York.

ConRunner.net (2009), Conrunner, a why and how-to reference for convention organizers, http://www.conrunner.net/wiki/index.php?journal=Hotel_Liaison.

Cross, R. (1998), *Revenue management: hard-core tactics for market domination*, Broadway.

Crovella, M., M. Harchol-Barter, and C. Murta (1998a), Task assignment in a distributed system: improving performance by unbalancing load, *Performance Evaluation Review*, *26*(1), 268–269.

Crovella, M., M. Taqqu, and A. Bestavroz (1998b), Heavy tailed probability distributions in the world wide web, in *A practical Guide to Heavy Tails: Statistical Techniques & Applications*, edited by R. Adler, R. Feldman, and M. Taqqu, pp. 3–26, Birkhäuser.

Davis, L. (1985), Job shop scheduling with genetic algorithms, in *Proceedings of the 1st International Conference on Genetic Algorithms*, pp. 136–140, L. Erlbaum Associates, Inc., Hillsdale, NJ, USA.

Dell'Amico, M., and M. Trubian (1993), Applying tabu search to the job-shop scheduling problem, *Annals of Operations Research*, *41*(1-4), 231–252.

Down, D., and M. Lewis (2006), Dynamic load balancing in parallel queueing systems: stability and optimal control, *The European Journal of Operational Research*, *168*(2), 509–519.

Duin, C. W., and E. van Sluis (2006), On the complexity of adjacent resource scheduling, *J. of Scheduling*, *9*(1), 49–62.

Elmohamed, M. A. S., P. D. Coddington, and G. Fox (1998), A comparison of annealing techniques for academic course scheduling, in *PATAT '97: Selected papers from the Second International Conference on Practice and Theory of Automated Timetabling II*, pp. 92–114, Springer-Verlag, London, UK.

Erben, W., and J. Keppler (1996), A genetic algorithm solving a weekly course-timetabling problem, in *Selected papers from the First International Conference on Practice and Theory of Automated Timetabling*, pp. 198–211, Springer-Verlag, London, UK.

eZee Technosys Pvt. Ltd. (2009), Assigning rooms to room owners - hotel software ezee frontdesk help, `http://www.ezeefrontdesk.com/`.

Garey, M. R., D. S. Johnson, and R. Sethi (1976), The complexity of flowshop and jobshop scheduling, *Mathematics of Operations Research*, *1*(2), 117–129.

Gross, D., J. F. Shortle, M. J. Fischer, and D. M. B. Masi (2000), Difficultied in simulating queues with Pareto service, in *Proceedings of the 2002 Winter Simulation Conference*, pp. 407–415.

Harchol-Balter, M., and A. Downey (1997), Exploiting process lifetime distributions for dynamic load balancing, *ACM Transactions on Computer Systems*, *15*(3), 253–285.

He, Q., and M. Neuts (2002), Two m/m/1 queues with transfers of customers, *Queueing systems*, *42*(4), 377–400.

Heyde, C., and K. S.G. (2004), On the controversy over tailweight of distributions, *Operations Research Letters*, *32*, 399–408.

Jain, A. S., and S. Meeran (1999), Deterministic job-shop scheduling: Past, present and future, *European Journal of Operational Research*, *113*(2), 390–434.

Jüttner, U., H. Peck, and M. Christopher (2003), Supply chain risk management: outlining an agenda for future research, *International Journal of Logistics*, *6*(4), 197–210.

Kolen, A., J. Lenstra, C. Papadimitriou, and C. Spieksma (2007), Interval scheduling: A survey, *Naval Research Logistics*, *54*(5), 530–543.

Kulkarni, V. (1999), *Modeling, Analysis, Design, and Control of Stochastic Systems*, Springer-Verlag, New York.

Kumar, V. (1992), Algorithms for constraint-satisfaction problems: A survey, *The AI Magazine*, *13*(1), 32–44.

Lewis, M. (2001), Average optimal policies in a controlled queuing system with dual admission control, *Journal of Applied Probability*, *38*, 369–385.

Lippman, S. (1975), Applying a new device in the optimization of exponential queueing systems, *Operations Research*, *23*, 687–710.

Luan, F., and X. Yao (1996), Solving real world lecture room assignment problems by genetic algorithms, *Complexity International*.

Lum, K., and B. Moyer (2001), Gross domestic product by industry for 1998-2000, *Bureau of Economic Activity, Survey of Current Business*, pp. 17–33.

Malhotra, R. (1997), *Encyclopaedia of Hotel Management and Tourism*, Anmol Publications PVT. LTD.

McGill, J., and G. van Ryzin (1999), Revenue management research overview and prospects, *Transportation Science*, *33*(2), 233–256.

Meyn, S., and R. Tweedie (1993), *Markov Chains and Stochastic Stability*, 1 ed., Springer–Verlag, New York.

NYC & Company, Inc. (2009), NYCGO/NYC statistics, `http://www.nycgo.com/`.

Papadimitriou, C. (1982), Private communication to J.K. Lenstra.

Paxson, V., and S. Floyd (1995), Wide-area traffic: The failure of Poisson modeling, *IEEE/ACM Transactions on Networking*, *3*(3), 226–244.

Puterman, M. (1994), *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, John Wiley & Sons, Inc., New York.

Resort Data Processing, Inc. (2009), Reservation software for timeshares, `http://www.resortdata.com/`.

Rice, J. B., and F. Caniato (2003a), Building a secure and resilient supply network, *Supply Chain Management Review*, pp. 22 – 30.

Rice, J. B., and F. Caniato (2003b), Supply chain response to terrorism: Creating resilient and secure supply chains, report by MIT Center for Transportation and Logistics.

Riska, A., E. Smirni, and G. Ciardo (2000), Analytic modeling of load balancing policies for tasks with heavy-tailed distributions, in *Proceedings of the Second International Workshop on Software and Performance – WOSP 2000*, pp. 147–157.

Russell, S., and P. Norvig (1995), *Artificial Intelligence: A Modern Approach*, Prentice Hall, New Jersey.

Sadeh, N., and M. Fox (1996), Variable and value ordering heuristics for the job shop constraint satisfaction problem, *Artificial Intelligence*, *86*(1), 1–41.

Sanlaville, E., and G. Schmidt (1998), Machine scheduling with availability constraints, *Acta Informatica*, *35*(9), 795–811.

Schiex, T., H. Fargier, and G. Verfaillie (1995), Valued constraint satisfaction problems: Hard and easy problems.

Schmitt, A., and L. Snyder (2009), Infinite-horizon models for inventory control under yield uncertainty and disruptions, `http://ssrn.com/abstract=1115443`.

Sennott, L. (1999), *Stochastic Dynamic Programming and the Control of Queueing Systems*, John Wiley & Sons, Inc., New York.

Sheffi, Y. (2001), Supply chain management under the threat of international terrorism, *International Journal of Logistics Management*, *12*(2), 1 – 11.

Shimkin, N., and A. Shwartz (1989), Control of admission and routing in parallel queues operating in a random environment, in *Proceedings of the 28th Conference on Decision and Control*, vol. 2, pp. 1064–1065, IEEE, Tampa.

Shirazi, B., A. Hurson, and K. Kavi (1995), *Scheduling and Load Balancing in Parallel and Distributed Systems*, 1 ed., Wiley–IEEE Computer Society Press.

Siggelkow, N. (2001), Change in the presence of fit: The rise, the fall, and the renaissance of liz claiborne, *The Academy of Management Journal*, *44*(4), 838–857.

Sigman, K. (1999), A primer on heavy-tailed distributions, *Queueing Systems*, *33*(1–3), 261–275.

SoftAcid (2009), Hotel reservation system script, `http://www.softacid.net/`.

Talluri, K., and G. van Ryzin (2004), *The Theory and Practice of Revenue Management*, Springer.

Tomlin, B. (2009), Impact of supply learning when suppliers are unreliable, *Manufacturing & Service Operations Management*, *11*(2), 192–209.

Tomlin, B., and Y. Wang (2005), On the value of mix flexibility and dual sourcing in unreliable newsvendor networks, *Manufacturing & Service Operations Management*, *7*(1), 37–57.

Turkington, D. (2007), *Mathematical Tools for Economics*, Blackwell Publishing, Malden, Massachusetts.

Ueda, H., D. Ouchi, K. Takahashi, and T. Miyahara (2001), A co-evolving timeslot/room assignment genetic algorithm technique for university timetabling, in *PATAT '00: Selected papers from the Third International Conference on Practice and Theory of Automated Timetabling III*, pp. 48–63, Springer-Verlag, London, UK.

Vaessens, R., E. Aarts, and J. Lenstra (1996), Job shop scheduling by local search, *INFORMS Journal on Computing*, *8*(3), 302–317.

van Laarhoven, P., E. Aarts, and J. Lenstra (1992), Job shop scheduling by simulated annealing, *OPERATIONS RESEARCH*, *40*(1), 113–125.

VCL Components (2009), IT Edge softwares, `http://www.resortdata.com/`.

Vogel, H. (2001), *Travel industry economics: a guide for financial analysis*, Cambridge University Press.

Wang, Y., and Y. Gerchak (1996), Periodic review production models with variable capacity, random yield, and uncertain demand, *Management Science*, *42*(1), 130 – 137.

Wang, Y., W. Gilland, and B. Tomlin (2009), Mitigating Supply Risk: Dual Sourcing or Process Improvement?, *MANUFACTURING SERVICE OPERATIONS MANAGEMENT*, p. msom.1090.0279, doi:10.1287/msom.1090.0279.

Wang, Y.-T., and R. Morris (1985), Load sharing in distributed systems, *IEEE Transactions on Computers*, *34*, 204–217.

Weatherford, L., and S. Bodily (1992), A taxonomy and research overview of perishable-asset revenue management: yield management, overbooking, and pricing, *Operations Research*, *40*(5), 831–844.

White, J. (2005), Ford to pay up to $1.8b on Visteon, Wall Street Journal, May 26, A3.

Xu, T., A. Desrochers, and R. Graves (2003), Hyperexponential-based network traffic model for distributed manufacturing, in *IEEE International Conference on Systems, Man and Cybernetics*, vol. 4, pp. 3452–3457, IEEE.

Yano, C. A., and H. L. Lee (1995), Lot sizing with random yields: a review, *Operations Research*, *43*(2), 311 – 334.

Yu, H., A. Zeng, and L. Zhao (2009), Single or dual sourcing: decision-making in the presence of supply chain disruption risks, *Omega*, *37*(4), 788–800.

Yum, T., and L. Hua-chun (1984), Adaptive load balancing for parallel queues with traffic constraints, *IEEE Transactions on Communications*, *32*(12), 1339–1342.

Yum, T.-S., and M. Schwartz (1981), The join-biased-queue rule and its application to routing in computer communication networks, *IEEE Transactions on Communications*, *29*(4), 505–511.