ABSTRACT

Proposed Title of Dissertation:	ANALYSIS OF DISTRIBUTION-FREE METHODS FOR REVENUE MANAGEMENT
	Huina Gao, Ph.D. candidate, 2008
Dissertation directed by:	Professor Michael O. Ball Professor Itir Karaesmen R.H. Smith School of Business

Revenue management (RM) is one area of research and practice that has gained significant attention in the past decade. The practice originated in the airline industry, where the idea was to maximize revenues obtained from a fixed amount of resources through differentiation/segmentation and strategic use of pricing and capacity. While many of the research models take into account uncertainty, the uncertainty is modeled using random variables and known probability distributions, which is often difficult to estimate and prone to error for a variety of reasons. For instance, demand patterns can fluctuate substantially from the past,and characterizing demand from censored data is challenging. This dissertation focuses on the multi-fare single resource (leg) problem in RM. We consider the "limited information" case where the demand information available consists of lower and upper bounds rather than a characterization of a particular probability distribution or stochastic process. We first investigate the value of the amount and type of information used in solving the single-leg RM problem. This is done via extensive computational experiments.

Our results indicate that new robust methods using limited information perform comparably to other well-known procedures. These robust policies are very effective and provide consistent results, even though they use no probabilistic information. Further, robust policies are less prone to errors in modeling demand. Results of our preliminary computations justify the use of robust methods in the multi-fare single-leg problem.

We next apply this distribution-free approach to a setting where progression of demand is available through time-dependent bounds. We do not make any further assumptions about the demand or the arrival process beyond these bounds and also do not impose a risk neutrality assumption. Our analytical approach relies on competitive analysis of online algorithms, which guarantee a certain performance level under *all* possible realizations within the given lower and upper bounds. We extend the robust model from a problem using static information into a dynamic setting, in which time-dependent information is utilized effectively. We develop heuristic solution procedures for the dynamic problem. Extensive computational experiments show that the proposed heuristics are very effective and provide gains over static ones.

The models and computations described above assume a single airline, disregarding competition. As an extension of robust decision-making, in the third part of this dissertation, we analyze a model with two airlines and two fare classes where the airlines engage in competition. The model does not use any probabilistic information and only the range of demand in each fare-class is known. We develop a game-theoretic model and use competitive analysis of online algorithms to study the model properties. We derive the booking control policies for both centralized and decentralized models and provide additional numerical results.

ANALYSIS OF DISTRIBUTION-FREE METHODS FOR REVENUE MANAGEMENT

by

Huina Gao

Dissertation submitted to the Faculty of the Graduate School of the University of Maryland, College Park in partial fulfillment of the requirements for the degree of Doctor of Philosophy
2008

Advisory Committee:
Professor Michael O. Ball, Chair/Advisor
Professor Itir Karaesmen, Co-Advisor
Professor Gilvan Souza
Professor Zhi-long Chen
Professor Steve A Gabriel

Dedicated to

Miles S. Gu, my lovely son

Liyan Gu, my husband

Peng Gao and Xiaoling Zhen, my parents

Huang Gu and Juying Huang, my parents-in-law

Acknowledgments

The writing of a dissertation is obviously not possible without the guidance and support of many people. I owe my gratitude to all the people who have made this thesis possible and because of whom my graduate experience has been one that I will cherish forever.

First and foremost I'd like to thank my advisor, Professor Michael Ball and my co-advisor, Professor Itir Karaesmen for their dedication to this dissertation. I have worked with them for over five years and they are always available for invaluable advice and support. It has been my honor to work under the guidance of them and to learn from such extraordinary individuals.

I am also grateful to Professor Zhi-long Chen, Professor Gilvan Souza and Professor Steve Gabriel for being my committee members. Their insightful comments on this dissertation are greatly appreciated.

I would like to specially thank Professor Kislaya Prasad and my colleagues, Yingjie Lan and Inbal Yahav for numerous debates and discussions that helped me improve my knowledge in this area.

Table of Contents

cn

List of Tables

List of Figures

Chapter 1

Introduction

Revenue management (RM) is viewed by many as among the most important management science and operation research concepts. Originally known as yield management and developed in the airline industry in the wake of deregulation in the late 1970s, the main idea was to maximize revenues obtained from a fixed amount of resources through differentiation/ segmentation and strategic use of pricing and capacity. The fare-class allocation (or seat inventory control) practice in RM became an industry standard after American Airlines launched Ultimate Super Saver fares in an effort to compete with low cost carrier PeopleExpress. To prevent severe revenue losses in competition, the sale of these highly discounted fares had to be carefully controlled by American Airlines. Other service industries in the hospitality and transportation sectors adapted similar practices quickly. Both industries have observed significant revenue increase using such tactics (Alstrup et al.,1989 Smith et.al,1992, Geragthy and Johnson,1996, Yeoman and Ingold,1997.) In fact, there are other well-known RM success stories from the broadcasting industry (Fox, 1992), freight/Cargo (Kasilingam, 1996), cruise ships (Hoseason, 2000), theaters and sporting venues (Leibs, 2001, Oberwetter, 2001). More recently, business-to-business services (Boyd and Bilegan, 2003) and manufacturing companies started applying RM to manage their resources effectively , where effective use of forecasting, pricing, and inventory control resulted in dramatic improvements.

The challenge of RM is to sell the right resources to the right customer at the right time for the right price. A RM system requires forecasts of quantities such as demand, price sensitivity, and cancelation probabilities, and its performance depends critically on the quality of these forecasts. Experts mention 80-90% of the time in any RM implementation is spent on data gathering, data analysis and forecasting. However, forecasting is challenging because the data collected is censored and not necessarily representative of the true demand. Furthermore, historical data are not available for new products such as new routes or fare products for an airline and new properties for a hotel company. Improving forecasting and estimation can significantly increase the quality of pricing and capacity control decisions, affecting the bottom line. Indeed, some industry estimates suggest that a 20% reduction of forecast error can translate into 1% incremental increase in revenue (Pölt,1998).

To-date, the majority of the successful RM systems rely heavily on the use of demand information and experts advise against rushing to optimization without first fully implementing accurate demand forecasting (Lahoti, 2002). Likewise, research models require a reasonably accurate characterization of demand to optimize revenues. In situations where demand patterns are stable and historical demand information is available, this requirement can be met. On the other hand, in the case of new products or situations where, for a variety of reasons, demand patterns might fluctuate substantially from past history, characterization of demand is difficult - if not impossible. Lennon (2004) mentions lack of data and naive forecasts based on inadequate data as two important factors limiting realistic application of RM to a new industry. Even in industries such as the airlines and hospitality where RM has been effectively used for decades, these two factors commonly cause problems for new products (e.g., new routes flown by an airline, new properties added to a hotel chain). On the other hand, RM of even old products in airlines and hotels is not easy; obtaining high quality aggregate or disaggregate forecasts from censored data remains a challenge (see for e.g., McGill and van Ryzin, 1999)

Despite the need for robust methods and approaches that do not rely heavily on demand information, research in that direction has been scarce. Traditional research models and analysis rely on several restrictive and possibly unrealistic assumptions about demand such as independence and stationarity (see survey articles of McGill and van Ryzin, 1999, Bitran and Caldentey, 2003, and book by Talluri and van Ryzin, 2004a). This dissertation consider the classical single resource (leg) multi-fare problem in RM and provides alternative forms of controlling the bookings when demand information is limited. No assumptions are made about the demand or the arrival process, and no risk neutrality assumption is enforced. The analytical approach is called "robust", because it relies on competitive analysis of online algorithms and can guarantee a certain performance level under all possible input sequences. The only information available consists of the lower and upper bounds on demand. In many practical cases, the lower and upper bounds on demand are easier to obtain rather than the probability distribution of demand or an estimate of an arrival process over time. In Chapter 2, using extensive computational experiments, we show that the average revenues obtained by such robust policies in simulation studies are comparable to other well-known procedures. The robust policies are very effective and provide consistent results, even though they use no probabilistic information. Further, robust policies are not prone to errors in modeling demand.

In air travel, the arrival rate of booking requests in different fare classes vary over time. Analysis of historical data has shown that demand for a particular class generally follows class-specific S-shaped (cumulative) booking curves. Actual booking patterns exhibit significant variations by season, time of day, day of week, and so on. In the hospitality industry, different customer segments have different booking curves (leisure travelers versus business travelers). In Chapter 3, we consider the case where uncertainty is incorporated in a distribution-free manner and the progression of demand information is conveyed through time-dependent bounds. The analysis extends robust booking control from a problem using aggregated static information into a dynamic setting, in which time-dependent information is utilized effectively.

All of the models and methods in Chapter 1 and 2 focus on a single decision-maker and a single flight. When we consider multiple flights belonging to competing airlines, then the decisions that arise out of the resulting game can differ significantly from seat allocations that would be optimal for a single decision-maker with control over a single-flight. There is limited research that considers the RM problem in a competitive framework, especially focusing on the robust policies. In Chapter 4, we study the robust capacity control problem of RM in a competitive framework, using game theory and competitive analysis. We compare decentralized versus centralized decisions and provide numerical experiments to illustrate the effect of competition when sellers use distribution-free methods. This is interesting both from a practical and theoretical point of view.

In summary, new approaches and models for the multi-fare single resource RM problem with limited information are studied in this dissertation. Extensive experiments show the effectiveness of these new methods compared with the traditional methods. Future research directions are discussed in Chapter 5.

Chapter 2

Value of Information in Single-leg Revenue Management

The single-leg multi-fare booking problem seeks to optimally allocate the capacity of a resource to different classes of demand. This allocation must be done dynamically as demand materializes and with considerable uncertainty about the quantity or composition of future demand. Two prototypical examples are controlling the sale of different fare classes on a single flight leg of an airline schedule and the sale of hotel rooms for a given date at different rate classes. In reality, many quantity-based RM problems fall into the mere general class of network RM problems, but in practice, they are still frequently solved as a collection of single-resource problems.

2.1 Literature Review of the Single-leg Booking Control Problem

The booking (or capacity) control problem for a single resource has long been studied in the RM literature. Littlewood (1972) considered two fare classes and assumed product is sold in a low-before-high **(LBH)** manner; i.e., demand in the lowest fare class arrives first. He showed how the booking limit for the low-fare class can be determined once the probability distribution of the demand for the high-fare class and the fares are given. Belobaba (1987,1989) discussed heuristic extensions of Littlewoods rule to multiple fare classes, again assuming (i) LBH arrivals (arrivals are monotonic in fare-classes) and (ii) the probability distribution of demand in each fare class is known. Curry (1990), Wollmer (1992) and Brumelle and McGill (1993) make similar assumptions on demand (and so do Li and Oum (2002) when they discuss the equivalence of the results obtained in these three papers). Robinson (1995) relaxed the LBH assumption in his analysis but assumed requests for different fare classes arrive at non-overlapping intervals. Lee and Hersh (1993) introduced a dynamic programming formulation for the multiple fare class problem by relaxing the LBH assumption and assuming the demand in each fare-class is characterized by a stochastic process. Lautenbacher and Stidham (1999) address the static (i.e. arrivals are LBH) and dynamic (i.e. the arrival sequence is not ordered by fare-class) problems for multiple fares by analyzing the underlying discrete time Markov Decision Process (MDP). More recently, Talluri and van Ryzin (2004b) analyzed the multiple fare class problem based on consumer choice. While this approach is more sophisticated and realistic than many, it requires information not only on the arrival process but also on the choice behavior. There are several other papers on the single-leg problem with various assumptions on demand or arrivals (e.g. Brumelle at al., 1990, Liang, 1999, van Slyke and Young, 2000). For further discussion please see Brumelle and Walczak (2003) as well as the unified treatment of the single-leg models in Talluri and van Ryzin (2004a).

The objective in all the above papers is to maximize revenues given the risk-neutrality of the decision-maker. Recently, this risk-neutrality assumption has been questioned (e.g., Feng and Xiao, 1999, Levin et al. 2006) and traditional assumptions on demand models, including availability of demand information, have been relaxed in the context of pricing (e.g., Farias and van Roy, 2006, Rusmevichientong et al., 2006, Lim and Shanthikumar, 2006). New perspectives have also emerged in the analysis of the single-leg, multi-fare booking control problem in the last decade, varying from use of customer-choice models (Talluri and van Ryzin, 2004b), to the use of adaptive methods to compute booking limits (e.g., McGill and van Ryzin, 2000), to the use of robust optimization (Birbil et al., 2006, Perakis and Roels, 2006), to methods that combine demand learning with optimization (Eren and Maglaras, 2006), to methods that use competitive analysis of online algorithms to determine policies that come with worst-case performance guarantees (Ball and Queyranne, 2006, Lan 2008).

In this new stream of research, Eren and Maglaras (2006) propose using the maximum entropy approach to update the booking limits while obtaining demand information. Perakis and Roels (2006) assume limited demand information in a manner similar to our work and provide a general approach to both the single-leg and the network RM problems where the objective is to minimize the regret. They also restrict their analysis to particular policies (e.g., nested policies for single-leg, partitioned allocations for network). Ball and Queyranne (2006), which we refer as BQ, were the first to adapt the notion of competitive analysis of online algorithms to RM context. They derive static, nested booking limits for the single leg problem and also consider bid-price controls where a booking request does not belong to a particular fare class but comes with a proposed fare. Their policies require no information on demand and come with performance guarantees (achieving maximum competitive ratio, which we define in the next section). They also show for two fare-classes how the booking limits can be updated during the booking period to improve the worst case performance. Their work is extended by Lan (2008) who considers both relative and absolute regret criteria and proposes new static and dynamic booking control policies for a multi-fare class problem when only upper/lower bounds on demand are available. The optimal policies derived in Lan (2008) reduce to BQ's in special cases (when bounds become zero and/or infinity). Lan (2008) also provides a model of overbooking to manage cancelations. Our focus in this dissertation is on models and methods assuming no cancelations or overbooking.

2.2 Research Questions and Problem Definition

As the literature review indicates, there are numerous models and methods proposed for the single-leg RM problem. These methods vary in their modeling assumptions but also on the amount of information used to derive the booking control policies. There is also a trend in using less information to obtain robust methods.

Clearly, more information enables better decision-making, as long as the information is correct and used in a model with correct assumptions. It is typical to make statistical assumptions about demand in single-leg RM models (e.g., stationarity, normally distributed). However, these assumptions may not portray reality. In that case, policies derived from models with these incorrect assumptions can lead to very poor performance.

Our goal in this chapter is to quantify the "value of information" in booking control and compare different booking policies in single-leg RM. To do this, we use simulations and extensive computational experiments to test the performance of existing models and methods, that vary in the amount of demand information. Before we introduce the models included in our study in Section 2.4, we first define the single-leg booking control problem:

A firm sells its capacity in m distinct classes that consume the same resource in a finite selling horizon. In the airline and hotel context, these classes represent different discount levels with differentiated sale conditions and restrictions. Throughout this chapter, let n denote the total capacity of the resource (seats, rooms, etc.) The fares f_i of the products are ordered $f_1 > f_2 > ... > f_m \geq 0$. The challenge in this problem is that demand in each fare class is unknown. The seller's goal is to maximize the total revenue obtained by selling n units to m fare classes. To increase his revenue, the seller can stop selling to a fare class during the booking horizon even if the capacity is not sold out. For instance, if the seller knows that business travelers book later and prefer less restricted fares, then he can limit the sales to classes preferred by leisure travelers, who book earlier and prefer discounted fares.

The question then becomes which of the fare classes should be open/closed at a particular time. It is customary to refer to the same question as "which booking requests should be accepted/rejected at a particular time?" Below is a review of the typical booking control policies and the models used to derive these types of policies. We classify the models and methods according to the amount of demand information they use. The models that require probability assumptions are reviewed in Section 2.4.1, models that use adaptive algorithms are discussed in Section 2.4.2, models based on competitive analysis are reviewed in Section 2.4.3.

2.3 Introduction to Booking Control Policies

There are different types of control mechanisms which are often dictated by a reservation system.

Booking Limits are controls that limit the amount of capacity that can be sold to any particular class at a given point in time. They are either partitioned or nested. A partitioned booking limit divides the available capacity into separate blocks that can be sold only to the designated class. When using nested booking limits, the capacity available to different classes overlaps in a hierarchical manner - with higher-ranked classes having access to all the capacity reserved for lower-ranked classes. Most reservation systems use nested rather than partitioned limits and the booking control we summarized earlier referred to standard nested booking limits. We use b_i to denote the i^{th} booking limit, i.e. the capacity reserved for classes i to m. (To read the discussion about standard nesting versus theft nesting, please refer to the Appendix of this dissertation).

Protection Levels specifies an amount of capacity to reserve (protect) for a particular class or set of classes. Again, protection levels can be nested or partitioned. A partitioned protection level is trivially equivalent to a partitioned booking limits. In the nested case, protection levels are again defined for sets of classes - ordered in a hierarchical manner according to class order. y_j is used to denote the protection level for class j. There is a simple relationship between nested booking limits and protection levels, namely, $y_k = n - b_{k+1}$, for $k = 1, ..., m-1$.

Bucket Size is the difference between adjacent booking limits. We use an additional variable x_i to define it. $x_m = b_m$ and $x_i = b_i - b_{i+1}$ for $i = 1, ..., m-1$. The relationship between bucket size, protection levels and booking limits is illustrated in Figure 2.1. The notation b, y and x are used for vectors of variables.

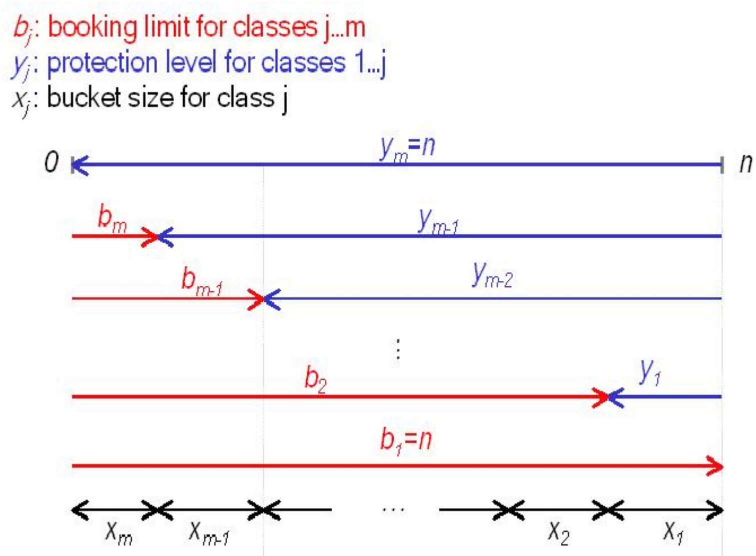

Figure 2.1: Relationship between the Different Control Variables

With a nested policy, a request in class j is accepted as long as the booking limit b_j has not been reached and capacity is available.

Note that nesting defines a static policy, which can be updated and used dynamically over time. There are also other types of control policies, which can be time-dependent, dynamic, and/or non-nested. We review the models commonly used in practice and studied in research in the next section.

2.4 Models Proposed in the Revenue Management Literature

In this section, we outline policies that have been proposed for booking control in single-leg RM. These models and methods are categorized according to the level of demand information used.

2.4.1 Static and Dynamic Models that Use Distributional Information

The models in this section use probability information on fare class demand. Let demand for class i be denoted as D_i and its distribution denoted by $F_i(\cdot)$. The earliest single-resource model is due to Littlewood (1972). The model is called **EMSR** (expected marginal seat revenue). It assumes there are only two product classes with $f_1 > f_2$ and that demand for class 2 arrives first. The optimal protection level y_1^* is given by the expression $y_1^* = F_1^{-1}(1 - f_2/f_1)$. This is a static solution. Given y_1^*, the booking limit for class 2 is $n - y_1^*$. If D_2 is the number of requests in class 2 and arrivals are LBH (i.e. class 2 requests arrive before class 1), then $\min(n - y_1^*, D_2)$ of class 2 requests are accepted. The remaining $n - \min(n - y_1^*, D_2)$ units are available for class 1 requests.

The multi-fare booking control problem can be formulated as a dynamic program, with remaining capacity n_r being the state variable. We introduce the model proposed by Brumelle and McGill (1993) here. Demand is independent across classes and over time and also independent of capacity controls. Demand for the m classes arrives in m stages, one for each class, with classes arriving in increasing order of their revenue values (i.e., LBH). At the start of stage i, the random demand D_i,D_{i-1},...,D_1 has not been realized. Let $V_i^{BM}(n_r)$ denote the value function at the start of stage i of the Brumelle and McGill model. The Bellman equation is:

$$V_{i+1}^{BM}(n_r) = E[\max_{0 \le u \le \min\{D_{i+1}, n_r\}} \{f_{i+1}u - V_i^{BM}(n_r - u)\}] \tag{2.1}$$

The optimal protection levels y_i^* for $i, i-1, \cdots 1$ by

$$y_i^* \equiv max\{n_r : f_{i+1} < \Delta V_i^{BM}(n_r)\}, \text{ for } i = 1, \cdots, m-1 \tag{2.2}$$

$$y_m^* \equiv n \tag{2.3}$$

where $\Delta V_i^{BM}(n_r)$ is defined as $V_i^{BM}(n_r) - V_i^{BM}(n_r - 1)$ and decision variable u is the quantity of demand D_{i+1} to accept. The optimal protection levels y_i^* requires computation of $\Delta V_i^{BM}(\cdot)$ for any value of remaining capacity.

Although computing optimal controls for the above model is not particularly difficult, exact optimization is not widely used in practice. Indeed, most airline RM systems use one of several heuristics to compute booking limits. One of those is called **EMSR-b** (expected marginal seat revenue-version b), proposed by Belobaba (1987,1989), which is an approximation that aggregates demand from future classes and treats them as one class with revenue equal to the weighted-average revenue. Specifically, define the aggregate future demand for classes i,$i-1, ...2$ by $S_i = \sum_{k=1}^{i} D_k$, which is a random variable, and let the weighted-average revenue from classes $1, ..., i$, be

$$\bar{f}_i = \frac{\sum_{k=1}^{i} f_k E[D_k]}{\sum_{k=1}^{i} E[D_k]} \tag{2.4}$$

Then EMSR-b protection level for class i and higher, y_i, is chosen by Littlewood's rule so that:

$$P(S_i > y_i) = \frac{f_{i+1}}{\bar{f}_i}. \tag{2.5}$$

EMSR-b policies are shown to have good performance in laboratory and industrial settings (See Talluri and Van Ryzin 2004a).

Lee and Hersh (1993) relaxed the LBH assumption. In their model, the demand for each fare class is modeled as a time-dependent stochastic process. The decision to accept or reject an incoming request takes into account the number of available capacity, the time remaining in the reservation horizon, and the fare class of the request. While the model is general, it is tractable when the arrival process is Markovian. Here is the notation for Lee and Hersh (1993) model. There are T periods and t is the time index (t=1 is the first period). In each period, at most one customer arrives. This requirement can be achieved for any system by sufficiently fine discretization of the time horizon. The probability of an arrival of class i in period t is denoted $\rho_i(t)$. Still, let n_r denote the remaining capacity and let $V_t^{LH}(n_r)$ denote the value function in period t in the Lee and Hersh model. The Bellman equation is therefore

$$V_t^{LH}(n_r) = E[\max_{u \in \{0,1\}} \{R(t)u - V_{t+1}^{LH}(n_r - u)\}], t = 1, \cdots, T. \tag{2.6}$$

where u denotes the decision to accept/reject the incoming request and $R(t) = f_j$ if a demand for class j arrives in period t and $R(t) = 0$ otherwise(i.e., if there is no arrival).

The MDP model of Lee and Hersh can be solved by backward recursion using

$$V_t^{LH}(n_r) = \sum_{i=1}^{m} \rho_i(t)(f_i - \Delta V_{t+1}^{LH}(n_r))^+, t = 1, \cdots, T. \tag{2.7}$$

staring with the boundary condition $V_{T+1}^{LH}(n_r) = 0, \forall n_r$, where $\Delta V_{t+1}^{LH}(n_r) = V_{t+1}^{LH}(n_r) - V_{t+1}^{LH}(n_r - 1)$.

2.4.2 Adaptive Algorithms that Require No Demand Information

Typically, applying the models described in the previous section require three steps. First, historical demand data are studied to determine a representative demand distribution. Second, forecasting techniques are applied to estimate the parameters of the distribution. Third, demand distribution characterization is passed to an optimization routine that solves for protection levels. However, Van Ryzin and McGill (2000) proposed a new method for directly updating booking policy parameters for the next resource usage based on observations of the performance on previous instances, without recourse to the above complex forecasting and optimization cycle. The method of Van Ryzin and McGill obtained the optimal solution to the Brumelle and McGill model: for y^* to be an optimal set of protection levels, it must satisfy:

$$P(B_i(y^*, \mathbf{D})) = \frac{f_{i+1}}{f_1}, \text{ for } i = 1, 2 \cdots, m-1. \tag{2.8}$$

where $\mathbf{D} = (D_1, D_2, .., D_m)$ is the random vector of fare requests and $B_i(y^*, \mathbf{D})$ is the *fill event*. A fill event occurs when demand in stages 1 through i exceeds the corresponding protection levels, i.e. $D_1 + D_2 + ... + D_i \geq y_i^*$. Here is a description of the adaptive algorithm:

Step 0: Choose a nested protection level vector $y^{(1)}$, set k=1.

Step 1: Observe demand vector $\mathbf{D}^{(k)}$.

Step 2: Determine fill events $B_i(y^{(k)}, \mathbf{D}^{(k)})$ for $i = 1, \cdots, m$. Compute

$$H_i(y^{(k)}, \mathbf{D}^{(k)}) = \frac{f_{i+1}}{f_1} - \mathbf{1}(B_i(y^{(k)}, \mathbf{D}^{(k)})) \tag{2.9}$$

where $\mathbf{1}$(E) is the indicator function of event E.

Step 3: Update booking limits

$$y^{(k+1)} = y^{(k)} - \gamma_k H(y^{(k)}, \mathbf{D}^{(k)}) \tag{2.10}$$

where γ_k is a step size (determined as priori).

Step 4: If $y^{(k)}$ and $y^{(k+1)}$ satisfy a convergence criteria, set $y^* = y^{(k+1)}$ and stop. Otherwise, set $k \leftarrow k+1$ and go to Step 1.

In this algorithm,γ_k is a sequence of nonnegative step sizes satisfying: $\sum_k \gamma_k = +\infty$ and $\sum_k \gamma_k^2 < +\infty$. As the algorithm progresses, if the i^{th} fill event occurs and $H_i(\cdot) = \frac{f_{i+1}}{f_1} - 1 < 0$, the protection level $y_i^{(k)}$ is increased by $\gamma_k(1 - \frac{f_{i+1}}{f_1})$. Otherwise, if $H_i(\cdot) = \frac{f_{i+1}}{f_1} > 0$, then y_i^k is reduced by $\gamma_k \frac{f_{i+1}}{f_1}$. Thus, protection levels are stepped up when high demand is observed and stepped down when low demand is observed. This algorithm is guaranteed to converge to the optimal solution with some mild assumptions. See Van Ryzin and McGill (2000) for details.

Van Ryzin and McGill's adaptive algorithm relies on no demand information and the nested booking limits computed using their adaptive procedure converge to the optimal. However this method requires LBH arrivals, learns the demand from flight to flight but not during the booking horizon of a particular flight. An important feature of this method is that it can easily be used in a simulation-based environment to compute y^* when the underlying demand distribution is available. Unfortunately, this method suffers from poor performance during the transient period when the booking limit is converging to the optimal. Besides, the choice of step-size and the initial vector of booking limits can affect the speed of convergence. Instead of using a DP method, we use Van Ryzin and McGill's algorithm in our experiments.

2.4.3 Robust Solutions Based on Competitive Analysis

In Ball and Queyranne (2006) and Lan's dissertation (2008), the classic single-leg RM problem is considered from the perspective of competitive analysis of online algorithms. (See Albers, 2003 for more information on competitive analysis of online algorithm.) This perspective evaluates the performance of a booking control policy relative to the performance of an offline algorithm that has perfect hindsight information. An offline optimal solution is a solution obtained by an offline algorithm (with hindsight) that optimizes the objective function of interest. In their competitive analysis, the competitive ratio (CR) is used as a measure of an algorithm's effectiveness. There is another performance metric of interest: absolute regret, which is the difference between the objective function values of the offline and online algorithms. (Please refer to BQ's paper and Lan's dissertation for more detailed discussion.)

CR is defined as the *minimum* of the ratio of revenues obtained by the online algorithm to the offline revenues. If we let Ω_Υ be the set of all possible input sequences to an online algorithm Υ and, for any $I \in \Omega_\Upsilon$, let $R(I;\Upsilon)$ be the objective value achieved by the online algorithm for input I and let $R^*(I)$ be the objective value achieved by an optimal offline algorithm, then, we can define CR as:

$$\text{CR of } \Upsilon \quad = \quad \inf_{I \in \Omega_\Upsilon} \frac{R(I;\Upsilon)}{R^*(I)}.$$

Using no information and considering a worst-case approach, BQ obtained the following static nested policy:

$$b^*_{i+1} \quad = \quad n - y^*_i \text{ for } i = 1, \cdots, m-1, b_1 = n. \tag{2.11}$$

where

$$y^*_i = \frac{n}{\theta}\left(i - \sum_{j=1}^{i} \frac{f_{j+1}}{f_j}\right), \text{ for } i = 1, 2, \cdots, m-1 \tag{2.12}$$

$$y^*_m \equiv n \tag{2.13}$$

where $\theta = m - \sum_{i=2}^{m} \frac{f_i}{f_{i-1}}$.

Lan (2008) extends this assuming demand in fare class i is no less than L_i and no more than U_i, for $i = 1, 2, \cdots, m$. They show that a Linear Program (LP) can be formulated to determine a nested policy that maximizes CR. We refer to this model as the GBM model.

The nested booking limits are defined by

$$b^{CR}_j \quad = \quad \sum_{i=j}^{m} x^{CR}_i \text{ for } j = 1, \cdots, m \tag{2.14}$$

where

$$\bar{z}^{CR} = \frac{R_u^+/f_u + N_u}{R^*(CAST_u)/f_u + \sum_{i=1}^{u-1} g_i} \tag{2.15}$$

$$x_j^{CR} = \begin{cases} g_j\bar{z}^{CR} + L_j & j < u \\ (1/f_u)(R^*(CAST_u)\bar{z}^{CR} - \sum_{i=1}^{u-1} f_i L_i) & j = u \\ 0 & j > u \end{cases} \tag{2.16}$$

$$u = \max\{j \leq m : (\sum_{i=1}^{j-1} f_i L_i)(\sum_{i=1}^{j-1} g_i) < R^*(CAST_j)(n - \sum_{i=1}^{j-1} L_i)\}. \tag{2.17}$$

The index u denotes the critical fare-class such that all classes $k > u$ are closed, and g_i is an auxiliary parameter defined as

$$g_i = \frac{R^*(CAST_i) - R^*(CAST_{i+1})}{f_i}, \quad i = 1, \cdots, m-1. \tag{2.18}$$

The $CAST_i$ refers to a non-dominated scenario such that any input scenarios I satisfy

$$\frac{R(I : b)}{R^*(I)} \geq \min \frac{R(CAST_i : b)}{R^*(CAST_i)}, i = 1, \cdots, m \tag{2.19}$$

given a nested booking limit policy b. Lan(2008) characterizes these non-dominated scenarios in his analysis and proves that there are only m scenarios of interest.

Lan (2008) also provides a dynamic policy where the booking limits can be updated over time. Consider a dynamic scenario where a dynamic policy has been executing to process the booking requests. Suppose h_i bookings have been accepted so far for fare class i for i=1,...,m. This accumulates revenue of $\sum_{i=1}^{m} h_i f_i$ from the $\sum_{i=1}^{m} h_i$ sold seats. And define $\hat{n} = n - \sum_{i=1}^{m} h_i$ as the remaining number of available seats. The question Lan (2008) asked is whether the booking limits can be adjusted to improve the CR. He shows that if demand and capacity information are updated as follows:

$$\hat{L}_j = \max(L_j, h_j) - h_j, \quad \hat{U}_j = U_j - h_j, \quad j = 1, ..., m,$$

then a new set of booking limits for allocating the remaining $\hat{n}$ seats can be computed to guarantee a better performance. The new policy is easily obtained by resolving the original model with the updated parameters

When $L_i = 0$ and $U_i \geq n$, $\forall i$ in Lan (2008), the booking limits in equation 2.14 is the same as in equation 2.11, which is the solution suggested by BQ. CR is a measure of relative regret. One can also consider "absolute regret" (AR). The maximum absolute regret of the online algorithm is defined as: MAR of $\gamma = sup_{I \in \Omega_r} |R^*(I) - R(I;\gamma)|$. Analysis and solution of AR problem is similar to CR. Lan (2008) also discusses an optimal policy that maximize AR can be obtained in closed form.

2.5 Numerical Experiments

We use simulation to evaluate the performance of methods introduced thus far. These methods are listed in Tables 2.1, 2.2, and 2.3.

We compute the performance of each policy relative to an "ideal solution". At the end of each simulation run, we compute the ratio of policy revenues to that of the offline optimal, which is the perfect hindsight solution for that simulation run. At the end of an experiment, we calculate the average of this performance ratio. In our experiments, we vary the fare values, demand parameters, arrival regimes with

Method	Underlying Model	Policy Structure	Introduced in
DP (or VRM)	DP model of Brumelle and McGill (1993), computed by Van Ryzin and McGill (2000) adaptive algorithm	nested BLs	Section 2.3 Section 2.4
STAT-CR	CR model of BQ and Lan(2008)	static nested BLs	Equation 2.11 and 2.12
STAT-AR	AR model of Lan(2008)	static nested BLs	Section 2.4.3
DYN-CR	CR model of Lan(2008)	nested BLs with dyanmic updates	Section 2.4.3
DYN-AR	AR model of Lan(2008)	nested BLs with dynamic updates	Section 2.4.3
FCFS	First-come First-serve	all accepted up to capacity	-

Table 2.1: Methods that require no information

Method	Underlying Model	Policy Structure	Introduced in
EMSR	stochastic model of Littlewood (1972)	static nested BLs	Section 2.3
EMSR-b	heuristic of Belobaba (1989)	static nested BLs	Section 2.3
MDP	stochastic model of Lee and Hersh (1993)	dynamic, time and inventory dependent	Section 2.3

Table 2.2: Probabilistic models

Method	Underlying Model	Policy Structure	Introduced in
BSTAT-CR	CR model of Lan(2008)	static nested BLs	Section 2.4.3 Equation 2.14
BSTAT-AR	AR model of Lan(2008)	static nested BLs	Section 2.4.3
BDYN-CR	CR model of Lan(2008)	nested BLs with dynamic updates	Section 2.4.3
BDYN-AR	AR model of Lan(2008)	nested BLs with dynamic updates	Section 2.4.3

Table 2.3: Distribution-free methods that use aggregate information

respect to fare classes (LBH vs. time-homogeneous arrivals), and the demand-mix (mean demand of a fare class relative to the demand of other classes). When arrivals are LBH, the sequence of arrivals is known, and the total number of arrivals in each fare class is computed by sampling from the demand distribution in each simulation run. When the arrivals are time-homogeneous, the total number of arrivals in each fare class is determined by sampling from the corresponding (aggregate) demand distribution. Then, the arrival times of requests within each fare class are randomly generated from a Uniform(0,1) distribution so that the arrivals in each fare class are requested during the entire booking horizon.

In all the experiments, the demands across each fare-class are independent. We denote the mean demand of fare class i as λ_i. Requests arrive one-by-one in all the experiments. The capacity is $n = 100$ in all the examples unless noted otherwise.

2.5.1 Models That Require No Demand Information

Example-2.1. In this two-fare example, we compare unbounded robust methods using CR criteria with VRM. The averages over 300 runs are reported in terms of the first 60 runs and next 240 runs for the LBH case, and for the total of 300 runs for the time-homogenous case. The fares are $f_1 = 200$, $f_2 = 100$. Demands are independent, stationary and normally distributed for both fare classes. Mean demands $[\lambda_2, \lambda_1]$ (we refer to this as the demand-mix or demand factor) range from [120,0] to [0,120].

Knowing the demand distribution and/or the arrival process, the EMSR policy (the DP solution) in the LBH case and the MDP policy in the time-homogeneous case are the online optimal algorithms, performing remarkably well compared to the offline. They are benchmarks to evaluate the performance of robust and VRM policies. When arrivals are low before high, the protection level under the VRM policy should converge to the mean demand of the high-fare class. This is observed in Figure 2.3 as VRM adapts to the online optimal after a certain warm-up period (warm-up period is illustrated in Figure 2.2). When arrivals are time-homogeneous, Figure 2.4 shows that, being a heuristic, VRM still gives relatively good performance. The key theoretical and practical problem associated with VRM is determining conditions under which such an adaptive algorithm will converge to optimal protection levels. The sequence needed should take large steps early (to speed up the warm-up period) and become smaller as the algorithm progress. In this experiment, we use a sequence of the form A/(n+B), where A=200 and B=10, chosen to effect larger

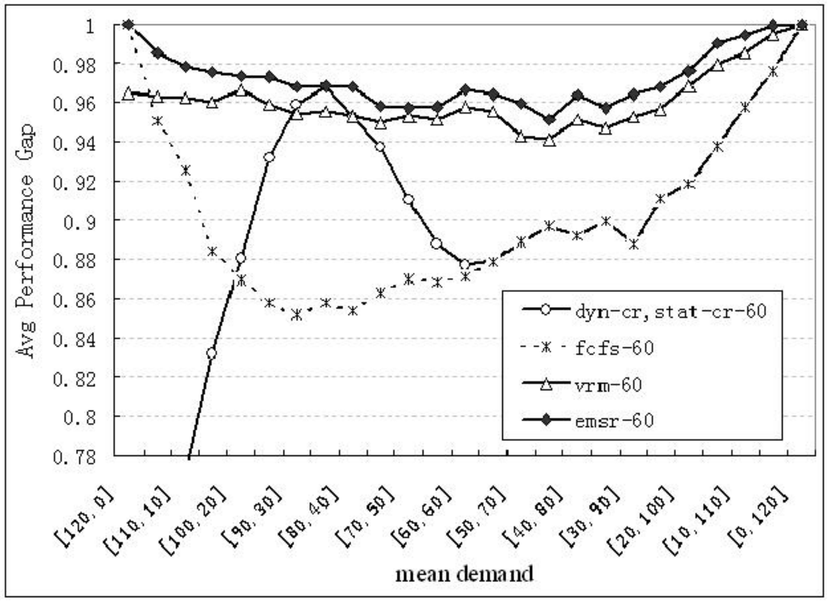

Figure 2.2: Average performance in Example 2.1 with LBH arrivals-first 60 days

early steps, which appeared to provide good performance on a range of examples.

The performance of robust policies first increases as demand shifts from class 2 to class 1 (because the protection levels are too high at first), followed with a slight decrease (where the protection levels are not high enough) and later a slight increase (when λ_2 is negligible). When arrivals are LBH, robust policies exhibit a similar behavior whereas FCFS does relatively poorly. While the best performance of robust policies is above 95%, the policies reach this peak performance with a lag. This is a direct effect of differing protection levels of the policies. For an LBH input regime, the dynamic policy DYN and static policy STAT are identical when there are two fares.

Example-2.2. In this experiment, we compare the robust policies with CR and AR criteria. The fares are $f_1 = 500$ and $f_2 = 100$. Demand for each fare-class is

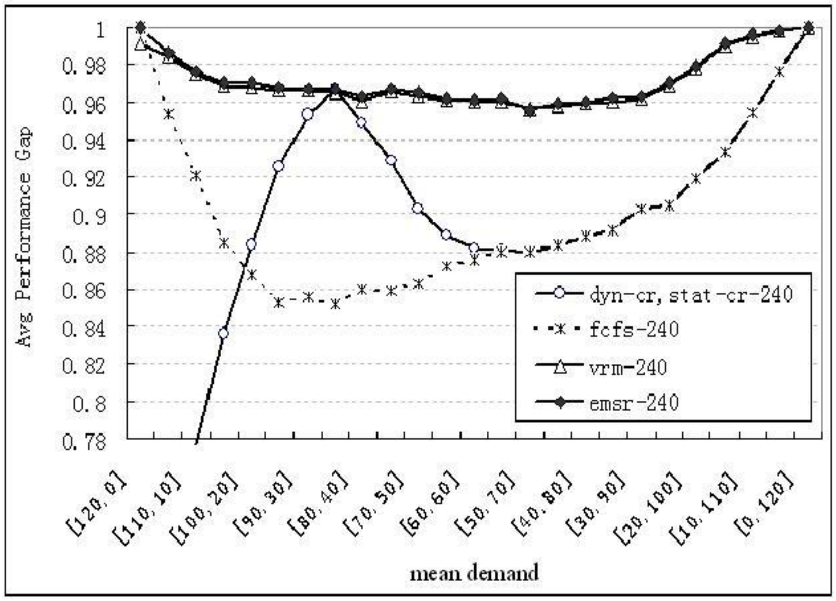

Figure 2.3: Average performance in Example 2.1 with LBH arrivals-next 240 days

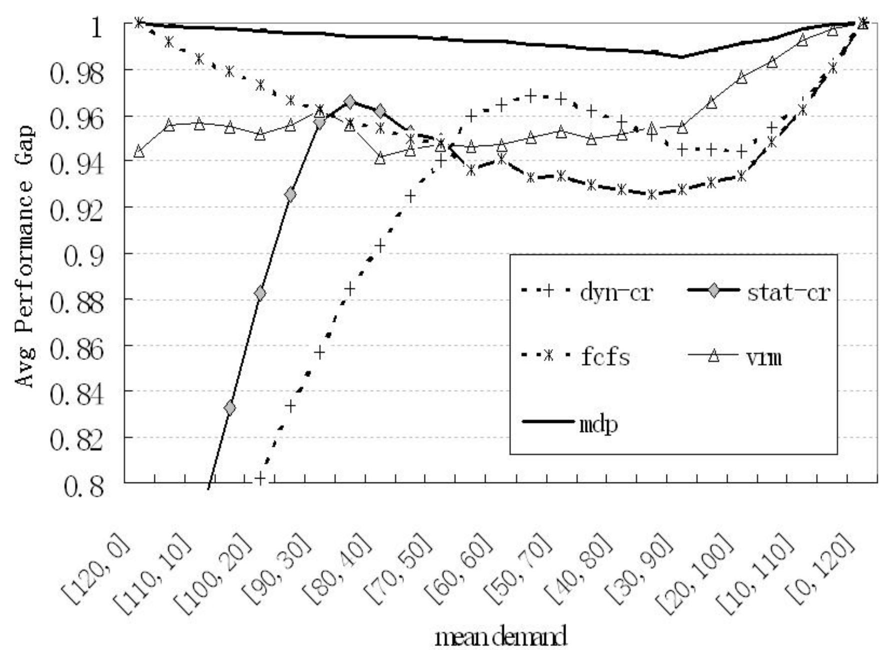

Figure 2.4: Average performance in Example-2.1 with time-homogeneous arrivals

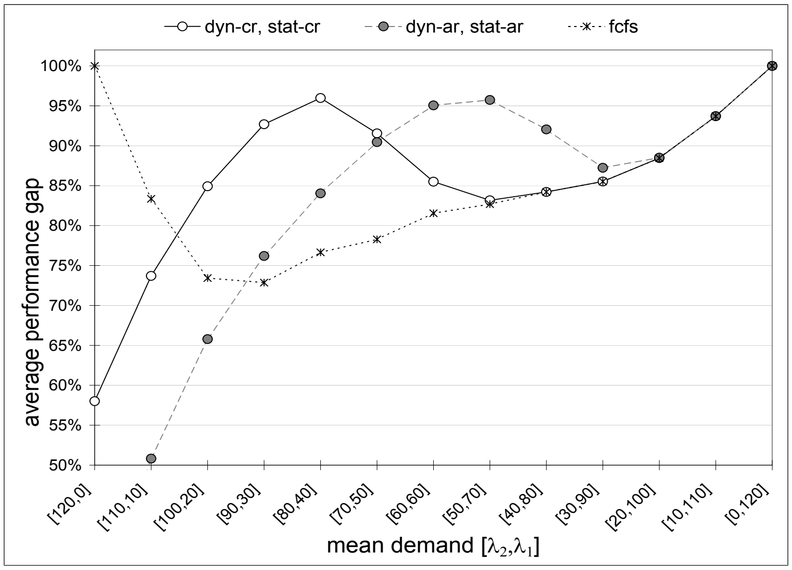

Figure 2.5: Average performance in Example-2.2 with LBH arrivals

independent and time-homogenous (Poisson distributed) with means ranging from [120,0] to [0,120]. In this example, static policies protect $y_1^{AR} = 80$ and $y_1^{CR} = 44.4$ seats for class 1. The theoretical worst-case ratio of FCFS, STAT-AR and STAT-CR revenues to offline revenues are 20%, 20% and 55.5%, respectively. However, the average performance can be significantly better as shown in Figures 2.5 and 2.6. Figure 2.6 gives the average performance gap when the arrivals are time-homogeneous. In Figure 2.7, the average number of unsold seats is reported.

The average performance of FCFS is over 90% in this case because r is not too low and arrivals are time-homogeneous. The unbounded robust policies do very poorly when λ_1 is low because they protect too many seats for the high-fare class, this is also evident in Figure 2.7 where the average number of unsold seats is reported. STAT-CR reaches the peak performance at [80,40] whereas STAT-AR reaches it at [40,80]. AR polices are more aggressive in protecting seats for the high fare class

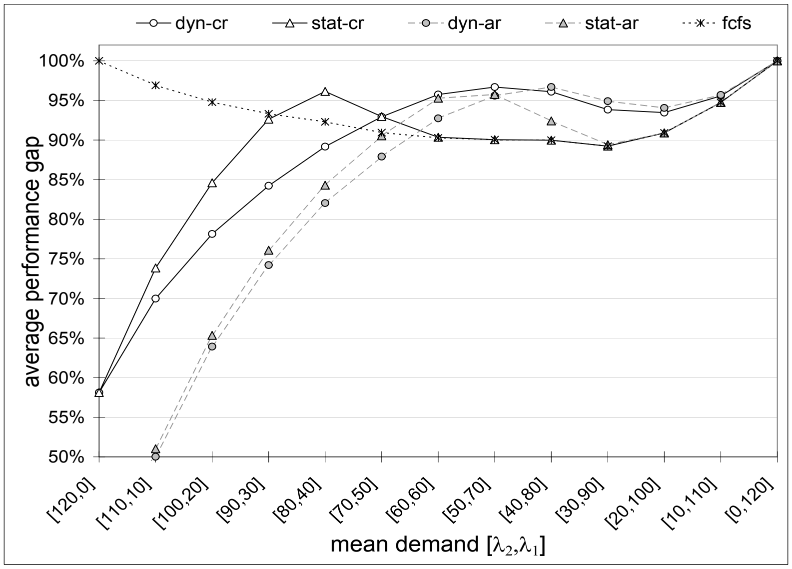

Figure 2.6: Average performance in Example-2.2 with time-homogenous arrivals

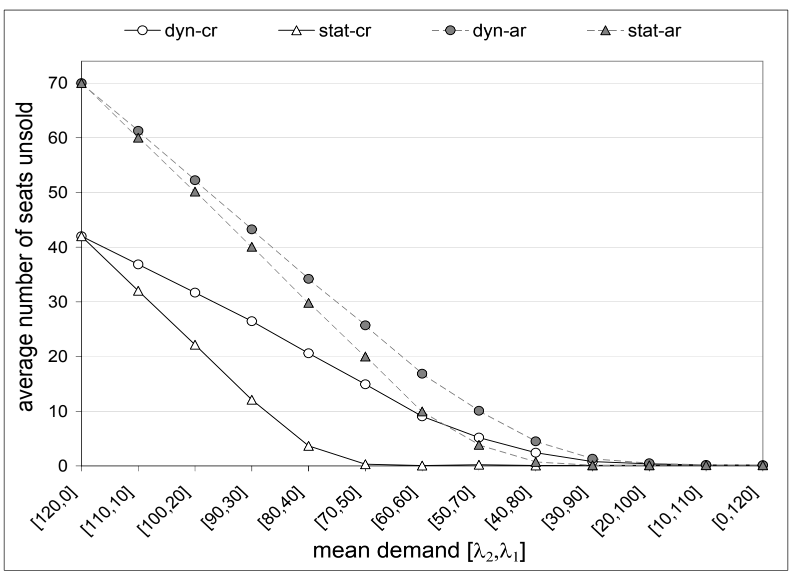

Figure 2.7: Number of Unsold Seats in Example-2.2 with time-homogeneous arrivals

($y_1^{AR} > y_1^{CR}$). Consequently, its revenues can be lower than STAT-CR when the demand for class 1 is low. As r $\longrightarrow$ 0, only 50% of the seats are protected by STAT-CR whereas 100% are protected by STAT-AR. In contrast, the FCFS policy that accepts all the incoming requests until capacity is reached protects 0 seats for class 1. It can easily be seen that as r $\longrightarrow$ 1, these three policies are equivalent.

We repeated this experiment with different values of r (Please refer to Example-A.3 and Example-A.4 in the Appendix) As expected, as r increases, the difference between policies obtained by CR and AR gets smaller. Here is a summary of our observations: Overall, we looked at the effect of expected demand and demand-mix, discount ratio and the arrival regime (time-homogenous vs. LBH arrivals) to gain insights on the behavior of online policies. We considered both practical and extreme cases. In some instances, robust policies perform very well. Yet, there is still room for improvement overall. For instance, when demand mix lies in the practical range of [90,30] to [60,60], the lowest average performance gap of the CR-based and AR-based policies are 85% and 75%, respectively. The robust policies do most poorly when λ_1 is low. One can argue that these algorithms are doing the job they were designed to do in the sense that they are protecting against the possibility of a high value of class-1 mean demand λ_1.

In summary, in this section, models that need no demand information are compared side by side. The protection level under the VRM policy converges to the mean-demand of the high-fare class and its performance adapts to the online optimal after a certain warm-up period. However determining conditions under which such an adaptive algorithm will converge to the optimal remains a challenge. The sequence of step sizes should be carefully chosen based on experimentation in order to speed up the warm-up period and quickly stabilize as the algorithm progresses. While the best performance of robust policies is above 95%, they do most poorly when there are fewer high-fare customers since they are designed to protect against possibility of later high-fare requests. AR policies are more aggressive than CR policies in protecting seats for high fare class. However, if information were available indicating high values of λ_1 are not possible, then improved performance should be possible in general, and in the extreme cases, in particular. These observations motivate models using limited demand information; discussion and analysis follow in the next section.

2.5.2 Models That Use Accurate Demand Information

In this section, we will show the performance of robust policies that use demand bounds: BSTAT-CR, BSTAT-AR, BDYN-CR and BDYN-AR. We first use the example with uniform demand to show how information improves the performance of online policies. Then we choose the same example used in the previous section to compare these robust methods (use limited demand information) with other well-known policies that require more information.

Example-2.3 We continue to employ fares $(f_1, f_2) = (500, 100)$, and assume demand in each fare class follows a discrete uniform distribution between $L_i = 40$ and $U_i = 80$, $i = 1, 2$. We use 6000 simulation runs with LBH arrivals in this example. The protection level of class 1 for each of the policies, the theoretical CR of each policy computed given the demand bounds, average policy revenues, relative performance, and the average number of seats sold are displayed in Table 2.4. Note that the differences between the protection levels of CR and AR policies are significant when no demand information is used. As expected, the worst-case CR is not an indication of the average performance of the policies. In this example, STAT-CR achieves significantly lower average revenues compared to other policies because of the protection level and the fare-ratio: STAT-CR achieves a higher load (as indicated by the average number of seats sold) compared to other policies, but other policies accept more of class 1 requests leading to higher average revenues. We monitored the performance of the policies more closely by studying the distribution of revenues (i.e., we computed estimates of percentiles). We split the 6000 simulation runs into 30 samples of size 200 each, we computed the percentiles in each sample, and we took the averages of the percentiles of the samples to obtain the estimate. This information is provided in Figure 2.8. Notice that the ranking of the policies with respect to the 10^{th}, 50^{th} and 90^{th} percentiles are different. In fact, the ranking of our policies in the 10^{th} percentile is reversed in the 90^{th}. Hence, no policy (OFFLINE and FCFS excluded) *stochastically dominates* the others. STAT-CR has the highest 10^{th} percentile value, hence the lowest downside risk. Use of demand information degrades the performance of CR policies in terms of the downside risk, but provides significant gains on the upside, i.e., BSTAT-CR has a significantly higher 90^{th} percentile value compared to STAT-CR. Likewise, STAT-AR has a higher downside risk and a lower upside risk compared to BSTAT-AR.

This example shows the main difference between CR and AR policies and also

Policy	*Protection level* y_1	*Theoretical CR (%)*	*Average revenues*	*Avg of ratio of revenue to OFFLINE (%)*	*Avg no.of seats sold*
BSTAT-BDYN-CR	68.49	89.04	129,235	95.37	89.3
BSTAT-BDYN-AR	72	87.69	129,501	95.28	86.93
STAT-DYN-CR	44.5	66.19	114,302	85.84	98.79
STAT-DYN-AR	80	84.61	128,013	93.82	79.58
FCFS	0	42.86	101,864	76.63	98.99
OFFLINE	$--$	100	135,663	100	98.99

Table 2.4: Protection level,Theoretical CR, Average performance in Example-2.3.

the effect of demand information. Because AR policies are aggressive in protecting seats for class 1, they have a lower upside risk (higher downside risk) compared to CR policies. The use of demand information increases the average revenues of both CR and MAR policies. The differences between CR and MAR policies are smaller when demand information is used. Demand information also affects the upside and downside risks: The downside risk is higher (lower) and the upside risk is lower (higher) for BSTAT-CR (BSTAT-AR) compared to STAT-CR (STAT-AR).

We repeated Example 2.3 with different demand parameters and arrival regimes. The results are reported in the Appendix (See Example-A.5 through Example-A.9 in the Appendix). When demand is distributed uniformly between 20 and 60 for class 1, and between 60 and 100 for class 2, and arrivals are LBH, no policy (except OFFLINE) is stochastically dominant. This particular demand-mix affects STAT-AR most: It has a significantly lower revenue at the 10^{th} percentile, and provides the second highest revenue value at the 90^{th} percentile, following BSTAT-AR. This is expected because STAT-AR protects the highest number of seats for class 1, for

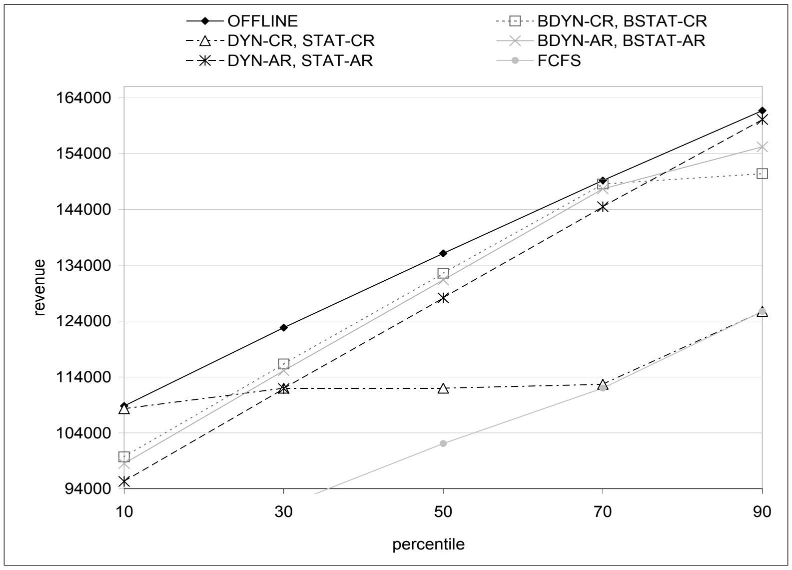

Figure 2.8: Average performance in Example-2.3

which the demand is lower in this case. The observations regarding the relative performances of the other policies remain the same. When demand is distributed uniformly between 20 and 60 for class 2, and between 60 and 100 for class 1, STAT-CR is stochastically dominated because its booking limit for class 2 is significantly higher than the other policies. When the arrivals occur homogeneously over time in both fare classes, dynamic and static policies are not equivalent. If demand-mix is balanced as in Example-1a, then the average revenues obtained by each of the policies except DYN-AR are higher when arrivals are time-homogeneous; DYN-AR sets a high protection level for class 1 initially and this protection level is updated with each class 1 request leading to rejecting far too many class 2 requests and having too many idle seats at the end of the booking horizon. However, DYN-AR is not stochastically dominated, i.e., has a higher 90^{th} percentile than, for e.g., BSTAT-AR. Our previous observations regarding stochastic dominance relations among the

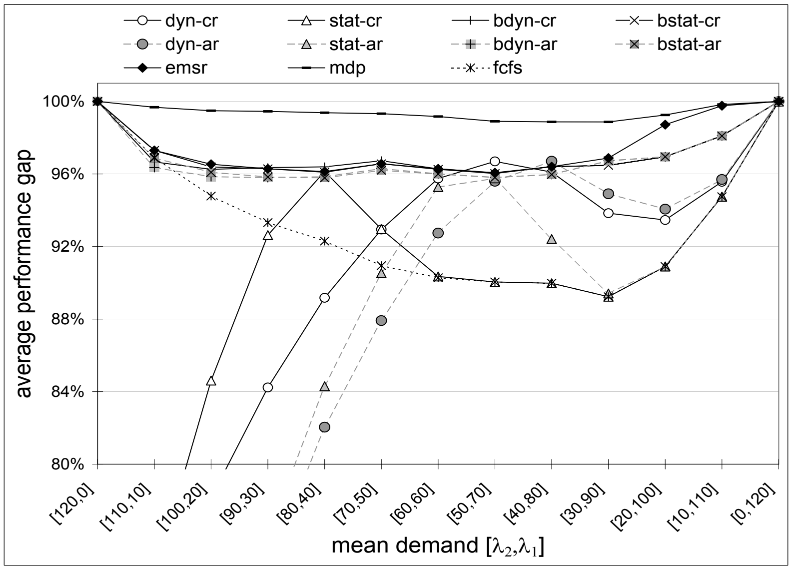

Figure 2.9: Average performance in Example-2.4, Time-homogeneous arrivals

policies do not change with the arrival regime.

Example-2.4 Same parameter setting as in Example 2.2. In addition, we choose lower and upper bounds of demand to be two standard deviations away from the true mean. When arrivals are time-homogeneous (see Figure 2.9), MDP is the optimal policy and it performs remarkably well compared to offline optimal. EMSR is only a heuristic in this case. The use of demand bounds improves the performance significantly and robust policies that use demand bounds do as well as EMSR except in the extreme case where class 2 demand is negligible and class 1 demand exceeds capacity (which is impractical from a RM perspective). When the arrivals follow the LBH regime (see Figure 2.10), EMSR is the optimal policy. Note that the robust policies with demand information are again indistinguishable from each other and also from EMSR for practical instances.

We repeated this example for different fare values. The main observations

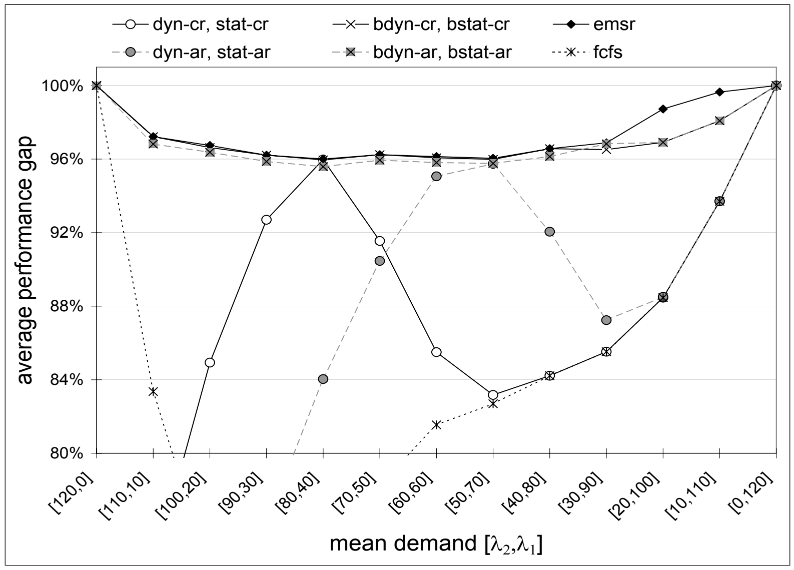

Figure 2.10: Average performance in Example-2.4, LBH arrivals

remain the same: MDP (as the ideal solution) is close to offline optimal when arrivals are time-homogeneous. Robust policies with demand information are as good as EMSR at all discount ratios regardless of the arrival regime (except when class 2 demand is negligible). The use of true bounds have a significant effect on the performance of robust policies and the average performance gap is no worse than 95% in any of the experiments. In these experiments, the additional benefit of using a dynamic robust policy is almost negligible when demand information is used. This is because the demand is stationary.

Next, we look at an example with more than two fare-classes.

Example-2.5 Here, we have m = 4. This is adapted from the example used in Talluri and van Ryzin (2004a), Section 2.2.3.4. The fares are f = (1050, 567, 527, 350), n = 124, the demand is Normal distributed and independent across fare-classes. The arrivals are LBH and mean demand is 17.3,45.1,73.6 and 19.8 for classes

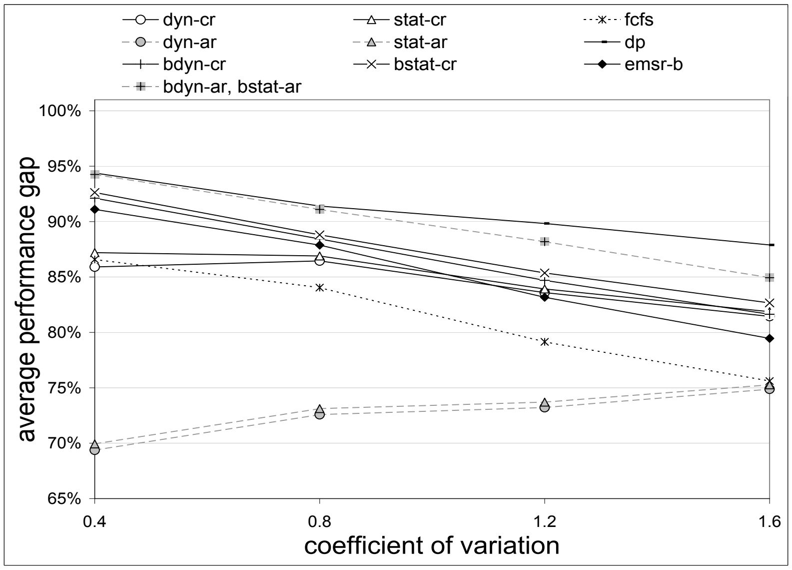

Figure 2.11: Average performance in Example-2.5

1 through 4. All fare-classes have the same coefficient of variation (CoV) and we vary the CoV in this experiment. Demand bounds are set to two standard deviations from the mean. The average performance gap is reported in Figure 2.11. The optimal policy in this case is given by DP and is significantly better than the other policies especially when the CoV is high. BSTAT-AR and BDYN-AR are indistinguishable in this experiment. They are almost as good as DP for low CoV. All robust policies with demand bounds dominate EMSR-b; performance of EMSR-b degrades more as the CoV increases. The benefit of demand information is significant on AR-based robust policies: DYN-AR and STAT-AR do very poorly.

The observations from the above set of experiments can be summarized as follows: (1) CR and static policies are better in terms of the 10th percentile of the revenue distribution, except when class 2 demand is low: These policies are more conservative and tend to protect fewer seats for class 1 compared to other policies.

(2) AR and dynamic policies are better in terms of the 90th percentile of revenues, i.e. the chance of achieving higher revenues is higher with these policies, except when class 1 demand is low. (3) Time-homogeneous arrivals are in general better for all the policies, but can hurt the performance of dynamic AR policies when class 1 demand is low. (4) The use of correct demand information decreases the variance in revenues of AR and CR policies, increases the average performance, and makes the relative performance of these policies less sensitive to changes in the fare ratio. (5) The difference between AR and CR policies is negligible when correct demand information is used and class 1 demand is low.

The results of the above experiments are very encouraging: When demand information is accurate, robust methods are practically as good as EMSR and even better than EMSR-b which is commonly used in airline RM practice. Examples 2.4 and 2.5 show the best results possible for robust policies because the bounds are computed with the knowledge of the underlying demand distribution. However, the main motivation for robust methods is lack of data and accurate forecasts. In real life, not only the demand forecasts will be wrong, but also some of the assumptions in models such as EMSR, DP or MDP may fail to hold. The next section show how these methods perform when booking limits of the policies are computed with incorrect demand models and/or data.

2.5.3 Effects of Inaccurate Information

In this section we look at the effect of quality of demand information on booking control policies.

Example-2.6 The arrivals are LBH and the true demand in each fare class is uniformly distributed between 40 and 80. In addition to the true range of the demand, we use robust policies with a *narrow* range of demand where the upper and lower bounds of demand are estimated to be 55 and 65, also with a *wide* range where the upper and lower bounds of demand are estimated to be 20 and 100. This setup results in three versions of each of our policies, and we denote them as T (for true bounds), N (for narrow range) and W (for wide range). The average relative performance of the policies are reported for different fare ratios in Figure 2.12. Note that the results with true bounds are only a repeat of what we had in Example-3, and AR and CR policies have almost exactly the same relative performance regardless of the fare ratio. This is still true when the range is narrow. However, the effect of wide range is different on AR vs. CR policies. When the range is wide and fare ratio is low, AR policies do worst because they expect higher class 1 demand and set higher protection levels. The effect of such aggressive protection levels on the policy revenues is low when fare ratio is high. The differences among the policies stem from the differences in protection levels. To highlight that, we computed the protection levels for all the policies at three different fare ratios. The results are presented in Table 2.5. One interesting observation in this example is the following: The policies do not perform perfectly when $r = 1$ if a narrow range is used. This is very

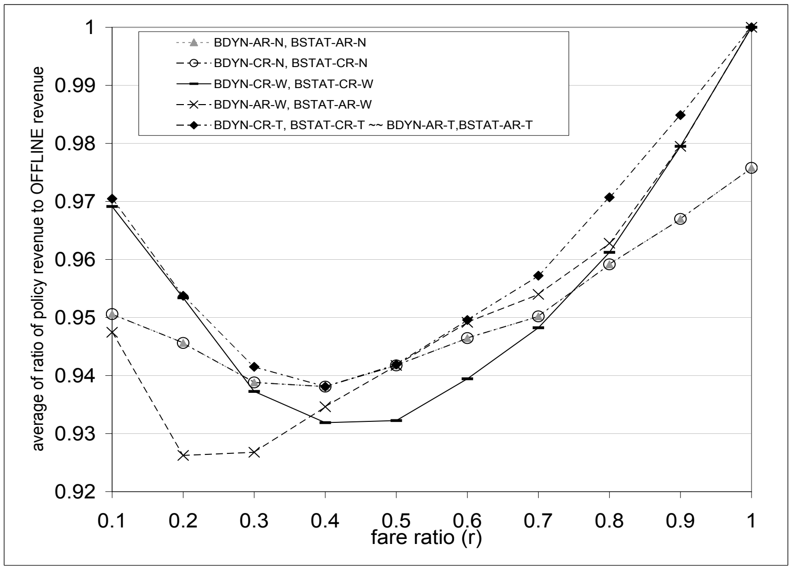

Figure 2.12: Performance of the policies with inaccurate demand information in Example-2.6

surprising given the fact that all requests have the same fare, and any reasonable policy (including FCFS) does as good as OFFLINE when $r = 1$. The main reason for this is that both AR and CR policies protect at least 55 seats for class 1 when the range is narrow, because that is the estimated lower bound on class 1 demand. Hence they reject too many class 2 requests.

r	$BSTAT^T_{CR}$, $BDYN^T_{CR}$	$BSTAT^N_{CR}$, $BDYN^N_{CR}$	$BSTAT^W_{CR}$, $BDYN^W_{CR}$	$BSTAT^T_{AR}$, $BDYN^T_{AR}$	$BSTAT^N_{AR}$, $BDYN^N_{AR}$	$BSTAT^W_{AR}$, $BDYN^W_{AR}$
0.2	68.5	62.81	67.22	72	63	84
0.5	57.5	59.85	50	60	60	60
0.9	43.85	55.99	27.42	44	56	28

Table 2.5: The protection level of class 1 for each policy in Example-2.6

We repeated this example with a different demand-mix (mean demand of classes 1 and 2 were 40 and 80). The main observations remained the same. Narrow bounds make AR and CR policies almost indistinguishable. While a wide range affects the performance of AR policy when the fare ratio is low, and a narrow range leads to conservative solutions when the fare ratio is high, robust policies remain very robust, i.e., with demand information achieve at least 92% of the OFFLINE revenues on the average in all the experiments.

Note that our last experiment only studied the effect of demand information when the demand distribution had the correct mean, but had a range that was either too narrow or too wide. The performance of robust policies can be very poor if the range of demand is chosen arbitrarily, e.g., when there is no overlap between the estimated range and the true range. However, such instances are not very realistic and/or are problematic for any method that relies on demand information.

Example-2.7 In this example, $m = 2$, demand is Poisson distributed, arrivals are time-homogeneous, and parameters of all control policies are computed assuming a demand-mix of [40,40] for the booking horizon. However, there is an unexpected surge and the demand doubles in the second half of the booking horizon, bringing the demand-mix to [60,60]. This would be the case on certain routes or at certain locations, e.g., during "March Madness" when the final four teams in the NCAA basketball championship are determined. The average relative performance of the policies as a function of the fare ratio r are given in Figure 2.13. EMSR performs significantly worse when r is small because the optimal protection level chosen by Littlewood's rule decreases with r, and EMSR is not able to benefit from the unexpected surge in the demand of class 1. There is not a significant difference among

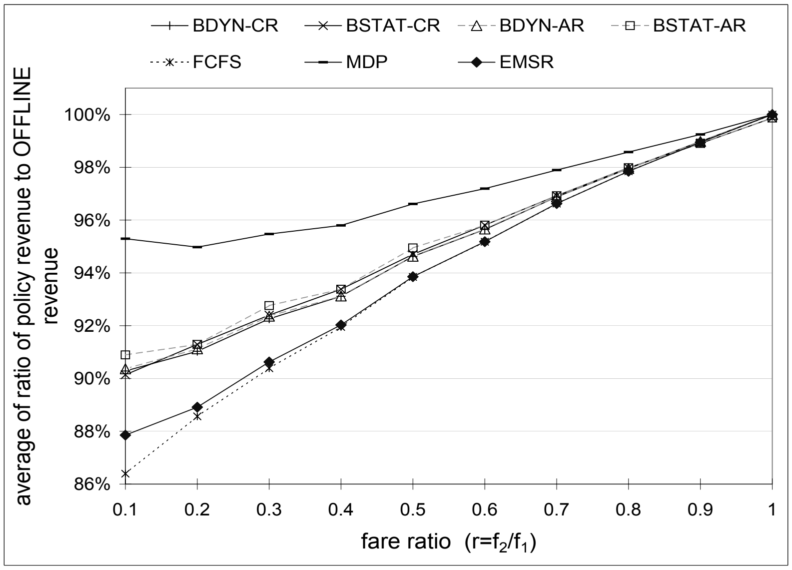

Figure 2.13: Average performance of policies in Example-2.7

our policies; they all dominate EMSR and are dominated by MDP which is able to adjust booking limits based on remaining time and capacity. Note that the policies that are most aggressive in protecting seats for class 1 (e.g., STAT-AR, DYN-AR) and can dynamically update the booking limits (e.g., DYN-AR) would do well in this experiment.(See Example-A.10 in Appendix for the performance of our policies with no demand information)

Next, we look at what happens when the demand distribution is misspecified.

Example-2.8 Here, $m = 2$ and the demand is again Poisson distributed. The true demand for each fare class is 60 for the entire booking horizon. While the true demand distribution of class 2 is known, only an estimate of the mean demand for class 1 is available. All policy parameters are determined based on this estimate and demand bounds are set to two standard deviations away from the estimated mean. The fares are $f_1 = 500$, $f_2 = 100$. The average performance

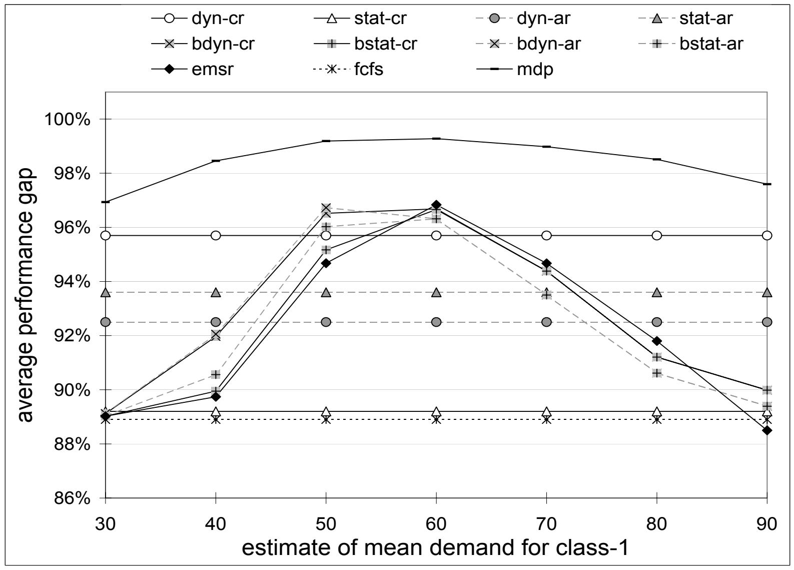

Figure 2.14: Average performance in Example-2.8 with time-homogeneous arrivals

gap for time-homogeneous and LBH arrivals are presented in Figure 2.15, where the estimate of mean demand varies from 30 to 90. Notice that robust policies with demand bounds behave similar to EMSR. They do relatively well given an acceptable forecast error (e.g. in the range 45 to 75). LBH arrivals amplify the effect of underestimating high-fare demand. Robust policies with demand information do slightly better than EMSR when class 1 demand is underestimated or highly overestimated. This is natural because (i) these policies tend to overprotect seats so adverse affect of underestimating the demand for class 1 is less, and (ii) the demand information is only valuable if the upper bound is no more than the capacity (i.e. overestimation becomes less detrimental to the performance beyond a certain point).

While robust policies are sensitive to bound information in this example, this is because the bounds used are relatively tight. The looser the bounds, the closer the performance of policies with bounds to the ones with no demand information.

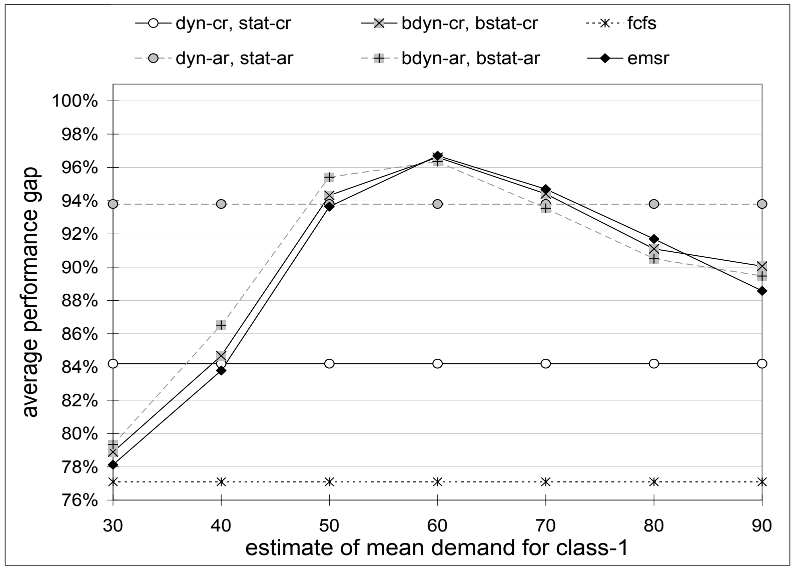

Figure 2.15: Average performance in Example-2.8 with LBH arrivals

The last two experiments confirm that our policies with limited demand information perform as good as EMSR under various scenarios. While MDP is quite robust to the error in parameter estimation in this example, the next example shows what happens when there are other misspecified parameters in the arrival process.

Example-2.9 The demand is Poisson distributed and demand-mix varies from [120,0] to [0,120]. Other parameter setting is similar to Example 2.3. The mean demand for the entire booking horizon is known. In this experiment, MDP assumes stationary arrivals for the entire booking horizon, whereas all the demand arrives in a rush only in the first half of the booking horizon. As seen in Figure 2.16, MDP has poor performance when λ_2 is high and λ_1 is low because it ends up rejecting too many class 2 requests. In contrast, robust policies and EMSR are not affected by modeling/estimation errors regarding inter-arrival times. While this experiment may be an extreme one, it illustrates that a control policy obtained via MDP is

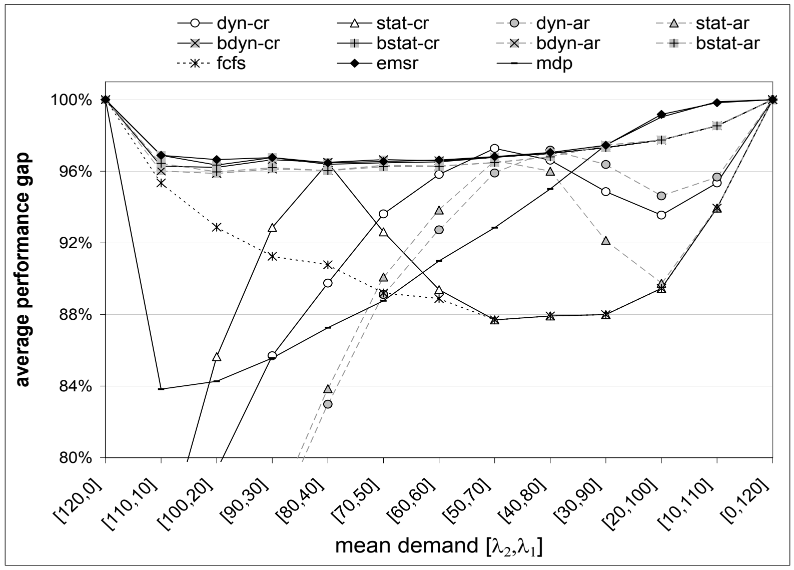

Figure 2.16: Average performance in Example-2.9

sensitive to specification of stationary vs. non-stationary arrivals even though it appears to be robust with respect to the rate of arrivals in a stationary regime. Example-A.11 in the Appendix illustrates another example when demand arrives until the second half of the booking horizon.

2.6 Summary of Chapter 2

These experiments demonstrated that the competitive analysis of online algorithms approach is very promising in RM. Our focus here was on comparing those robust policies with other popular policies given demand bounds. The average revenues obtained by such robust policies in simulation studies are comparable to other well-known procedures. They are very effective and provide consistent results, even though they use no probabilistic information. Further, robust policies are not prone to errors in modeling demand.

Our numerical results indicate that the performance of policies do depend on the choice of demand bounds. Furthermore, any policy that assumes a static demand distribution can perform poorly if there are significant changes in demand characteristics over time. This observation leads to our development of policies based on time-dependent bounds in the next section.

Chapter 3

Robust Dynamic Decision Making in Revenue Management

A revenue management system requires forecasts of quantities such as demand, price sensitivity, and cancelation probabilities, and its performance depends critically on the quality of these forecasts. Although there's little doubt that good forecasting is a vital step, it is usually a high-profile task, consuming huge resources and is error-prone.

In addition to availability of data or credible information, another important issue in RM is the decision makers attitude towards risk. A hotel property manager may be more concerned about breaking-even in the first year and increasing his/her chances of attaining a certain revenue level as opposed to increasing average revenues. A risk neutrality assumption is also questionable in applications of RM to one-time events such as concerts. While recent research in RM proposes models for robust decisions and relaxes the risk neutrality assumption as we reviewed in Chapter 2, the majority of that work proposes static policies that cannot be updated during the booking horizon. These static policies typically rely on demand information that is aggregated over the entire booking horizon. Our work in this chapter extends this literature (i) by proposing *multi-period models* that use time-variant information and (ii) by developing *dynamic booking control policies.* In the airline and hotel industries, booking curves show significant differences across fare classes and over time (see for example Swan, 2002, Liu and Smith, 2002), motivating multi-period models. Furthermore, being able to change the booking limits dynamically over time has clear advantages in a multi-period setting.

In this chapter, we study the single-leg fare-class allocation (booking control) problem in RM. We do not characterize uncertainty using probability distributions. The only information we use about the demand is upper and lower bounds in a given time interval. We develop methods that come with worst-case performance guarantees: we use *competitive analysis of online algorithms* to obtain new booking control policies. Our contributions can be summarized as follows: We propose a new formulation of the single-leg, multiple-fare booking control problem when only time-varying bounds on demand are available. Using competitive analysis, we first analyze a multi-period, static decision problem, and show that booking limits can be determined by solving a mixed-integer program (MIP). We then extend our analysis to the booking control problem in a dynamic setting when the seller can update his/her policy based on history. This is a challenging problem where characterization of the optimal policy is difficult due to the curse of dimensionality. We show that the structural properties of the decision problem in the single-period and the multi-period static models do not generalize to the multi-period dynamic one. We design efficient heuristic methods to obtain booking limits in this setting. Our heuristics provide closed-form solutions. Hence, the computational burden is minimal. Through computational experiments, we reveal the benefit of these new dynamic policies. The average revenues obtained by dynamic heuristic policies can be significantly higher than those of static ones or ones that are commonly used in practice.

The chapter is organized as follows. A literature review is provided in Section 3.1. Section 3.2 introduces competitive analysis and provides a brief description of the multi-period problem. The static, multi-period problem is analyzed in Section 3.3. The dynamic model is analyzed and two effective heuristics to solve the dynamic problem are proposed in Section 3.4. Section 3.5 presents results of computational experiments. Conclusions and suggestions for further research are discussed in Section 3.7.

3.1 Literature Review

We refer the reader to Section 2.1 for a review of research in single-leg RM. In this section, we will only discuss the work that is closet to ours:

Ball and Queyranne (2006) and Lan (2008) are closest to ours in terms of methodology. Ball and Queyranne (2006) use no information about the fare class demand and obtain closed-form optimal solutions (achieving the maximum competitive ratio as defined in the next section) for the static problem. We show their policies are too conservative in our simulation studies in Chapter 2. To increase the effectiveness of these policies, Lan (2008) propose using (estimated) lower and upper bounds on demand. Lan (2008) use static, aggregate information on fare-class demand. Their policies are designed to provide the best performance under the worst-case scenario, and they discuss how the policies can be updated during the booking horizon when the worst-case scenario is not realized, as would be the case in practical situations. Their dynamic policies are designed to use the same, aggregate information as the static ones. An observation on the dynamic policies of Lan (2008) is that a fare-class that has been closed at any time during the booking horizon is never re-opened and these policies aggressively protect seats for higher fare-classes. In contrast to Ball and Queyranne (2006) and Lan (2008), we use time-variant information on fare class demand and extend their work to two different settings: (i) a static setting, where booking limits are determined once at the beginning of the booking horizon given time-variant bounds on fare-class demand, and (ii) a dynamic setting, where booking limits can be dynamically updated to improve the performance given time-variant bounds on fare-class demand. The former shows the benefit of using additional information and serves as a benchmark for the dynamic policies. The latter, as we show in our computational experiments, provides flexible policies, with the ability to re-open a fare class that has been closed, resulting in superior performance.

Recently, robust dynamic programming (RDP) has been proposed to solve dynamic decision problems with ambiguities in problem parameters. We refer the reader to Iyengar (2004) and the references therein for more information. RDP is powerful in analyzing and solving the 'robust' version of a recursive, dynamic programming model. In fact, the dynamic model analyzed by Birbil et al. (2006) falls into the RDP category. In our framework, the objective function for the decision maker is maximizing the competitive ratio, which is defined relative to an offline optimal solution, obtained with hindsight information. This sets our work apart from RDP models.

3.2 Problem Definition

Our approach is similar to that of BQ and Lan (2008). We focus on the CR policy, which is introduced in Section 2.4.3. We use the same notation as in Chapter 2. For the sake of completeness, we repeat the notation and the CR approach in this section, while we introduce the specifics of our multi-period problem.

We study the booking control problem in single-leg RM with m fare classes and n seats (units of capacity). The seller has to make seat allocation decisions across m fares where f_j is the unit revenue from class j, with $f_1 \geq f_2 \geq \cdots \geq f_m$. The planning horizon consists of $T > 1$ time periods. In each period, the seller does not know the actual number of booking requests, nor their arrival sequence. The only information that is available to the seller is lower and upper bounds of demand in each fare class in each period. We use the notation L_j^t and U_j^t for the lower and upper bound, respectively, of demand for fare class j in period t, $j = 1, ..., m$ and $t = 1, .., T$. We use the vectors $L^t = (L_1^t, ..., L_m^t)$ and $U^t = (U_1^t, ..., U_m^t)$ for the demand information in period t.

In practice, and also in research, it is common to use nested booking limits as a booking control policy, where fare classes are nested based on revenue order. These policies are shown to be optimal for the single-leg RM problem in many settings, including Lan (2008). In this paper, we restrict our attention to nested booking limit policies. Nested booking limit policies are based on a booking limit vector $(b_1, ..., b_m)$ with $b_1 \geq ... \geq b_m \geq 0$. Here, b_j denotes the maximum number of booking requests in classes $j, j+1, ..., m$ that can be accepted. In our multi-period setting, where demand characteristics can change by time period, we explicitly model as decision variables a revision to the booking limits at the beginning of each time period. The vector of nested booking limits set at the beginning of period t is denoted $b^t = (b_1^t, ..., b_m^t)$. We only use the standard implementation of nested policies; theft-nesting is excluded from our analysis (see Appendix A for information on standard vs. theft-nesting). We assume the booking limits take continuous values.

In this distribution-free setting, we characterize the arrival of booking requests using *input sequences*. An input sequence consists of a finite stream of fare requests during the booking horizon. We use the notation $\vec{I} = [I^1, ..., I^T]$ to denote the entire stream of booking requests from the beginning of period 1 to the end of period T where I^t is the specific input for period t. An input I^t includes information on the order and the amount of each request in period t. While 'bulk' requests are allowed, we assume each request belongs to a single fare product (this is without loss of generality as will become clear in our analysis of worst-case input streams). We assume any request can be split and accepted partially, i.e., group/bulk reservations cannot be enforced in our model.

We use the simpler notation I for a generic input, when the context does not require time-specific information on the booking requests. Given any input I, we use the notation $I[j]$ to denote the total number of class j requests in input I. Similarly, $I^t[j]$ denotes the total number of class j requests given input I^t in period t. For any input I, we call the vector $(I[1], ..., I[m])$ the *profile* of input I.

3.2.1 Nested booking limits

We now describe the application of standard nesting in a multi-period setting from an algorithmic point of view: The execution of the algorithm (booking control policy) in each period is driven by the period t nested booking limits $b^t = (b_1^t, ..., b_m^t)$, which are specified at the beginning of the period. The algorithm processes the period t input I^t, where requests arrive, requesting units of only one fare class. Let $C(s)$ and $\delta(s)$ be the fare class of and the amount of capacity requested by the s^{th} request, respectively, for $s = 1, ..., ||I^t||$ where $||I^t||$ denotes the total number of requests in input I^t. When standard nesting is applied, each request in I^t is processed using an effective booking limit initialized as $\hat{b}^t = b^t$. The effective booking limits are continuously updated as the algorithm executes to process the input requests in period t. Algorithm **PNEST(t)** below shows the evolution of the effective booking limits in period t.

PNEST(t)

Step 0: Let $n \geq b_1^t \geq \cdots \geq b_m^t$ be the period t nested booking limits, and set $\hat{b}^t = b^t$. Let $s = 1$.

Step 1: After receiving request s, $min(\hat{b}^1_{C(s)}, \delta(s))$ units of it is accepted.

Step 2: Update the effective booking limit starting from class 1

$$\hat{b}_1^1 \leftarrow \hat{b}_1^t - min(\hat{b}^t_{C(s)}, \delta(s)), \tag{3.1}$$

and iteratively computing

$$\hat{b}_j^t \leftarrow min(\hat{b}_{j-1}^t, \hat{b}_j^t - min(\hat{b}^t_{C(s)}, \delta(s))) \quad j = 2, ..., s. \tag{3.2}$$

$$\hat{b}_j^t \leftarrow min(\hat{b}_{j-1}^t, \hat{b}_j^t) \quad j = s+1, ..., m. \tag{3.3}$$

Step 3: If $s < ||I||$, then $s \leftarrow s + 1$ and go to Step 1. Otherwise, stop.

Of course, this procedure would be executed iteratively for each period. The subject matter of this chapter is the determination of the booking limits that drive each execution. One could view each time period as an individual application of a booking control policy so that b^t is set based on the characteristics of that one time period. Alternatively b^1 could be viewed as a booking control policy for all T periods, which might be dynamically adjusted (or not) as one moved from period to period. A static policy that did not make any adjustments would simply set b^t to $\hat{b}^{t-1}$ at the beginning of each period t for $t > 1$.

3.2.2 Performance criteria: Competitive ratio

We employ competitive analysis of online algorithms to determine the nested booking limits in this paper. In competitive analysis, we are interested in determining an algorithm that maximizes the competitive ratio (see Albers, 2003, for more information on competitive analysis and its use in algorithm design and analysis). Competitive ratio (CR) is defined as the *minimum* of the ratio of revenues obtained by an online algorithm to the offline optimal revenues, obtained with hindsight information. For a generic problem, if we let Ω be the set of all feasible (possible) input sequences to an online algorithm Υ and, for any $I \in \Omega$, let $R(I; \Upsilon)$ be the objective value achieved by the online algorithm for input I and let $R^*(I)$ be the objective value achieved by an optimal offline algorithm. Then, CR is defined as:

$$\text{CR of } \Upsilon \;=\; \inf_{I \in \Omega} \frac{R(I;\Upsilon)}{R^*(I)}.$$

Note that the CR as defined above applies to a deterministic algorithm, i.e., algorithm that applies the same decision rule and yields the same output for a given input sequence as opposed to "randomized" algorithms that make some choices based on the draw of a random number. Our focus is on deterministic algorithms in this paper. Given the CR, the seller's problem is to find an algorithm that maximizes the CR

$$\sup_{\Upsilon \in \Pi_\Upsilon} \inf_{I \in \Omega} \frac{R(I;\Upsilon)}{R^*(I)}$$

where the set Π_Υ is the set of feasible/admissable algorithms.

Using the above definitions as starting points, one can view the problem of determining CR and an associated optimal policy as solving an optimization problem defined relative to a very large constraint set (based on all input streams in Ω). On the other hand, it is often instructive to view the problem as a competition between an algorithm designer and an adversary in charge of generating booking requests. The *adversary* is aware of the seller's algorithm (nested booking limits in our case) and chooses an *input sequence* (the number of requests and the arrival sequence) to minimize the algorithm performance (i.e., so that the algorithm achieves the lowest CR). We will use this paradigm to provide intuition and motivate proofs, however, most proofs appeal directly to the definition.

Note that our definition of CR so far does not use specifics of a multi-period problem and does not refer to dynamic decision making. This will be done in more detail in Sections 3.3 and 3.4. In the static model analyzed in Section 3.3, the seller has to commit a policy at the beginning of the booking horizon. This particular model is more of theoretical interest: First, it extends the analysis in Lan (2008) to a multi-period setting and shows the complications associated with multi-period problems. Second, the analysis of the static model serves as a stepping stone for the analysis of the dynamic model, which is introduced in Section 3.4. The dynamic model is designed such that the seller, after observing the performance in prior periods, can choose the booking limits at the beginning of each period. Analysis of the dynamic model allows us to develop effective heuristic policies which are tested computationally in Section 3.5.

3.3 Static Booking Control Policies using Time-varying Demand Information

In this section, we analyze a problem where the seller is restricted to a static booking limit policy. The seller uses multi-period demand information as outlined in the previous section, but does not update his/her policy anytime during the booking horizon. Specifically, for $t > 1$, when PNEST(t) is executed, b^t is set equal to the effective booking limit vector at the end of the prior period, $\hat{b}^{t-1}$. Note that under this approach the execution of the policy is effectively blind to time period boundaries, although the characteristics of demand in each time period could certainly be taken into account in setting the initial policy. Thus, the seller's only decision is the vector b^1 at the beginning of the booking horizon. The optimization problem of the seller in this case can be expressed as

$$\max_{n \geq b_j^1 \geq 0, j=1,..,m} \quad z: \;\; z \leq \frac{R(\vec{I}, b^1)}{R^*(\vec{I})}, \forall \vec{I} = [I^1, ..., I^T], \;\; I^t \in \Omega^t(L^t, U^t) \;\; t = 1, .., T \tag{3.4}$$

where $\Omega^t(L^t, U^t)$ is the set of feasible sequences for period t such that for any $I \in \Omega^t(L^t, U^t)$, we have $L_j^t \leq I[j] \leq U_j^t$ for all $j = 1, ..., m$. $R(\vec{I}; b^1)$ is the revenue obtained by the seller when policy b^1 is used to process input $\vec{I}$ and $R^*(\vec{I})$ is the offline optimal revenue, obtained with perfect hindsight, after observing input $\vec{I}$. We sometimes refer to the policy revenue $R(\vec{I}; b^1)$ as 'the online revenue' or the 'online policy revenue'. Note that the offline optimal revenue $R^*(\vec{I})$ depends only on the profile of the input, but not on the sequence in which requests arrive. Given $\vec{I}[j]$ for $j = 1, .., m$, the offline optimal revenue is easily determined by solving a continuous knapsack problem.

This model differs from that of Lan (2008) as input sequences have period-by-period feasibility requirements. The Lan (2008) model is a special case with $T = 1$. Notice that the optimization problem in (3.4) has infinitely many constraints, which correspond to the feasible input sequences. We will now develop some properties of the static policy, which will lead to an approach to determining an optimal policy.

3.3.1 Sequence reduction in the static model

In order to make the underlying optimization problem (3.4) tractable, we will reduce its size: specifically the number of constraints. We first focus on the order in which fare requests arrive.

Proposition 1 *Given a nested booking limit policy b^1, a period-wise LBH sequence* *minimizes the CR.*

Proof   Consider an input $\vec{I} = [I^1, ..., I^T]$ where I^t is not LBH for some t, $1 \leq t \leq T$. There is a corresponding LBH input I'^t with exactly the same profile as I^t. Define $\vec{I'} = [I^1, ..., I'^t, ..., I^T]$, which replaces I'^t of $\vec{I}$ with I'^t. The offline optimal revenues from sequences $\vec{I}$ and $\vec{I'}$ are the same. Given a static nested booking limit vector b^1, both sequences result in the same decisions/performance prior to period t. Since a nested policy is employed, (i) the online revenues with input I^t are at least as much as the online revenues with I'^t at the end of period t while the remaining capacity, n^t, is the same, and (ii) effective booking limits of fare classes after processing inputs I'^t and I^t are equal (Proposition 1 of Lan (2008)). Therefore, the online revenues, hence the CR, is lower when input $\vec{I'}$ is used. The proof is completed by replacing an input I^t that is not LBH with the LBH input of the same profile for $t = 1, ...T$.   •

While the result in Proposition 1 eliminates many possible input sequences, the constraint space in (3.4) is still huge (actually it is still unbounded). We now introduce the notion of an *extreme sequence* which we show is non-dominated from the adversary's point of view: Inputs with arbitrary profiles result in higher CR compared to the extreme sequences.

Definition 1 *(Extreme input sequence) The j^{th} extreme input sequence for period t is the LBH input with the profile $I^t[k] = L^t_k$ for $k < j$ and U^t_k for $k \geq j$, defined for $j = 1, ..., m$.*

We denote the set of extreme input sequences for the entire booking horizon as Q such that if $\vec{I} = [I^1, ..., I^T] \in Q$, then I^t is one of the m extreme sequences defined for period t, $t = 1, ..., T$. The proof of the next result is in Appendix B.1.

Proposition 2 *Given a static, nested booking limit policy b^1 and the set Q of extreme input sequences, the lowest CR is achieved by one of the inputs in set Q. That is, for any feasible input sequence $\vec{I}$, we have*

$$\frac{R(\vec{I}, b^1)}{R^*(\vec{I})} \geq min\{\frac{R(\vec{I}^*, b^1)}{R^*(\vec{I}^*)}, \quad \forall \vec{I}^* \in Q\}. \tag{3.5}$$

We provide intuition for the proof for $m = 2$ using the adversary paradigm: Consider a one-period problem first. Suppose the seller chooses a booking limit that protects too many seats for class 1. In that case, the adversary chooses to send the minimum feasible amount of class-1 requests so that the seller ends up with idle seats and regrets having rejected class-2 requests early on. Sending the minimum quantity of class-1 requests also reduces the offline optimal revenue, but the difference in offline and optimal revenues favors the adversary's decision; minimum CR in this case is achieved by an input that has fewer class 1 requests. If the booking limit chosen by the seller protects too few seats for class 1, then the adversary prefers to send as many class-1 requests as possible to maximize the offline optimal while minimally affecting the online revenues. The adversary always sends the maximum feasible amount of class-2 requests in this problem. Given that the effective booking limits remain nested after processing each request and that only LBH inputs are of concern, the argument above is easily extended to multiple fares and periods.

There is one important observation about the multi-period problem, though. The number of input sequences can further be reduced from m^T because some of the extreme input sequences for periods that come later in the booking horizon are dominated. This is implicit in the proof of Proposition 2. Given a policy b^1, if it is beneficial for the adversary (increasing the difference between online and offline revenues) to send as many class k requests as possible in period t so that the seller rejects as many of the class k requests in that period, then clearly the adversary prefers to send as many class k requests in the remainder of the booking horizon to ensure the CR is minimized. Additional requests of class k in later periods can increase the offline optimal revenue but not online revenues because they will be rejected by the seller who is using a static nested booking limit policy where booking limits are non-increasing over time.

Corollary 1 *If the minimum CR is achieved by a period-wise LBH input sequence* $\vec{I} = [I^1 \cdots I^T] \in Q$ *with profile* $I^t[k^*] = U^t_{k^*}$ *in period* t *for some* k^**, then* $\vec{I}$ *has profile* $I^s[k] = U^s_k$ *for all* $k \geq k^*$ *and* $s > t$*.*

Using this last observation, the total number of input sequences in a two-period, m-fare problem reduces from m^T to $(m+1)m/2$. However, for longer booking horizons, the number of constraints in problem (3.4) is still in the order of m^T. In practical problems in airline RM, m is typically no less than 3 and no more than 15. When the number of periods is chosen carefully, the number of constraints in problem (3.4) can be kept to a manageable level. Using the properties of nested booking limits, the optimization problem (3.4) can be expressed as a mixed integer programming (MIP) model to determine the optimal static booking limits in this multi-period setting. We provide the MIP model in Appendix B.2, where the differences between the policies obtained by our multi-period static model and the policies derived from the single-period static model of Lan (2008) are also discussed.

3.4 Dynamic Booking Control Policies with Time-varying Demand Information

In this section, we extend the seller's policy to change dynamically over time based on the history of orders processed so far. The decision for the seller at the beginning of period t is the vector of nested booking limits b^t. Let R^t be the seller's accumulated revenue and n^t be the remaining capacity at the beginning of period t. Both R^t and n^t depend on the booking limit vectors $b^1, .., b^{t-1}$ that were used to process input sequences $I^1, ..., I^{t-1}$, respectively, in the previous periods. We call the pair (R^t, n^t) the state of the system at the beginning of period t. Introducing new notation, we use $\vec{b} = (b^1, ..., b^T)$ for the decisions the seller makes throughout the booking horizon. We still use $\vec{I}$ for the entire input sequence consisting of $I^1, ..., I^T$. At the beginning of period t, the seller chooses a new vector of nested booking limits, knowing the state (R^t, n^t). The goal for the seller is to determine the booking limit policy that will maximize the CR for the entire booking horizon.

The seller's CR with multi-period policy vector $\vec{b}$ on sequence $\vec{I}$ is $\frac{R^d(\vec{I},\vec{b})}{R^*(\vec{I})}$ where $R^d(\vec{I}, \vec{b}) = \sum_{t=1}^{T} R(I^t, b^t)$. The seller makes T decisions during the booking horizon with the end-of-horizon goals of maximizing $\frac{R^d(\vec{I},\vec{b})}{R^*(\vec{I})}$.

Similar to the static problem introduced in Section 3.3, one complication in this dynamic problem is its size due to the feasible number of input sequences in each period. Note also that the optimization problem for the seller cannot be expressed recursively in any simple way. The offline optimal revenue is not additive per period, but requires knowledge of the entire T-period input. In the remainder of this section, we use input sequence reduction and linear programming to analyze and provide solutions to this complex decision problem. Our analysis shows that we only need to consider LBH sequences (i.e., these dominate from the adversary's perspective). While we show that the number of non-dominated input sequences is $m+1$ in the last period, the set of non-dominating sequences in any period other than the last cannot be reduced to a manageable level; the decisions are state-dependent. Therefore, we resort to effective heuristic methods in solving the dynamic model.

The heuristic methods that we develop after analyzing the dynamic model allow the seller to re-open a fare class (if needed) and update his/her booking limits in any period based on the current performance of the policy and the state of the system at that point in time. The heuristics rely on solutions of the problem for the last period of the booking horizon, for which sequence reduction is successfully applied and the optimal nested booking limits are computed in closed form.

3.4.1 Analysis of the decisions made in period T when $T > 1$

Given the state (R^T, n^T) and the inputs $I^1, ..., I^{T-1}$ at the beginning of period $t = T$, the CR problem can be stated as follows:

$$\max_{n^T \geq b_j^T \geq 0, j=1,...,m} \quad z \quad : \quad z \leq \frac{R^T + R(I^T, b^T)}{R^*(I^1 \cdots I^{T-1} I^T)} \quad \forall I^T \in \Omega_T(L^T, U^T). \tag{3.6}$$

Similar to our analysis of the static model, we first prove that the input sequence that minimizes the CR in the last period is LBH.

Proposition 3 *The minimum CR is achieved by a LBH input sequence in period t=T.*

Proof Proof is similar to that of Proposition 1 and is omitted. •

We next investigate the adversary's choices for the amount of requests that will arrive. In Lan (2008) for $T = 1$ and in our static model for $T > 1$, the lowest CR is shown to be achieved by one of the m extreme input sequences in a given period. Unfortunately, this set of extreme sequences does not completely characterize the adversary's actions in a multi-period, dynamic problem.

Definition 2 *(Complete set of extreme sequences) The complete set of extreme sequences, denoted Q^{*T}, consists of $m + 1$ LBH sequences, where the j^{th} sequence in Q^{*T}, denoted $I^{*T,j}$ has profile $I^{*T,j}[k] = L_k$ for $k < j$ and $I^{*T,j}[k] = U_k$ for $k \geq j$ for $j = 1, ..., m$, and the $(m + 1)^{st}$ sequence has profile $I^{*T,m+1}[i] = L_i^T$ for all $i = 1, ..., m$.*

Proposition 4 *Given the state* (R^T, n^T)*, the inputs* $I^1, ..., I^{T-1}$*, and a nested booking limit policy* b^T*, the lowest competitive ratio in the multi-period, dynamic problem is achieved by an LBH input belonging to the complete set of extreme sequences* Q^{*T}*, i.e., for any feasible input sequence* $I \in \Omega_T(L^T, U^T)$*,*

$$\frac{R^T + R(I, b^T)}{R^*(I^1 \cdots I^{T-1} I)} \geq min\{\frac{R^T + R(I^T, b^T)}{R^*(I^1 \cdots I^{T-1} I^T)}, \forall I^T \in Q^{*T}\}. \tag{3.7}$$

Proof See Appendix B.3. •

In the single-period and multi-period static formulations, the extreme sequences are such that the adversary always sends the maximum number of requests of class m. Because the booking limits are nested and effective booking limits remain monotonically non-increasing over time (if a class is closed, it remains closed) in the static models, the adversary can improve his offline performance and reduce the CR by sending the highest possible amount of class m in each period. In contrast, the adversary can choose to send the minimum number of class m requests as a way to minimize the CR in the dynamic model. Consider the following scenario: If the seller reaches the last period with a high number of unsold seats and having rejected requests in the previous periods (indicating some fare classes were closed in those periods), then he can revise his decision and reopen some of those classes in the last period. To guard against a situation where the seller significantly increases his revenues by reopening fare classes, the adversary chooses to send the minimum number of requests possible in period T. In this case, the difference between the online revenue and offline optimal revenue is mainly due to differences in decisions made in periods prior to the last: the seller regrets having rejected too many requests in the previous periods because the adversary causes him to have unsold seats at the end of the booking horizon.

3.4.2 Optimal booking limits in period T for $T > 1$

Proposition 4 reduces the number of input sequences to be considered in period T to only $m+1$. This is a significant reduction in the problem size and allows us to develop a linear programming model (LPM) to solve the CR problem given in (3.6).

In LPM, we use the following auxiliary parameters:

$$N_1^T = n^T, N_k^T = [n^T - \sum_{i=1}^{k-1} L_i^T]^+ \quad for\ k = 2, ..., m$$

Instead of the nested booking limits $b_1^T, ..., b_m^T$, we use $x_1^T, ..., x_m^T$ as the decision variables in LPM. We call x_j^T the bucket size for class j in period T and define $x_m^T = b_m^T$, $x_i^T = b_i^T - b_{i+1}^T$ for $i = 1, ..., m-1$.

$$(LPM) \quad \max \quad z$$

s.t.

$$R^*(I^1 \cdots I^{T-1} I^{*T,k})z \quad \leq \quad R^T + \sum_{i=1}^{k-1} f_i L_i^T + \sum_{i=k}^{m} f_i x_i^T, \quad k = 1, ..., m, \tag{3.8}$$

$$R^*(I^1 \cdots I^{T-1} I^{*T,m+1})z \quad \leq \quad R^T + \sum_{i=1}^{m} f_i \min(L_i^T, N_i^T) \tag{3.9}$$

$$\sum_{j=1}^{m} x_j^T \quad \leq \quad n^T \tag{3.10}$$

$$0 \leq \quad x_j^T \quad \leq U_j^T, \quad j = 1, ..., m \tag{3.11}$$

Constraints (3.10) and (3.11) are natural for any booking limit policy: the former assures total seats allocated to fare classes does not exceed the remaining capacity and the latter assures that there is no 'slack' allocation (i.e., the number of seats allocated to a fare class is no more than the maximum demand for that class). LPM has one constraint for each of the extreme sequences to represent the CR achieved by the booking control policy, stated in inequalities (3.8) and (3.9). The closed-form solution of LPM is stated below.

Proposition 5 *Given state (R^T, n^T) and inputs $I^1, ..., I^{T-1}$, the optimal solution to LPM is computed as follows: If $u < m$ or $u = m$ and $\theta_u \leq \pi$, then*

$$z^* = \theta_u \tag{3.12}$$

$$x_k^{*T} = \begin{cases} g_k z^* + L_k^T & 1 \leq k < u \\ (R_k^* z^* - \sum_{i=1}^{k-1} f_i L_i^T - R^T)/f_k & k = u \\ 0 & k > u, \end{cases} \tag{3.13}$$

otherwise

$$z^* = \pi \tag{3.14}$$

$$x_k^{*T} = \begin{cases} g_k z^* + L_k^T & 1 < k \leq m \\ \min(U_1^T, n^T - \sum_{i=2}^{m}(L_i^T + g_i \pi)) & k = 1 \end{cases} \tag{3.15}$$

where the auxiliary parameters are computed as

$$g_m = R_m^*/f_m, \quad g_k = (R_k^* - R_{k+1}^*)/f_k \geq 0, \quad k = 1, ..., m-1,$$

$$u = \max\{j : (R^T + \sum_{i=1}^{j-1} f_i L_i^T) \sum_{i=1}^{j-1} g_i < N_j^T R_j^*\},$$

$$\theta_u = \frac{(R^T + \sum_{i=1}^{u-1} f_i L_i^T)/f_u + N_u^T}{R_u^*/f_u + \sum_{i=1}^{u-1} g_i},$$

$$\pi = \frac{R^T + \sum_{i=1}^{m-1} f_i L_i^T + f_m * min(L_m^T, N_m^T)}{R_{m+1}^*},$$

and R_k^* is the offline optimal revenue of the T-period input $\vec{I} = [I^1, ..., I^{T-1}, I^{*T,k}]$ with extreme sequence $I^{*T,k}$ for $k = 1, ..., m+1$.

Proof See Appendix B.4. •

While the left hand side of the constraints (3.8) and (3.9) represent the offline optimal revenues for the given extreme input, the right hand side provides only an upper bound on the online revenues. Therefore, the formulation in LPM provides an upper bound on the true objective function of the CR problem introduced in (3.6). However, when the binding constraints are taken into consideration only, the optimal solution to LPM yields the correct CR for the given bucket sizes.

Proposition 6 *Given state (R^T, n^T) and inputs $I^1, ..., I^{T-1}$, if z^* is the CR achieved by the nested booking limit policy b^{*T} defined as*

$$b_k^{*T} = \sum_{i=k}^{m} x_i^{*T}, \quad k = 1, ..., m \tag{3.16}$$

where x_k^{*T}, $k = 1, ..., m$ and z^* are the optimal solution to LPM, then z^* is the maximum last period CR as defined in equation (3.6) and b^{*T} is a policy that achieves this maximum.

Proof See Appendix B.5. •

This last result gives credibility to use of LPM, for which closed-form optimal solutions exist.

3.4.3 Analysis of decisions in periods $t < T$

Next, we focus on the decisions made prior to period T. Ideally, one should be able to solve the dynamic problem using backward recursion. While the curse of dimensionality is expected, one can hope to reduce the number of extreme sequences in each period to a manageable level, similar to the case in the static problem. Given the properties of the multi-period static solution, it is not surprising that CR is minimized by period-wise LBH inputs in the dynamic problem.

Proposition 7 *The minimum CR is achieved when input sequences in each period are LBH.*

Proof Proof is similar to that of Proposition 1 and is omitted. •

Unfortunately, further sequence reduction is not possible for a general, dynamic problem because the adversary's non-dominated choices in period $t < T$ depend on (R^t, n^t) and cannot be characterized by a predefined set of extreme sequences. This is in contrast to what is observed in single- and multi-period static models or in the last period of the dynamic model, where the number of extreme sequences are reduced to m or $m + 1$, respectively.

Proposition 8 *Given that the seller restricts her choices to a nested booking limit policy b^t in period $t < T$, the adversary can choose an LBH input I^t in period t with the profile $L^t_k < I^t[k] < U^t_k$ to minimize the CR.*

Proof We provide an example to prove this. Consider $T = 2$, $m = 2$, $n = 14$, $f_1 = 2$, and $f_2 = 1$. The demand information is $U^1_1 = 4$,$L^1_1 = 2$,$U^1_2 = L^1_2 = 5$,

$U_1^2 = 3$,$L_1^2 = 1$,$U_2^2 = 5$, and $L_2^2 = 2$. By Proposition 4, there are three extreme input sequences in the second period, denoted as $I^{*2,1}, I^{*2,2}, I^{*2,3}$, respectively. Consider an LBH input I^1 with the profile $I^1[2] = L_2^1 = 5$ and $I^1[1] = L_1^1 = 2$ for $t = 1$. We call this Input-1. Consider two other alternative LBH inputs for $t = 1$, with profiles $I^1[2] = U_2^1 = 5$, $I^1[1] = U_1^1 = 4$ and $I^1[2] = 5$,$I^1[1] = 3$, called Input-2 and Input-3, respectively. Note that Input-1 and Input-2 represent all the inputs for $t = 1$ where the demand for each fare class takes either the lowest or the highest possible value. (If one were to define the set of extreme sequences for $t = 1$, these two inputs would constitute its superset.)

In Table 1, we fix the seller's and adversary's decisions in period $t = 1$, and compute the optimal decisions in the second period. When we start with $b_2^1 = 0$ and Input-1 in $t = 1$, the resulting CR is 0.6154 after computing optimal decisions in period $t = 2$. When $b_2^1 = 0$ is used to process Input-2 and Input-3 in $t = 1$, the CR is 0.7059 and 0.6666, respectively, after second period optimal decisions are computed for the seller and the adversary. Therefore, if seller chooses $b_2^1 = 0$ in $t = 1$, the adversary would choose Input-1 over Input-2 and Input-3 to achieve the lowest CR. However, if the seller chooses $b_2^1 = 5$ in $t = 1$, the CR are 0.9474, 0.9474, and 0.9444, for Input-1, Input-2 and Input-3, respectively. In this case, Input-3 is clearly a better choice for the adversary: It yields a lower CR compared to Input-2 and Input-3. This shows that given a nested booking limit policy in period $t = 1$, the adversary has non-dominated inputs where the number of requests in each fare class is not equal to the minimum or the maximum feasible amount.

Input-1 in $t = 1$: $I^1[1] = L_1^1 = 2$, $I^1[2] = U_2^1 = L_2^1 = 5$				Input-2 in $t = 1$: $I^1[1] = U_1^1 = 4$, $I^1[2] = U_2^1 = L_2^1 = 5$				Input-3 in $t = 1$: $L_1^1 < I^1[1] = 3 < U_1^1$, $I^1[2] = U_2^1 = L_2^1 = 5$			
	Optimal in $t = 2$				Optimal in $t = 2$				Optimal in $t = 2$		
b_2^1	b_2^2	input	CR	b_2^1	b_2^2	input	CR	b_2^1	b_2^2	input	CR
0	12	$I^{*2,3}$	0.6154	0	10	$I^{*2,3}$	0.7059	0	11	$I^{*2,3}$	0.6666
1	11	$I^{*2,3}$	0.6923	1	9	$I^{*2,3}$	0.7649	1	10	$I^{*2,3}$	0.7333
2	10	$I^{*2,3}$	0.7692	2	8	$I^{*2,3}$	0.8235	2	9	$I^{*2,3}$	0.8000
3	9	$I^{*2,3}$	0.8461	3	7	$I^{*2,3}$	0.8825	3	8	$I^{*2,3}$	0.8666
4	8	$I^{*2,3}$	0.9231	4	4	$I^{*2,3}$	0.9412	4	7	$I^{*2,3}$	0.9333
≥ 5	7	$I^{*2,1}$	0.9474	≥ 5	3	$I^{*2,2}$	0.9474	≥ 5	4	$I^{*2,2}$	0.9444

Table 3.1: An example to characterize the solution to the two-period dynamic problem

•

This property is expected in dynamic decision making. Thus, non-dominated sequences appear difficult to characterize for $t < T$. Therefore, we are unable to reduce the general, dynamic model to a compact form that is amenable to optimization. However, effective heuristics can be developed.

3.4.4 Heuristics methods for the dynamic model

We introduce two heuristic methods that can be used in a rolling-horizon fashion to solve the dynamic, distribution-free, single-leg RM problem. Each of the heuristics takes advantage of the properties of the last period solution.

Rolling Horizon Heuristic (ROH): This heuristic first aggregates bounds of all the periods for each fare class, solves the single-period model of Lan (2008), and determines a static booking limit vector which is used to control the bookings in the first period. Starting from the second period, reoptimization is done by aggregating the bounds of each fare class for the remaining periods, and solving the last-period problem LPM. The solution of LPM is then used to control bookings for one period. At the end of that period new aggregate bounds are computed and the process iterates. That is, LPM is resolved in a rolling-horizon fashion starting from the second period.

Rolling Horizon, Lower-Bounds Adjusted Heuristic (ROL): This approach can be viewed as a slight adjustment to ROH. As in ROH, aggregate bounds are used to obtain the static solution from Lan (2008) prior to the start of the first period. The solution to LPM is used in periods $t > 1$. However, the booking limits provided by Lan (2008) for $t = 1$ or by LPM for $t > 1$, denoted $b^{*t} = (b_1^{*t}, ..., b_m^{*t})$, are adjusted to protect more seats for higher-fare classes early in the booking horizon. The adjustment results in new booking limits, denoted $b_j^{*t(a)}$, and are computed starting with class m,

$$b_m^{*t(a)} = max(0, (b_m^{*t} - \sum_{s=t+1}^{T} L_m^s)^+),$$

$$b_j^{*t(a)} = max(b_{j+1}^{*t(a)}, (b_j^{*t} - \sum_{s=t+1}^{T} L_i^s)^+), \quad j = m-1, m-2, ..., 1. \tag{3.17}$$

Equation (3.17) can be interpreted in this way: LPM suggests a booking limit of b_j^{*t} for class j while the guaranteed minimum number of requests in class j <u>after</u> period t is $\sum_{s=t+1}^{T} L_j^s$. In this case, one may be able to delay selling to class j: the booking limit of class j in period t is reduced by the quantity of class j requests guaranteed to arrive in later periods. This has the effect of reserving more capacity for higher fare orders in time period t, knowing that if the seats that are protected for higher fares are unsold at the end of period t, they can be sold to class j in the remaining periods.

3.5 Computational Results

The models we developed in this chapter improve the worst-case performance compared to existing methods that use competitive analysis. However, it is typical for worst-case analysis to provide very conservative solutions to a decision problem and the average performance of suggested policies may not be superior in practical situations. We designed computational experiments (i) to quantify the performance of our dynamic policies, (ii) to compare our policies to other well-known procedures in multi-period, single-leg RM, and (iii) to show the practical value of our heuristics under various demand scenarios and parameter settings.

We evaluate several booking control methods using simulation. Below is a complete list. Note that some of methods require probabilistic characterization of demand.

- OFFLINE is the policy that has hindsight information. After each simulation run, the offline optimal revenue and offline optimal policy are determined with perfect information.

- FCFS is the first-come-first-serve policy that accepts any booking request up to the capacity. Effectively, all booking limits are equal to the (remaining) capacity in this policy.

- ROH is our rolling-horizon heuristic as described in Section 3.4.4.
- ROL is our rolling-horizon heuristic that protects more seats for higher fares early in the booking horizon; see Section 3.4.4.
- STSE is the static, nested booking limit policy derived from our static, multi-period model by solving the MIP provided in Appendix B.2. The booking limits of this policy are not updated during the booking horizon.
- ASTAT is the nested booking limit policy derived from the static, single-period model of Lan (2008). The demand information is aggregated, i.e. $\sum_{t=1}^{T} L_i^t$ and $\sum_{t=1}^{T} U_i^t$ are used as the lower and upper bound, respectively, of fare-class demand in this model.
- TROH is the rolling horizon application of ASTAT, i.e. ASTAT is resolved at the beginning of every period based on the aggregate future demand bounds and the remaining capacity to obtain a nested booking limit policy.
- REMSR is the rolling horizon application of the nested booking limits suggested by of Littlewood (1972) for $m = 2$, see Chapter 2.
- REMSR-b is the rolling horizon application of the EMSR-b method of Belobaba (1989), see Chapter 2.
- ADYN is the dynamic method suggested by Lan (2008) that uses aggregate demand bounds at the beginning of the booking horizon to determine nested

booking limits and updates the policy parameters after each accepted request.

- MDP is the Markov Decision Process model of Lee and Hersh (1993), see Chapter 2. We evaluate this policy in the experiments where the demand is Poisson distributed.

In several experiments, we report the performance of the policies relative to OFFLINE. We call this measure *performance gap.* The average performance gap for a policy is obtained by taking the ratio of the policy revenue to OFFLINE revenue at the end of each simulation run, and then computing average of this ratio across all simulation runs in an experiment.

Each instance in each experiment involves 6000 simulation runs. The parameters of our policies are computed by setting the lower and upper bounds of demand to two standard deviations away from the mean demand of each fare class. Requests of each fare class arrive in random order, with arrival times distributed uniformly, in each period.

3.5.1 Experiments with only two fare classes

We use an example with $m = 2$ and $T = 3$. Demand of for each fare class is normally distributed in this section unless noted otherwise.

Example 3.1 In this example, capacity is 200 seats and the demand factor is 1.2. The fare ratio, defined as $r = f_2/f_1$, ranges from 0.1 to 1.0. Mean demand for each fare class in each period is listed in Table 3.2, the coefficient of variation (CV) of demand is 1/3 for each fare class in each period, except for class 1 in the second period, which has a CV of 1. This parameter setting represents the practical situation where demand for class 1 arrives later, and is more volatile early in the booking horizon.

Class	Period-1	Period-2	Period-3	Aggregate
1	0	40.0	80.0	120.0
2	40.0	40.0	40.0	120.0

Table 3.2: Expected demand in Example 3.1

Figure 3.1 shows how the average performance gap of each policy varies with r. Note that the use of dynamic policies and time-dependent demand information significantly improves average performance, especially when r is in the middle range in this experiment: The performance gap of ADYN, ROL or ROH is about 2% higher than that of ASTAT and STSE. This improvement illustrates the value of re-optimizing the booking limits considering revenue on-hand and remaining capacity at the end of each period. REMSR is dominated by heuristics ROH and ROL, even though REMSR uses a probabilistic characterization of demand and updates the booking limits in each period. TROH continuously tries to guard against the worst-case scenario without taking into account accumulated revenue in each period; therefore its performance is inferior.

Example 3.2 The set up is the same as in Example 1 here, except for changes in the demand characteristics. The expected demand values are listed in Table 3.3, where the expected total demand for class 1 is only half of the expected demand for class 2. Figure 2 displays the average performance gap of the policies.

In this example, our heuristics ROH and ROL are again almost indistinguish-

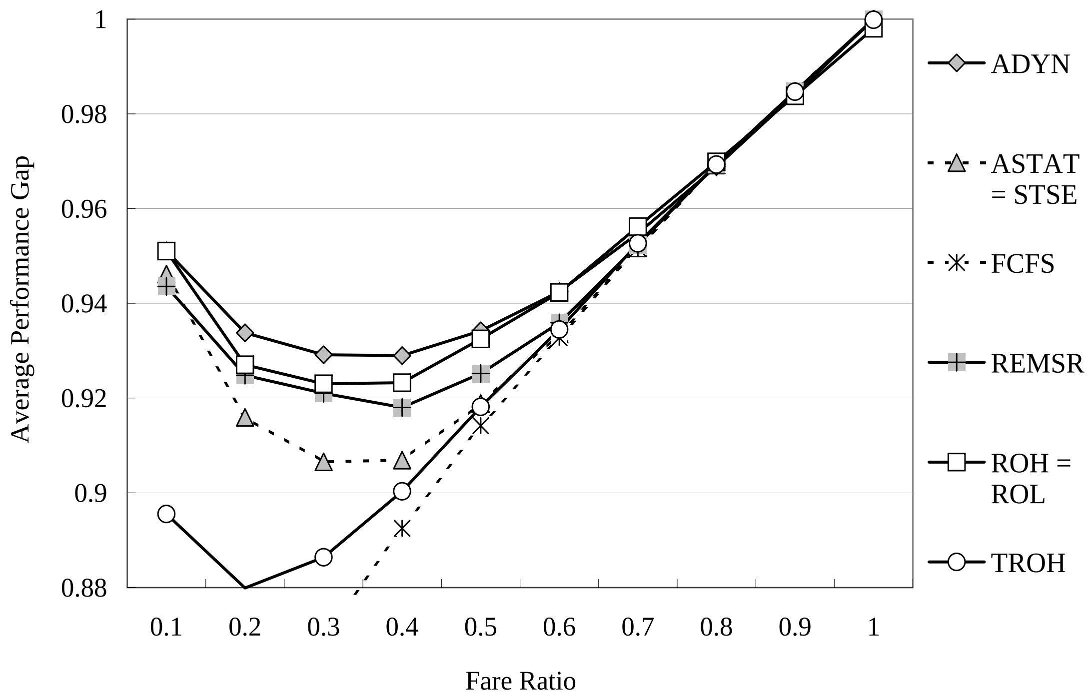

Figure 3.1: Average performance of policies in Example 3.1

Class	Period 1	Period 2	Period 3	Aggregate
1	0	20.0	60.0	80.0
2	56.0	52.0	52.0	160.0

Table 3.3: Expected demand in Example 3.2

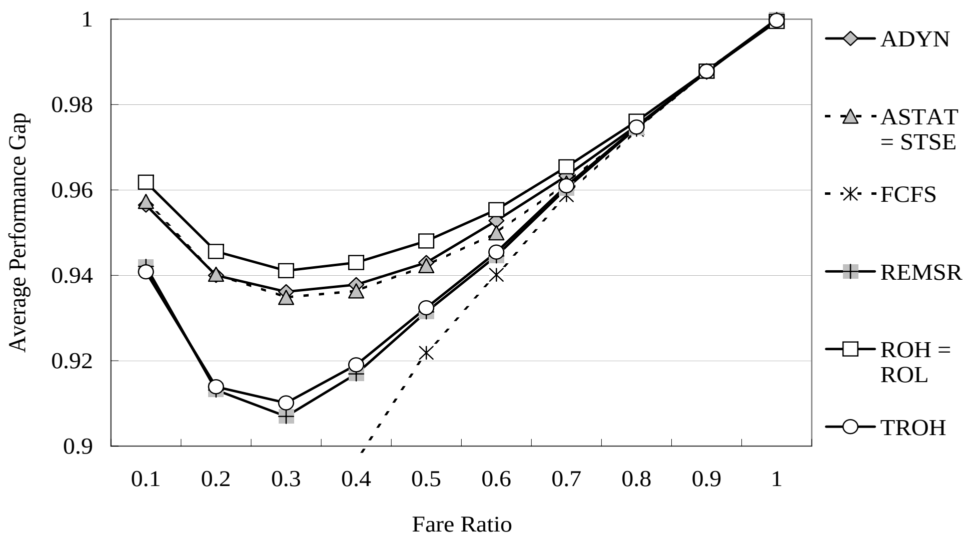

Figure 3.2: Average performance of policies in Example 3.2

able from each other and they dominate policies that do not use time-dependent information. Notice that their performance is also significantly better than REMSR, which has inferior performance because the demand for class-1 varies significantly across time periods. In our experience, REMSR typically performs well when the mean demand is stable over time and the variance is small.

Example 3.3. In this experiment, we increase the total capacity in Example 1 to 240 seats, while keeping the demand parameters the same. This reduces the expected demand factor to 1.0. The main observations from Figure 3 are similar to those of Example 1, with the notable difference being in the relative performance of ADYN, which is dominated in this case. Our heuristics ROL and ROH still have higher performance than ASTAT, STSE, REMSR and TROH. Both ROL and ROH policies close and then reopen class 2 in this experiment when the remaining capacity at the beginning of period 3 is high.

The results of these experiments are very encouraging: The proposed two-stage re-optimization heuristics ROL and ROH perform better than ASTAT and STSE using aggregate demand information, much better than the REMSR, which is commonly used in airline RM practice, and considerably better than a policy that naively applies the methods of Lan (2008) in a rolling horizon framework. Besides, ROL and ROH are more robust to changes in fare ratio, load factor, demand mix. Finally, these dynamic policies close and reopen lowest fare class (if favorable).

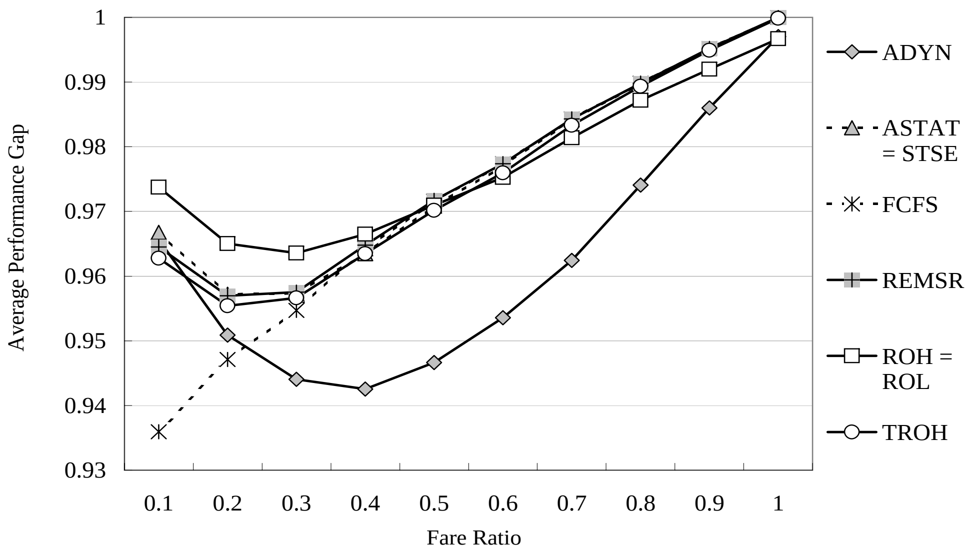

Figure 3.3: Average performance of policies in Example 3.3

3.5.2 Experiments with $m > 2$

The above examples are limited to two fare classes, each with moderate demand variability (CV is at most 1). This second group of experiments are implemented for examples with $m = 4$ and $T = 3$. TROH is excluded from the experiments in this section due to its inferior performance.

Example 3.4 Fare-class demand is normally distributed in this example. The fares and mean and standard deviation of demand are listed in Table 3.4. We experiment with three different choices of demand variability, differentiated as L (low), M (medium), and H (high) in Table 3.4. The capacity is 250 and the average demand factor is 1.23.

		Period-1 Demand				Period-2 Demand				Period-3 Demand			
			stdev				stdev				stdev		
class	fare	mean	L	M	H	mean	L	M	H	mean	L	M	H
4	350	13.2	5	10	20	13	5	10	20	13.2	5	10	20
3	527	30	10	20	30	60	20	30	73.6	56.6	20	30	40
2	567	0	0	0	0	45.1	18	35.1	57.1	45.1	10	20	30
1	1030	0	0	0	0	10	5	20.6	37.7	24.6	8	12.7	24.7

Table 3.4: Parameters used in the experiments in Example 3.4

As observed in Figure 3.4, when the variance of demand is low, all policies, except FCFS, perform well, achieving at least 92% of the OFFLINE. When demand variance increases, the average performance of all methods decreases. However the performance gap of REMSR-b degrades more than other methods, and it can only achieve 81% of OFFLINE revenues under high variance. This is expected because EMSR-b is not able to accurately pool the demand information. The key take-

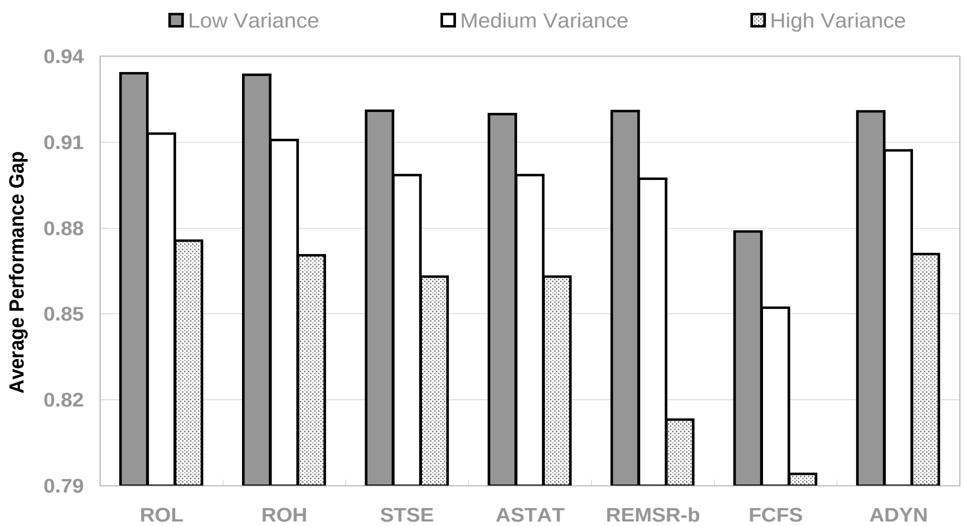

Figure 3.4: Average performance of policies in Example 3.4

away from this experiment is that both ROL and ROH have the highest average performance gap, outperforming all others by more than 1.5% when demand variance is low, and ROL, ROH, ADYN outperform others by over 1% when variance is medium or high. This shows the practical value of heuristics ROH and ROL. While the differences are small, they are statistically significant: Table 3.5 reports the standard error of the performance gap for each policy when demand variance is high.

ROL	ROH	STSE	ASTAT	REMSR-b	FCFS	ADYN
0.038%	0.039%	0.040%	0.041%	0.0375%	0.047%	0.032%

Table 3.5: Standard error of the estimates of performance gap in Example 3.4

Example 3.5 In this example, the demand in each period is Poisson distributed. The problem parameters are listed in Table 6. We experiment with different demand factors in this example. When the capacity is 225, 250, and 275, the average demand factor is 1.30, 1.17, and 1.07, respectively. This is one experiment where MDP provides the optimal policy for the underlying stochastic model: MDP uses complete information about the arrival process and achieves approximately 99% of the OFFLINE revenues as seen in Figure 5. ROL and ROH outperform all other heuristics based on the performance gap. In addition, ROL and ROH have more stable performance compared to others. When the load factor is lower (capacity is 250 or 275), both ROL and ROH re-open fare classes that are closed, which explains the difference between these heuristics and ADYN in this experiment for the case where variance of demand is not high.

class	fare	Mean demand in		
		period 1	period 2	period 3
4	350	23.2	23.2	23.2
3	527	20.0	40.6	40.6
2	567	0	35.1	35.1
1	1030	0	20.0	24.6

Table 3.6: Parameters used in the experiments in Example 3.5

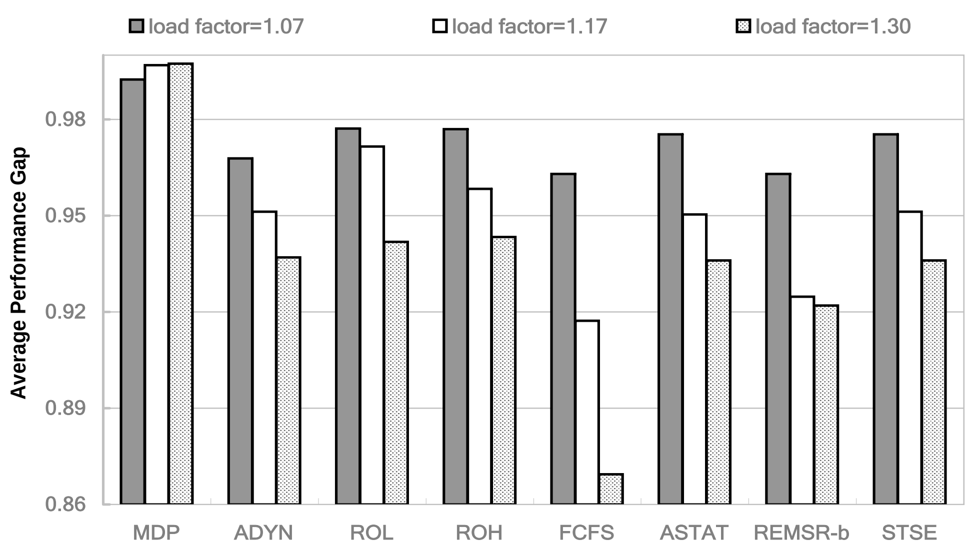

Figure 3.5: Average performance of policies in Example 3.5

3.6 Conclusions

In this paper, we have analyzed the traditional single-leg, multi-fare seat inventory control problem in RM from the perspective of competitive analysis of online algorithms. Our models make use of time-varying bounds on demand in each fare class as opposed to static, aggregate demand information. We derived a static problem formulation and proposed dynamic heuristics when the seller's objective is to maximize the competitive ratio. Our two-stage re-optimization heuristics have significant practical advantages: the computational burden is minimal, the policies provide chances for reopening fare classes (if needed) that are previously closed, the average performance gain over other heuristic policies are significant, and policies exhibit more stable and robust performance across different problem scenarios. Since we employ very few assumptions when modeling the time-variant demands, our policies are less prone to problems associated with modeling errors. They are viable and effective solutions for the single-leg RM problem when demand exhibits variability over time and across fare classes.

We discuss possible extensions and future research directions below.

Other Relative Performance Measures: Our analysis focused on determining dynamic booking limits in a distribution-free environment where the goal is to maximize the competitive ratio, which is computed relative to a solution that has hindsight information. Another alternative is use *absolute regret* (AR) in determining the booking limit policies. The analysis of the static and dynamic problems with the AR measure would be done in the same way as the CR measure. In that case, the adversary's goal is to maximize the difference between the offline optimal revenues and the policy revenue, while the seller tries to minimize this maximum absolute regret. The static and dynamic problems are analyzed as we did in this paper. The results on sequence reduction are the same for both the CR or AR criterion. The solution of a MIP model provides the static solution to the AR problem, and heuristics are used to get the dynamic solution. Details are omitted. We refer the reader to Gao (2008) for experimental performance of policies that are derived using the AR criterion.

Design of Dynamic Model: In our static model both the seller and the adversary make decisions only once at the beginning of the booking horizon, while they update their decisions periodically in the dynamic model. An alternative way to get dynamic policies is to focus on a hybrid model that merges features of the static and dynamic models in order to favor the seller: The adversary commits to the entire input sequence a priori and cannot update the inputs anytime during the booking horizon, while the seller can dynamically change his/her booking limits after observing the state of the system at the beginning of a period. While this is a viable approach from a modeling perspective, it should be clear from our analysis in this chapter that such an alternative model is still technically challenging, and would not lead to an optimization problem of manageable size.

Choice of Time Intervals and Demand Bounds: The analysis in this paper assumed that all demand information was known a priori and that the time period boundaries where demand characteristics could change were fixed. An interesting future research direction is to model the situation where demand information is acquired over time, leading to changes in the demand bounds for future periods. Other areas worth investigating are the choice of bounds to use with our models given the existing forecasting information from RM systems. Similarly, it would be useful to investigate how one would choose the number and length of periods for the booking horizon, given available information on how demand changes over time.

3.7 Summary

In this chapter, we analyzed the distribution-free single-leg RM problem from the perspective of competitive analysis of online algorithms. Our models make use of time-varying bounds on demand in each fare class as opposed to static, aggregate demand information. We propose both static and dynamic booking limit policies to maximize CR or minimize the maximum AR. Our two-stage re-optimization heuristics for the dynamic problem have significant practical advantages: the computational burden is minimal, policies provide chances of reopening fare classes, the average performances gain over policies that only use aggregate information are impressive and policies are robust to demand fluctuation. Besides, since we make very little assumptions on modeling the time-variant demands, our policies are less prone to problems associated with demand modeling errors. Finally, these solutions perform significantly better than the rolling horizon methods used in current practice. They are viable and effective solutions for the single-leg RM problem when demand changes overtime, but cannot be characterized with as much detail or accuracy. In this paper, however analysis was carried out assuming the demand information is given and static. It will be interesting to combine the heuristics we proposed with demand bounds estimation procedures in a entirely dynamic manner, investigating how demand bounds can be estimated, what types of optimization-estimation procedures make our policies most effective, and how distribution-free methods adapt when there's changes in demand information. investigating how demand bounds can be estimated, what types of optimization-estimation procedures make our policies most effective, and how distribution-free methods adapt when there's changes in demand information.

Chapter 4

Robust Decision-Making and Competition

4.1 Introduction

In the presence of limited demand information, the effectiveness of the control decisions can be increased by developing models and methods based on relative regret as we have seen in Chapter 2 and 3. Although there is a recent stream of operations research literature on robust pricing and capacity control, work on competitive inventory management in RM is relatively rare. In reality, what we usually see is two competing airlines offer flights on the same route, departing within minutes of each other, and having similar fares. For example, if a traveler is looking for a late morning flight from New York City (JFK) to San Francisco (SFO) in August 2008, he can choose between two airlines, Virgin America and JetBlue, which both offer non-stop flights 10:55AM and 10:45AM, respectively, for nearly the same price (around \$420 for the lowest non-refundable fare and approximately \$650 for the refundable fare). Now suppose that a customer wishes to purchase a non-refundable ticket. If the seats in that fare class have sold out on JetBlue, it is very likely that the customer will attempt to purchase a ticket in the same fare class on the Virgin America flight that departs 10 minutes later. A similar situation exists when a major carrier has back-to-back flights on the same route with identical fares. For example, from Chicago (ORD) to Washington DC (IAD), United Airlines offers non-stop flights at 8:00AM and 9:00AM with exactly same fare classes. If a customer can not be accommodated by the lowest fare class on the 8:00AM flight, he may look for seats of the lowest fare on s the 9:00AM flight on this route. In this latter case, there is only one carrier. However, if the airline independently manages the bookings on these flights, with the objective of maximizing returns on each flight (as opposed to centralized planning to maximize the total return on this route), the RM problem is very similar to the competitive one we described above.

If we explicitly consider that there is competition and no collaboration on seat allocation decisions of multiple flights, then the problem can be viewed as a game and the resulting decisions that arise out of the resulting game can differ significantly from the seat allocations that would be optimal for a single flight or for a single decision-maker. Hence, our main focus in this chapter is to analyze the robust capacity control problem in a competitive framework. We also compare the effect of centralized versus decentralized decision-making

4.2 Literature Review

We reviewed the literature on single-leg RM in Section 2.1 and Section 3.1. In this section, we provide a short discussion on research that focuses on competition in airline RM.

Netessine and Shumsky (2005) examine the seat allocation problem of two airlines for two fare classes, taking into account explicitly the overflow of demand for a class of one airline to the same class of the other airline. Li et al (2006) Li and

Zhang (2007) studied a similar problem except that they choose a homogeneous approach: two airlines face a common market demand and demand will be split between the two airlines. The objective in these studies, however, is still focused on expected revenues given risk-neutrality of the decision-maker and given a probabilistic demand distribution. The overflow model in Netessine and Shumsky (2005) is similar to the newsvendor model of Lippman and McCardle (1997), who study the classical newsboy problem in a competitive setting and show that competition can lead to higher inventories. Note that there is considerable research on competitive pricing decisions (e.g., Kleywegt and Cooper 2006), but our focus in this work is on capacity control given the fares.

More recently, Jiang and Netessine (2007) analyze competition among newsvendors when the only information competitors possess about the nature of future demand realizations is the support of demand distributions. In their analysis, they focus on several alternative criteria used in the robust optimization literature, such as relative and absolute regret, as well as worst-case performance. Using these robust criteria, they establish the unique Nash equilibrium solution for a (symmetric) game with an arbitrary number of players. Their work motivates our study on robust booking control problem in a decentralized competitive market.

4.3 Model with Two Fare Classes

Suppose there are two direct flights between the same origin and destination, with departures and arrivals at similar times. We use i,j to distinguish these two flights. Flights have the seat capacity n^i and n^j and there are two fare classes available for passengers: 'class-1' and 'class-2.' We assume the two flights offer the same fare structure. Specifically, for the base model, we consider a two-fare setting where $f_2^i = f_2^j$, and $f_1^i = f_1^j$, where f_k^t is the unit revenue obtained by selling in class k on flight t, $k = 1, 2$, $t = 1, 2$.

Demand follows the basic assumption of the "independent demand" model that is commonly used in revenue management: Each customer requests reservation in only one fare class, but does not buy-up or down to another fare class if their first choice is not available. However, customers whose request is not fulfilled by airline i "overflow" to airline j (in the same fare-class) in our problem. In our model, we also assume that class-2 requests occur before class-1. This is the case when the seller uses advance purchase restrictions to distinguish between low-valued and high-valued customers.(e.g., 14-day advance purchase tickets typically have a lower fare than unrestricted tickets.) This is the LBH assumption is commonly used in the classical literature in RM; see the review in Chapter 2.

Our approach is distribution-free and the demand is characterized by an input for each flight. A sequence consists of a finite stream of fare requests during the booking horizon. Notation I^i and I^j are used to denote the input sequences for flights i and j respectively. Each input has a specific number of requests for each class. Without loss of generality, we assume each request in a sequence demands one unit. Let $I^i[k]$, $I^j[k]$ be the total number of class k requests in sequence I^i and I^j. Because of our LBH assumption, class-2 requests in I^i are followed by class-1 requests.

Both flights use nested booking limits to process incoming booking requests. This means class-1 has access to all the capacity while class-2 requests are accepted up to the booking limit, denoted b_2^i and b_2^j for flight i and j respectively. Demand for class-k in any input sequence is assumed to lie in the range $[L_k, U_k]$, namely, $L_k^i \leq I^i[k] \leq U_k^i$,$L_k^j \leq I^j[k] \leq U_k^j$, $k = 1, 2$.

The sequence of events is as follows and occur simultaneously for both flights:

1) Set booking limits b_2^i and b_2^j.
2) Flight i and j observe $I^i[2]$ and $I^j[2]$ respectively. Class-2 requests are accepted/denied based on standard nesting given the booking limits.
3) Any rejected class-2 request in flight i (j) overflow to alternative flight j (i). This overflow demand is again processed according to nesting rules and booking limits.
4) $I^i[1]$ and $I^j[1]$ requests arrive to flights i and j respectively, and are accepted based on the remaining capacity.
5) Any rejected class-1 request in flight i (j) overflows to the alternative flight j (i) and are accepted if there is any remaining capacity.

Figure 4.1 summarizes the notation, the demand and overflow processes.

The objective for flight i given the booking limit of the alternate flight j is:

$$\max_{b_2^i} z : z \leq \frac{R^i(I^i, b_2^i | I^j, b_2^j)}{R^{i*}(I^i, I^j, b_2^j)} \qquad I_k^i \in \Omega(L_k^i, U_k^i), I_k^j \in \Omega(L_k^j, U_k^j) \tag{4.1}$$

where $R^i(\cdot)$ is revenue for flight i when I^i is processed by b_2^i and overflow de-

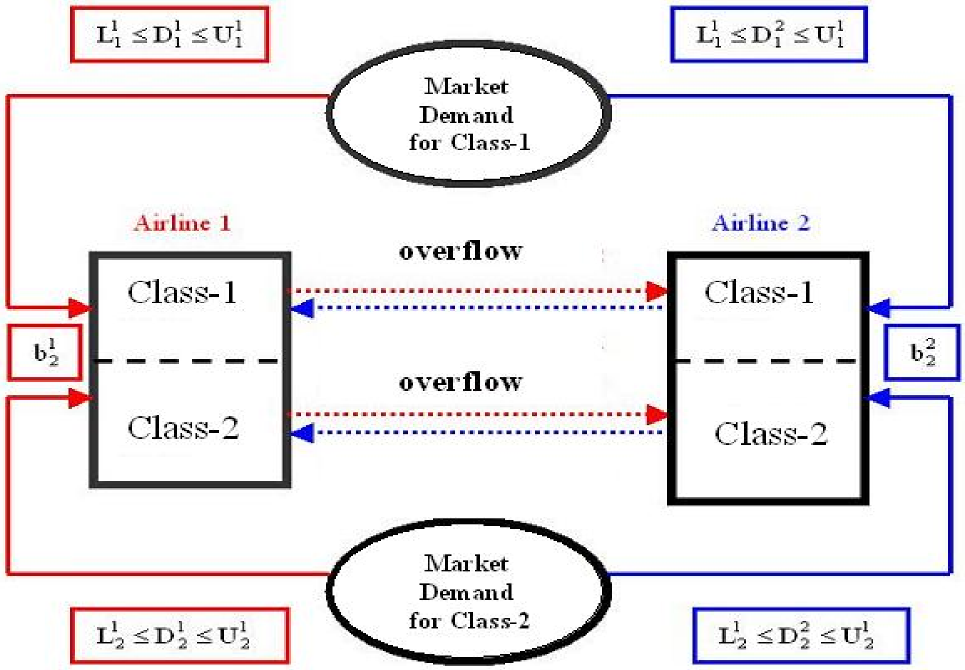

Figure 4.1: Horizontal Competition

mand occurs, given I^j and b_2^j for flight j. We refer to this as "conditional online revenue". $R^{i*}(I^i, I^j, b_2^j)$ is the offline revenue, obtained with perfect hindsight information on the aggregate demand (including the overflow). Note that the seller is interested in maximizing the competitive ratio (CR) in flight i. A similar formulation can be developed for the AR criterion. Note that by solving problem in (4.1), the decision-maker chooses a policy that has a worst-case guarantee.

4.4 Analysis of Decentralized Decision-making

Instead of a single carrier, suppose two competing airlines offer these two flights. The unaccommodated customer, who wishes to purchase a ticket in airline i, will attempt to request a ticket in the same fare class on competing airline j. The class-2 (class-1) demand for airline j can expand if airline i set a low (high) booking limit. So in the presence of competition on the same route, an airline might set its seat inventory control rule differently from the single carrier case.

In our problem, both airlines simultaneously solve (4.1) to determine the best booking limits in a competitive setting with perfect information. Our goal is to determine each airline's strategy, whether or not equilibrium exists and how the airline's booking control policies under this setting can be computed.

4.4.1 Extreme Input Sequences

Note that airline i's optimal decision based on (4.1) is determined by the worst case scenario(s). Naturally, not all the inputs yield lowest CR given booking limits, i.e., constraints associated with some inputs are redundant. We refer to the inputs that result in non-redundant constraints as "extreme inputs". We provide a characterization of the extreme inputs in the market with two competitors below. We refer to the vector $(I[2], I[1])$ as the *profile* of the LBH input I.

Proposition 9 *Given the LBH input I^j and nested booking limits b_2^i and b_2^j, the LBH input extreme sequences ES_1^i,ES_2^i with the following profile*

$(ES_1^i[2], ES_1^i[1]) = (U_2^i, U_1^i)$ and $(ES_2^i[2], ES_2^i[1]) = (U_2^i, L_1^i)$ respectively

provide the lowest CR for airline i. That is,

$$\min\{\frac{R^i(ES_k^i, b_2^i|I^j, b_2^j)}{R^{i*}(ES_k^i, I^j, b_2^j)}\ k = 1,2\} \leq \frac{R^i(I^i, b_2^i|I^j, b_2^j)}{R^{i*}(I^i, I^j, b_2^j)}, \quad \forall I^i \in \Omega^i(L, U).$$

Similarly, ES_1^j and ES_2^j with profile $(ES_1^j[2], ES_1^j[1]) = (U_2^j, U_1^j)$ and $(ES_2^j[2], ES_2^j[1]) = (U_2^j, L_1^j)$ provide the lowest CR for airline j, given an arbitrary input I^i and nested *booking limits b_2^j and b_2^i.*

Proof Given I^j and b_2^j, the overflow amount from flight j to flight i is fixed. So we seek the worst-case input within $\Omega^i(L,U)$. The proof is immediate following the argument in Proposition 2 of Chapter 3 and can be shown using the technique of Lan(2008) who shows that these extreme inputs provide the lowest CR for a single decision-maker with a single flight. $\bullet$

Based on this result, we have to consider only a total of four extreme input sequences for flights i and j. However, the best response function of an airline can be characterized by focusing on even a smaller set of inputs:

Proposition 10 *Given the booking limits b_2^i and b_2^j, either input pair $\{ES_1^i, ES_1^j\}$ or input pair $\{ES_2^i, ES_2^j\}$ provide the lowest CR for airline i.*

$$\min\{\frac{R^i(ES_k^i, b_2^i|ES_k^j, b_2^j)}{R^{i*}(ES_k^i, ES_k^j, b_2^j)},\ \ k=1,2\} \leq \frac{R^i(I^i, b_2^i|I^j, b_2^j)}{R^{i*}(I^i, I^j, b_2^j)}\ \ \forall I^i \in \Omega^i(L,U),\ \forall I^j \in \Omega^j(L,U)$$

.

Proof Let us denote the "aggregate demand" of class-k for flight i as $AD_k^i(I^i, I^j, b_2^j) = I^i[k] + OF^i(I^j[k], b_2^j)$, where $OF^i(I^j[k], b_2^j)$ is the overflow in class-k from flight j to flight i. Let $I^{AD(i)}$ be a pseudo input for airline i such that it represents the aggregate demand observed by the airline including the overflow. That is, $I^{AD(i)}[k] = AD_k^i(I^i, I^j, b_2^j)$. Then airline i's problem is:

$$\max_{b_2^i \geq 0} \quad z \leq \frac{\hat{R}^i(I^{AD(i)}, b_2^i)}{\hat{R}^{i*}(I^{AD(i)})}$$

$$\forall I^{AD(i)} \in \Omega^{AD(i)}(L^{AD(i)}, U^{AD(i)}).$$

where $\hat{R}^i(\cdot)$ is the online revenue, $\hat{R}^{i*}(\cdot)$ is the offline revenue, and $\Omega^{AD(i)}$ is the set of all LBH inputs with demand bounds $L^{AD(i)}$ and $U^{AD(i)}$ correspond to the aggregate demand to be observed by flight i. Lan(2008) and Proposition 9 above show that the two extreme inputs for this problem have the profile $(U_2^{AD(i)}, U_1^{AD(i)})$ and $(U_2^{AD(i)}, L_1^{AD(i)})$.

By construction, $U_k^{AD(i)} = \max\{I^{AD(i)}[k] : \forall I^{AD(i)} \in \Omega^{AD(i)}(L^{AD(i)}, U^{AD(i)})\}$, which is equal to: $U_k^{AD(i)} = U_k^i + OF^i(U_k^j, b_2^j), k = 1, 2$ in our problem. Similarly, $L_k^{AD(i)} = \min\{I^{AD(i)}[k] : \forall I^{AD(i)} \in \Omega^{AD(i)}(L^{AD(i)}, U^{AD(i)})\}$ is equal to: $L_k^{AD(i)} = L_k^i + OF^i(L_k^j, b_2^j), k = 1, 2$ in our problem.

Therefore, the extreme inputs in our problem are represented by an input pair with profile (U_2^i, U_1^i) and (U_2^j, L_1^j) or an input pair with profile (U_2^i, L_1^i) and (U_2^j, L_1^j).

•

4.4.2 The Best Response Function

Based on the characterization of the extreme inputs, airline i's best response function can be easily expressed as a LP in the following way. We use bucket size x_1^i, x_2^i as the decision variable, and $x_1^i = b_1^i - b_2^i$, $x_2^i = b_2^i$. When airline j uses booking limit b_2^j, airline i's best nested booking limit policy that maximizes CR is obtained by solving the problem below.

$$z^{i*}(b_2^j) = \max_{x_1^i, x_2^i \geq 0} \gamma \tag{4.2}$$

$$R^{i*}(ES_1^i, ES_1^j, b_2^j)\gamma \leq f_1^i x_1^i + f_2^i x_2^i \tag{4.3}$$

$$R^{i*}(ES_2^i, ES_2^j, b_2^*)\gamma \leq f_1^i L_1^i + f_2^i x_2^i \tag{4.4}$$

$$x_1^i + x_2^i \leq n^i \tag{4.5}$$

$$0 \leq x_k^i \leq U_k^i + OF^i(U_k^j, b_2^j), k = 1, 2. \tag{4.6}$$

Notice that $OF^i(\cdot)$ and $OF^j(\cdot)$ terms can be easily computed given b_2^j and the demand bounds.

Proposition 11 *The best response of airline i given b_2^j is:*

$$z^{i*}(b_2^j) = \frac{f_1^i L_1^i + f_2^i(n^i - L_1^i)}{R_2^{i*} + f_2^i/f_1^i(R_1^{i*} - R_2^{i*})} \tag{4.7}$$

$$x_1^i(b_2^j) = (R_1^{i*} - R_2^{i*})/f_1^i * z^{i*} + L_1^i \tag{4.8}$$

$$b_2^{i*}(b_2^j) = x_2^i(b_2^j) = \frac{1}{f_2^i}[\frac{(n^i - L_1^i)f_2^i + f_1^i L_1^i}{1 + f_2^i/f_1^i(R_1^{i*}/R_2^{i*} - 1)} - f_1^i L_1^i] \tag{4.9}$$

where, $R_k^{i*} = R^{i*}(ES_k^i, ES_k^j, b_2^j)$.

Proof   Follows from the solution of the single-leg problem defined in Lan(2008); see equation 2.15 and 2.16 of Chapter 2. •

So the best response booking limit b_2^{i*} as a function of b_2^j is computed with minimal computational requirements. We identify further properties of the best response functions:

Proposition 12 *The best response functions satisfy:* $b_2^{i*}(b_2^j) \geq n^i - L_1^i$ *and* $b_2^{j*}(b_2^i) \geq n^j - L_1^j$.

Proof

$$(L_1^i + n^i f_2^i/f_1^i - f_2^i/f_1^i)R_1^{i*} \geq (L_1^i + n^i f_2^i/f_1^i - f_2^i/f_1^i)R_2^{i*}$$

$$n^i - L_1^i + n^i \frac{f_2^i}{f_1^i}\frac{R_1^{i*}}{R_2^{i*}} - n^i \frac{f_2^i}{f_1^i} - \frac{f_2^i}{f_1^i}\frac{R_1^{i*}}{R_2^{i*}} L_1^i + \frac{f_2^i}{f_1^i} L_1^i \geq n^i - L_1^i \frac{R_1^{i*}}{R_2^{i*}}$$

$$\frac{n^i - L_1^i \frac{R_1^{i*}}{R_2^{i*}}}{1 + \frac{f_2^i}{f_1^i}(\frac{R_1^{i*}}{R_2^{i*}} - 1)} \leq n^i - L_1^i$$

$$b_2^{i*}(b_2^j) \geq n^i - L_1^i$$

Similarly, the best response of airline j given b_2^i satisfies: $b_2^{j*}(b_2^i) \geq n^j - L_1^j$. •

This is intuitive because airline i is guaranteed to receive at least L_1^i requests of class-1 (regardless of b_2^j) and any reasonable booking limit policy should protect at least L_1^i units for class-1.

4.4.3 Existence of Nash Equilibrium

We now investigate the properties of the seat inventory control game with two airlines that are interested in maximizing CR using limited demand information. The first observation on the Nash Equilibrium follows immediately from existing results on game theory.

Proposition 13 *A mixed-strategy Nash Equilibrium in this seat inventory control game always exists.*

Proof For any game with multiple players, if the strategy set of each player is a compact set of the Euclidean space and each player's payoff function is continuous, then a mixed-strategy Nash Equilibrium exists (Theorem 2.9 of Vives 1999). Evidently, the strategy set of the airlines in this problem are $[0, n^i]$ and $[0, n^j]$, which are compact (closed and bounded).

Let $\alpha(b_2^i) = \frac{R^i(I^i, b_2^i | I^j, b_2^j)}{R^{*i}(I^i, I^j, b_2^j)}$, be the payoff of airline i as a function of booking limit (strategy) b_2^i. Because the inputs are LBH, the payoff can be expressed as:

$$\alpha(b_2^i) = \frac{1}{R^{*i}(I^i, I^j, b_2^j)} \{ f_2^i \min(b_2^i, AD_2^i(I^i, I^j, b_2^j)) + f_1^i \min(n^i - \min(b_2^i, AD_2^i(I^i, I^j, b_2^j)) \}$$

for arbitrary inputs I^i and I^j. Notice that, $\alpha(b_2^i)$ is a piecewise continuous function of b_2^i. This completes the proof. $\bullet$

While the result above guarantees Nash Equilibrium, it may be achieved by mixed-strategy or it may not be unique. First, we state the following observation and then we will investigate the pure strategy Nash Equilibrium exists and the uniqueness.

Next, we provide the following theorem on the pure-strategy Nash Equilibrium.

Theorem 1 *A pure-strategy Nash Equilibrium always exists in this seat inventory control game.*

Proof According to equation 4.8 and 4.9, we analyze how the optimal solution to the LP in (4.2) to (4.6) changes when demand bounds change.

Case A: If L_1^i is increased by an amount $\theta > 0$ (other bounds remains the same), then R_2^{i*} increases and $(R_1^{i*} - R_2^{i*})/f_1^i$ decreases. $(R_1^{i*} - R_2^{i*})/f_1^i$ decreases by less than θ because R_1^{i*} is unchanged and change in R_2^{i*} is less than $f_1\theta$. $(R_1^{i*} -$

$R_2^{i*})/f_1^i * z^{i*}$ decreases by less than θ as $z^{i*} \leq 1$. So if L_1^i is increased by θ in the optimal solution to the LP, $x_1^i(b_2^j)$ increases and, $x_2^i(b_2^j)$ decreases.

Case B: If U_1^i is increased by θ (other bounds remains the same), then R_1^{i*} increases and R_2^{i*} unchanged. Then $x_2^i(b_2^j)$ decreases.

Case C: If U_2^i is decreased by θ (other bounds remains the same), then R_1^{i*} and R_2^{i*} both can decrease. If part of θ is originally counted in the R_2^{i*} but not counted in R_1^{i*}, then $\frac{R_1^{i*}-X}{R_2^{i*}-Y} \geq \frac{R_1^{i*}}{R_2^{i*}}$ as long as scalars $X \leq Y$. So $\frac{R_1^{i*}}{R_2^{i*}}$ increases and in the optimal solution to the LP, $x_2^i(b_2^j)$ decreases.

Next, we consider the change in the best response functions. Let $\Delta^{*i} = b_2^{i*}(b_2^j + \Delta^j) - b_2^{i*}(b_2^j)$ be the change in the best response of i given a $\Delta^j \geq 0$ increase in the booking limit of airline j. Below we analyze the properties of Δ^{*i}. Notice that when b_2^j increases by Δ^j, the overflow from j to i in class-2 can decrease or remain unchanged, and the overflow in class-1 can increase or remain unchanged. We next investigate the overflow and aggregate demand observed by i under extreme input pairs $\{ES_2^i, ES_2^j\}$ and$\{ES_1^i, ES_1^j\}$.

Case 1: If $AD_2^i(U_2^i, U_2^j, b_2^j) = AD_2^i(U_2^i, U_2^j, b_2^j + \Delta^j)$, $AD_1^i(U_1^i, U_1^j, b_2^j) = AD_1^i(U_1^i, U_1^j, b_2^j + \Delta^j)$, $AD_1^i(L_1^i, L_1^j, b_2^j) = AD_1^i(L_1^i, L_1^j, b_2^j + \Delta^j)$, then airline i's best response is unchanged: $\Delta^{*i} = 0$.

Case 2: If $AD_2^i(U_2^i, U_2^j, b_2^j) > AD_2^i(U_2^i, U_2^j, b_2^j + \Delta^j)$, $AD_1^i(U_1^i, U_1^j, b_2^j) = AD_1^i(U_1^i, U_1^j, b_2^j + \Delta^j)$, $AD_1^i(L_1^i, L_1^j, b_2^j) = AD_1^i(L_1^i, L_1^j, b_2^j + \Delta^j)$. This corresponds to Case C above. Therefore, $\Delta^{*i} < 0$.

Case 3: If $AD_2^i(U_2^i, U_2^j, b_2^j) = AD_2^i(U_2^i, U_2^j, b_2^j + \Delta^j)$, $AD_1^i(U_1^i, U_1^j, b_2^j) < AD_1^i(U_1^i, U_1^j, b_2^j + \Delta^j)$, $AD_1^i(L_1^i, L_1^j, b_2^j) = AD_1^i(L_1^i, L_1^j, b_2^j + \Delta^j)$. This corresponds to Case B above.

Therefore, $\Delta^{*i} < 0$.

Case 4: If $AD_2^i(U_2^i, U_2^j, b_2^j) = AD_2^i(U_2^i, U_2^j, b_2^j + \Delta^j)$, $AD_1^i(U_1^i, U_1^j, b_2^j) < AD_1^i(U_1^i, U_1^j, b_2^j + \Delta^j)$, $AD_1^i(L_1^i, L_1^j, b_2^j) < AD_1^i(L_1^i, L_1^j, b_2^j + \Delta^j)$. This requires Cases B and A. Therefore, $\Delta^{*i} < 0$.

Case 5: If $AD_2^i(U_2^i, U_2^j, b_2^j) > AD_2^i(U_2^i, U_2^j, b_2^j + \Delta^j)$, $AD_1^i(U_1^i, U_1^j, b_2^j) < AD_1^i(U_1^i, U_1^j, b_2^j + \Delta^j)$, $AD_1^i(L_1^i, L_1^j, b_2^j) = AD_1^i(L_1^i, L_1^j, b_2^j + \Delta^j)$. This requires Cases C and B. Therefore, $\Delta^{*i} < 0$.

Case 6: If $AD_2^i(U_2^i, U_2^j, b_2^j) > AD_2^i(U_2^i, U_2^j, b_2^j + \Delta^j)$, $AD_1^i(U_1^i, U_1^j, b_2^j) < AD_1^i(U_1^i, U_1^j, b_2^j + \Delta^j)$, $AD_1^i(L_1^i, L_1^j, b_2^j) < AD_1^i(L_1^i, L_1^j, b_2^j + \Delta^j)$. This requires Cases A, B and C. Therefore, $\Delta^{*i} < 0$.

This prove the best response b_2^{i*} is non-increasing in b_2^j. In duopoly games, non-increasing best response function guarantees existence of Nash Equilibrium (see Vives,1999). This completes the proof. •

Theorem 2 *A Pure-strategy Nash Equilibrium in this seat inventory control game is unique.*

Proof

We show this by contradiction. First, assume $(b^{i*}, b^{j*}), (b^{i*} + \theta_i, b^{j*} - \theta_j)$ are two distinct best response pairs. Considering airline i's problem, these two pairs should give $z^{i*}(b^{i*}, b^{j*}) = z^{i*}(b^{i*} + \theta_i, b^{j*} - \theta_j)$ by definition. According to Equation 4.9 and 4.7:

$$b^{i*} = \frac{1}{f_2^i}(R_2^{i*}(b^{j*}) z^{i*}(b^{i*}, b^{j*}) - f_1^i L_1^i)$$

$$b^{i*} + \theta_i = \frac{1}{f_2^i}(R_2^{i*}(b^{j*} - \theta_j) * z^{i*}(b^{i*} + \theta_i, b^{j*} - \theta_j) - f_1^i L_1^i)$$

$$\theta_i f_2^i = z^{i*}(R_2^{i*}(b^{j*} - \theta_j) - R_2^{i*}(b^{j*}))$$

Thus, $\theta_i f_2^i < R_2^{i*}(b^{j*} - \theta_j) - R_2^{i*}(b^{j*})$.

We can show:

$$R_2^{i*}(b^{j*} - \theta_j) - R_2^{i*}(b^{j*}) \leq \theta_j f_2^i$$

which implies

$$\theta_j > \theta_i$$

.

Similarly, if we focus on airline j's problem, based on the same argument, we will get the contradictory result as $\theta_i > \theta_j$. Hence, there's only a unique pure-strategy Nash equilibrium in the seat allocation game.

•

And because of the existence of a unique pure-strategy Nash Equilibrium and because of the simplicity of the best response function, a simple iterative search process can be constructed to find this unique equilibrium very quickly.

4.5 Analysis of Centralized Decision Making

Consider the above model with the same carrier that has back-to-back flights on the same route. There is one decision-maker who will set the booking limits for both flights. The question is how to choose the booking limits on both flights to maximize the revenues on that route. Mathematically, for such centralized decision-making, it means the following optimization problem.

$$(CEN) \qquad \max_{(b_2^i, b_2^j)} z^{CEN} \leq \frac{R^i(I^i, b_2^i|I^j, b_2^j) + R^j(I^j, b_2^j|I^i, b_2^i)}{R^*(I^i, I^j)} \tag{4.10}$$

$$\forall I^i \in \Omega(L^i, U^i), \forall I^j \in \Omega(L^j, U^j) \tag{4.11}$$

where the numerator is the sum of the conditional online revenue and the denominator is the offline revenue.

We can define a single-leg problem by aggregating capacity, $\tilde{n} = n^i + n^j$. The input sequences for this single-leg problem is also LBH and number of requests in each fare class should be within the range of $\tilde{L_k} = L_k^i + L_k^j$ and $\tilde{U_k} = U_k^i + U_k^j$, $k = 1, 2$. The objective of this single-leg problem can be expressed as:

$$(SIN) \qquad \max_{\tilde{b}} z^{SIN} \leq \frac{R^{SIN}(\tilde{I}, \tilde{b})}{R^{*SIN}(\tilde{I})} \tag{4.12}$$

$$\tilde{I} \in \Omega(\tilde{L}, \tilde{U}) \tag{4.13}$$

where R^{*SIN} is the online revenue given input sequence $\tilde{I}$ and booking limit $\tilde{b}$ and R^{*SIN} is the corresponding offline revenue.

Proposition 14 *For any arbitrary feasible* (b_2^i, b_2^j) *for model CEN, we can define* $\tilde{b} = b_2^i + b_2^j$ *that* $\tilde{b}$ *is feasible to model SIN.*

<u>Proof</u> We show that for any feasible solution (b_2^i, b_2^j) to model CEN with ratio $\frac{R(I^i, b_2^i|I^j, b_2^j) + R(I^j, b_2^j|I^i, b_2^i)}{R^*(I^i, I^j)}$, the feasible solution $\tilde{b_2} = b_2^i + b_2^j$ of model SIN has $\frac{R^{SIN}(\cdot)}{R^{*SIN}(\cdot)} = \frac{R^i(\cdot) + R^j(\cdot)}{R^*(\cdot)}$.

We first show $R^{*SIN} = R^*$. The offline revenue $R^*(I^i, I^j)$ is computed based on the aggregated capacity and the number of requests of each fare class in I^i,I^j. Hence, $R^*(I^i, I^j)$ in model CEN will be equal to $R^*(I^i) + R^*(I^j)$, which equals to R^{*SIN}.

Then, we will show the equivalence between $R^i(I^i, b_2^i | I^j, b_2^j) + R^j(I^j, b_2^j | I^i, b_2^i)$ and $R^{SIN}(\tilde{I}, \tilde{b})$.

As we assumed LBH arrivals, the following relationship holds: $\tilde{I}[2] = I^i[2] + I^j[2]$, $\tilde{I}[1] = I^i[1] + I^j[1]$. The number of class-2 requests accepted in model SIN is $\min(\tilde{b}, \tilde{I}[2])$. In model CEN, the unaccommodated requests of one flight can overflow to the other, and it will get accepted by the other flight as long as there's availability. So the number of class-2 requests accepted in model CEN is $\min(b_2^i, I^i[2] + [I^j[2] - b_2^j]^+)$ on flight i and $\min(b_2^j, I^j[2] + [I^i[2] - b_2^i]^+)$ on flight j.

We compare online revenues from class-2 requests in the following cases.

<u>Case 1:</u> $\tilde{I}[2] \leq b_2^i + b_2^j$. The online revenue from class-2 requests in model SIN is the number of accepted multiply the fare, which is $\tilde{I}[2] * f_2$. For model CEN, because $I^i[2] + I^j[2] \leq b_2^i + b_2^j$, all class-2 fare requests can get their reservation on one of the flight, which yields the total of $(I^i[2] + I^j[2]) * f_2$ as the sum of conditional revenue.

<u>Case 2:</u> $\tilde{I}[2] > b_2^i + b_2^j$. For model SIN, in this case, the online revenue from class-2 request is $(b_2^i + b_2^j) * f_2$. In model CEN, given b_2^i and b_2^j, the sum of conditional online revenues is $(b_2^i + b_2^j) * f_2$.

Similarly, we next compare online revenues from class-1 requests in the following cases.

Case 1: $\tilde{I}[1] \leq \tilde{n} - b_2^i - b_2^j$. The online revenue from class-1 requests in model SIN is the number of accepted multiply the fare, which is $\tilde{I}[1] * f_1$. For model CEN, because $I^i[1] + I^j[1] \leq \tilde{n} - b_2^i - b_2^j$, all class-1 fare requests can get their reservation on one of the flight, which yields the total of $(I^i[2] + I^j[2]) * f_1$ as the sum of the conditional revenue.

Case 2: $\tilde{I}[1] > \tilde{n} - b_2^i - b_2^j$. For model SIN, in this case, the online revenue from class-1 requests is $(\tilde{n} - b_2^i - b_2^j) * f_1$. In model CEN, given b_2^i and b_2^j, the sum of conditional online revenues is $(\tilde{n} - b_2^i - b_2^j) * f_1$.

Note also that for any input pair I^i and I^j in CEN, there is a corresponding unique $\tilde{I}$ in SIN. For an arbitrary $\tilde{I}$ in SIN, there could be more than on feasible I^i,I^j pair in CEN. However, if $\tilde{I}$ is the worst-case input for SIN, one of the corresponding feasible inputs in CEN is the worst- case input. If I^i and I^j are worst-case for CEN, then $\tilde{I} = I^i + I^j$ is the worst case input for SIN. This completes the proof.

•

Proposition 15 *The optimal solution to the centralized model is as follows:*

$$\tilde{b} = \frac{1}{f_2}\left[\frac{(\tilde{n} - L_1^i - L_1^j)f_2 + (L_1^i + L_1^j)f_1}{1 + f_2/f_1(\tilde{R}_1^*/\tilde{R}_2^* - 1)} - f_1(L_1^i + L_1^j)\right] \tag{4.14}$$

where the offline revenues $\tilde{R}_1^* = R^*(\tilde{U_2}, \tilde{U_1})$ *and* $\tilde{R}_2^* = R^*(\tilde{U_2}, \tilde{L_1})$

Proof Same as single-leg model in Lan(2008) with parameter $\tilde{n}$ and $\Omega(\tilde{L}, \tilde{U})$. See equation 2.14 in Chapter 2. •

4.6 Decentralized versus Centralized Solutions

We now compare the solution in the decentralized model versus the centralized model in the following proposition.

Proposition 16 *In a symmetric case when two airlines (two flights) are equivalent in fares, capacity, and market share (i.e. demand bounds), the sum of booking limits of class-2 in the Nash Equilibrium solution of the decentralized solution is no more than $\tilde{b}$ of equation 4.14.*

Proof We use $n, U_k, L_k, k = 1, 2$ to denote the capacity and demand bounds for the single airline(flight). In this symmetric case, the booking limit in problem CEN in equation 4.14 equals to:

$$\tilde{b} = \frac{1}{f_2}\left[\frac{(2n - 2L_1)f_2 + 2f_1L_1}{1 + f_2/f_1(\tilde{R}_1^*/\tilde{R}_2^* - 1)} - 2f_1L_1\right]$$

where $\tilde{R}_1^* = R^*(2U_2, 2U_1)$ and $\tilde{R}_2^* = R^*(2U_2, 2L_1)$.

Because of symmetry, we can denote input pair as $\{ES_1, ES_1\}$ and $\{ES_2, ES_2\}$. We use $\dot{b}$ to denote the mutual best response for both airline i and j. Based on equation 4.9, we derive the decentralized solution:

$$\dot{b} = b_2^{i*} = b_2^{j*} = \frac{1}{f_2}\left[\frac{(n - L_1)f_2 + f_1L_1}{1 + f_2/f_1(R_1^*/R_2^* - 1)} - f_1L_1\right] \tag{4.15}$$

where $R_1^* = R^*(U_2 + [U_2 - \dot{b}]^+, U_1 + [U_1 - n + \dot{b}]^+)$ and $R_2^* = R^*(U_2 + [U_2 - \dot{b}]^+, L_1)$.

Case 1: $U_2 + U_1 \leq n$. In this case, $\dot{b} = n$. $R_1^* = U_1f_1 + U_2f_2$. $R_2^* = L_1f_1 + U_2f_2$. $\tilde{R}_1^* = 2U_1f_1 + 2U_2f_2$, $\tilde{R}_2^* = 2L_1f_1 + 2U_2f_2$. $R_1^*/R_2^* = \tilde{R}_1^*/\tilde{R}_2^*$, so $2\dot{b} = \tilde{b}$.

Case 2: $U_2 + U_1 > n$ and $U_2 + L_1 \geq n$. In this case, $R_1^* = (U_1 + [U_1 - n + \dot{b}]^+)f_1 + (n - U_1 - [U_1 - n + \dot{b}]^+)f_2$. $R_2^* = L_1f_1 + (n - L_1)f_2$. $\tilde{R}_1^* = 2U_1f_1 + (2n - 2U_1)f_2$,

$\tilde{R}_2^* = 2L_1 f_1 + (2n - 2L_1) f_2$. $R_1^*/R_2^* \geq \tilde{R}_1^*/\tilde{R}_2^*$, $2\dot{b} \leq \tilde{b}$

<u>Case 3:</u> $U_2 + U_1 > n$ and $U_2 + L_1 < n$. In this case, $n - L_1 \geq U_2 + [U_2 - \dot{b}]^+$ due to symmetry. Because otherwise for input pair $\{ES_2, ES_2\}$, both airlines will get L_1 of class-1 requests and $(n - L_1)$ of class-2 requests, which will be contradictory to $U_2 + L_1 < n$. Thus, $R_1^* = (U_1 + [U_1 - n + \dot{b}]^+) f_1 + (n - U_1 - [U_1 - n + \dot{b}]^+) f_2$ and $R_2^* = L_1 f_1 + (U_2 + [U_2 - \dot{b}]^+) f_2$. $\tilde{R}_1^* = 2U_1 f_1 + (2n - 2U_1) f_2$, $\tilde{R}_2^* = 2L_1 f_1 + 2U_1 f_2$.

We rearrange equation 4.15 in the following way:

$$\dot{b}(R_1^* f_2 - R_2^* f_2 + R_2^* f_1) = R_2^* n f_1 - R_1^* L_1 f_1$$

$$\frac{R_1^*}{R_2^*} = \frac{\dot{b} f_2 + (n - \dot{b}) f_1}{\dot{b} f_2 + L_1 f_1}$$

If $[U_1 - n + \dot{b}]^+ = 0$, $\tilde{R}_1^* = 2R_1^*$ and $\tilde{R}_2^* = 2R_2^*$. Otherwise, $\tilde{R}_1^* = R_1^* + \dot{b} f_2 + (n - \dot{b}) f_1$, $\tilde{R}_2^* = R_2^* + \dot{b} f_2 + L_1 f_1$. $\tilde{R}_1^*/\tilde{R}_2^* = R_1^*/R_2^*$. Hence, in this case,

$$2\dot{b} = \tilde{b}$$

. •

This result shows that competition favors class-1, i.e. the number of seats sold at the higher fare is more. In other words, competition actually increases the (average) price for the good/service. This is in contrast with the findings in the economics literature, where prices go down with competition and competition is good for the consumers. The reason for our result is as follows: In our model, both airlines(flights) are equivalent in their price structure, and the prices are exogenous. Airlines only compete on quantity (seat inventory) but not on price. If the airlines were in price <u>and</u> quantity competition, the price for class-1 may decrease, even though airlines allocate more seats to class-1, leading to an overall decrease in the average price per unit sold. Another way to think of the problem under competition is that, the centralized problem represents a monopolist, and a monopolist can increase the prices for the route. Since we are limiting our decision to only fare-class allocation, the pricing effects are not captured. Therefore, the result on competition leading to higher prices in the market should be interpreted within the realm of our model.

In asymmetric cases, the actual difference between booking limits of decentralized and centralized models is parameter-specific. For example, suppose $f_1 = 3$, $f_2 = 2$. Airline i has capacity 10 with $U_2^i = 8$, $U_1^i = 8$, $L_1^i = 1$, and airline j has capacity 12 with $U_2^j = 10$, $U_1^j = 10$, $L_1^j = 0$. Solution to the decentralized model is $b_2^i = 6.666$, $b_2^j = 8.620$. The centralized solution $\tilde{b}$ is 14.650. Therefore, in this case, $\tilde{b} < b_2^i + b_2^j$. If fares are $f_1 = 10$, $f_2 = 1$ and $L_1^j = 1$ instead of 0, the decentralized solution is $b_2^i = 3.321$,$b_2^j = 4.559$, and the centralized model yields $\tilde{b} = 8.897$. In this case, $\tilde{b} > b_2^i + b_2^j$. Hence, in an asymmetric game, the total number of seats protected for class-1 may increase or decrease with competition; the actual outcome depends on the capacity, bounds and fares.

4.7 Numerical Examples

In this section, we use several examples with two flights to illustrate the differences between the centralized and decentralized problems.

Example 4.1 In this example,$n^i = 80$,$n^j = 100$. Fares are $f_1 = 2$, $f_2 = 1$. Demand bounds are: $U_2^i = U_2^j = 80$, $U_1^i = U_1^j = 60$, $L_2^i = L_2^j = 20$, $L_1^i = L_1^j = 40$.

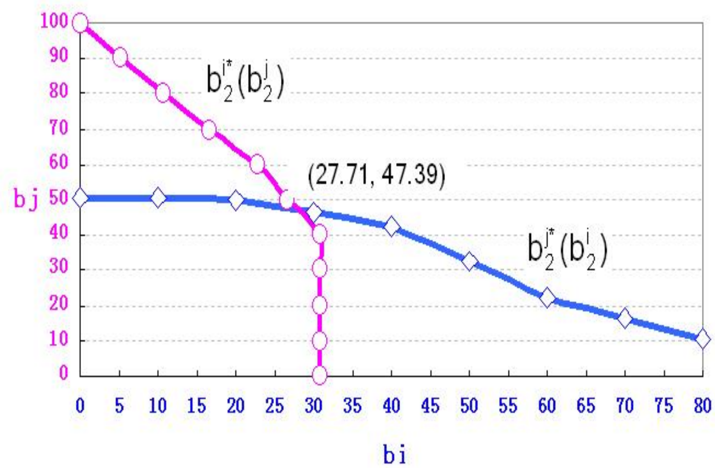

Figure 4.2: Airlines' Best Response Function in Example 4.1

Figure 4.2 plots the best response functions $b_2^{i*}(b_2^j)$ and $b_2^{j*}(b_2^i)$ accordingly. Observe that when one airline increases her booking limit b_2^i, the best response of the other airline is to decrease his booking limit b_2^j. This game has a pure-strategy Nash Equilibrium at booking limits $b_2^{i*} = 27.71$ and $b_2^{j*} = 47.39$, shown as the crossing point in the figure. Compared to this decentralized solution, the centralized one computed for this example is $\tilde{b} = 125.21$.

Example 4.2 In this example, we consider two symmetric airlines with capacity of 100 each. For both airlines, class-1 demand lies between [40,60] and class-2 demand lies between [20,80]. Both class-1 and class-2 demand are uniformly distributed and demand arrives in a LBH sequence. The demand factor is 1.0 and there are 6000 instances for each experiment.

First, we investigate the effect of fare ratios. Figure 4.3 illustrates the difference of booking limits between centralized versus decentralized solutions. The x-axis is the fare ratio $f_2 : f_1$, varying from $1 : 1$ to $1 : 3$. The y-axis the difference $\tilde{b} - (b_2^{i*} + b_2^{j*})$.

Note that the difference starts at 0 when fare ratio is $1 : 1$, rapidly increases and reaches the maximum around $1 : 1.8$, and then gradually decreases. This is because when fare ratio is $1 : 1$, the booking limits equals the capacity in both centralized and decentralized model. As f_1 starts to increase, the potential overflow in class-1 gives incentives for competing airlines to book less in class-2 requests in the competitive situation compared to the centralized model. When f_1 continues to increases, centralized model also reduces the booking limits for class-2. Hence, the discrepancy between centralized and decentralized solution becomes smaller when f_1 is high.

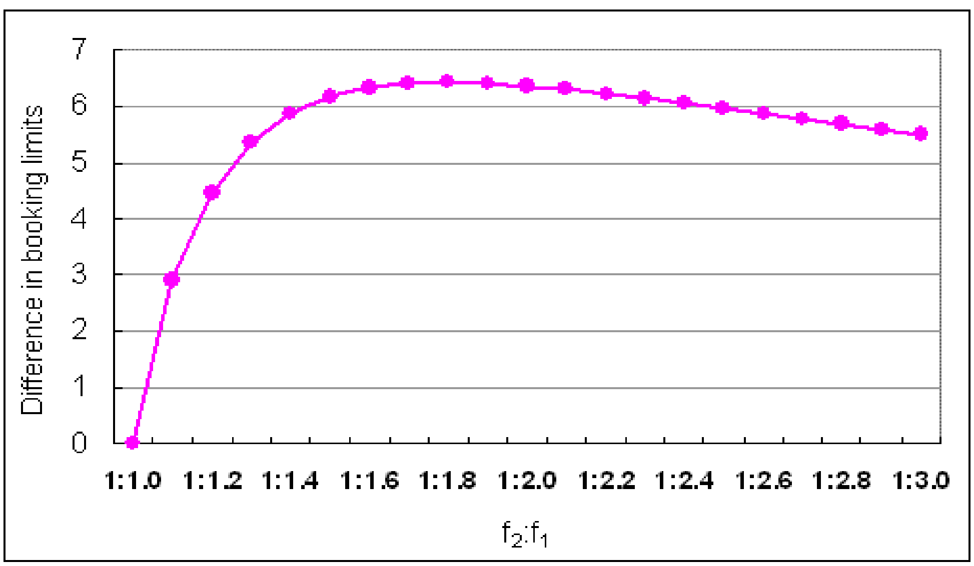

Figure 4.3: Difference of booking limits in Example 4.2

In the same example, we report the revenues and booking limits of the airlines

	$f_1 = 1.2,\ f_2 = 1$		$f_1 = 2,\ f_2 = 1$		$f_1 = 5,\ f_2 = 1$	
	cent	decent	cent	decent	cent	decent
b_2^i	-	54.53	-	47.48	-	42.89
b_2^j	-	54.53	-	47.48	-	42.89
Total booking limits of class-2	113.53	109.06	101.55	94.96	89.95	85.78
Avg Revenue	207.43	207.51	283.89	283.08	578.25	576.30
% Diff in Avg Rev	-0.038		0.286		0.338	

Table 4.1: Effect of fare ratio in Example 4.2

for a few instances. Table 4.7 summaries the results for both centralized and decentralized solutions. Notice that the models we investigate come with worst-case objectives. Therefore, centralized decision-making need not improve the average revenues over decentralized based on the results in Table 4.7. For the same instances, Figures 4.4, 4.5 and 4.6 plot the distribution of revenues. The total of 6000 runs are split into 30 samples of 200 each to come up with percentile estimates. We observe that when fare ratio is very small or relatively large, the performance of centralized and decentralized solution are very close. When fare ratio is moderate, they have similar performance in the worst case. But the decentralized solution has better performance in 90^{th} percentile. It yields much higher revenues because the competition makes airlines more aggressive in protecting seats for class-1 and being aggressive leads to higher chances of making higher revenues. Compared to a centralized solution, competition increases availability to class-1 and increases variance of revenues without compromising the worst-case.

We next consider an example with two asymmetric airlines in the following

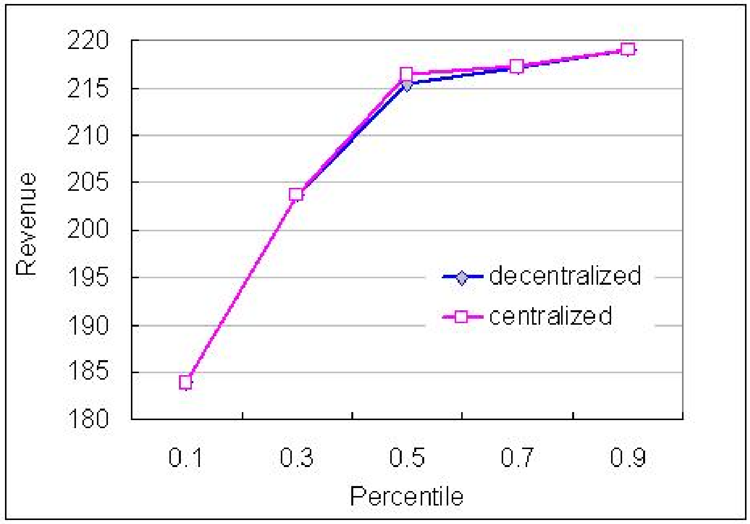

Figure 4.4: Distribution of Revenues in Example 4.2 with fare ratio 1:1.2

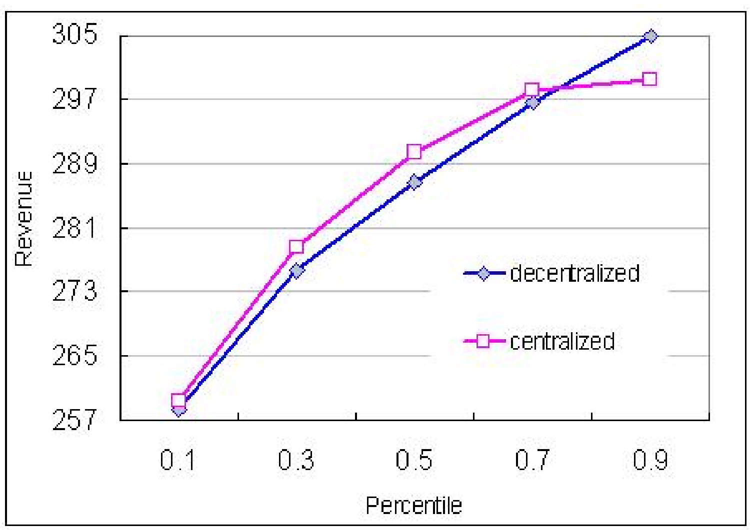

Figure 4.5: Distribution of Revenues in Example 4.2 with fare ratio 1:2

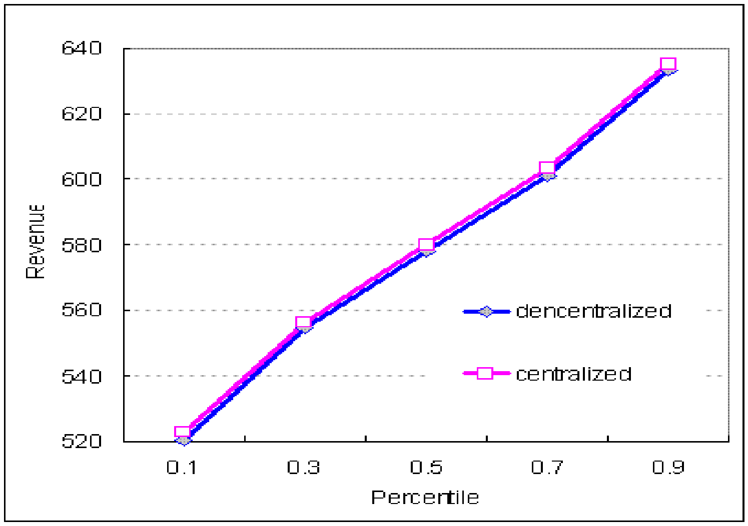

Figure 4.6: Distribution of Revenues in Example 4.2 with fare ratio 1:5

example.

Example 4.3: We still use the same parameters of Example 4.2 . But in this example, these two airlines differ in terms of their capacity , $n^i = 80$ and $n^j = 100$. Table 4.7 shows the corresponding booking limits and average performance.

	$f_1 = 1.2,\ f_2 = 1$		$f_1 = 2,\ f_2 = 1$		$f_1 = 5,\ f_2 = 1$	
	cent	decent	cent	decent	cent	decent
b_2^i	-	34.60	-	27.71	-	23.07
b_2^j	-	54.52	-	47.39	-	42.76
Total booking limits of class-2	93.55	89.12	81.43	75.10	69.92	65.83
Avg Revenue	194.07	194.21	270.13	268.75	562.24	559.76
% Diff in Avg Rev	-0.071		0.516		0.443	

Table 4.2: Effect of fare ratio in Example 4.2

For this asymmetric example, the observation is similar to the symmetric case. When fare ratio is relatively small or large, the performances of centralized and

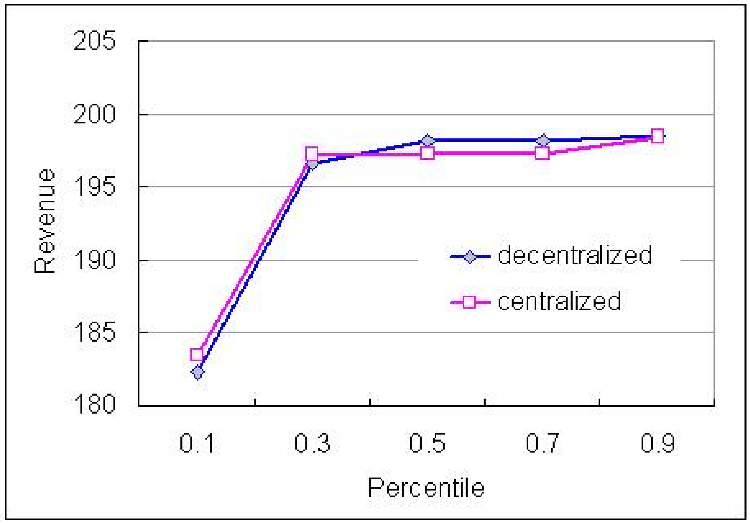

Figure 4.7: Distribution of Revenues in Example 4.3 with fare ratio 1:1.2

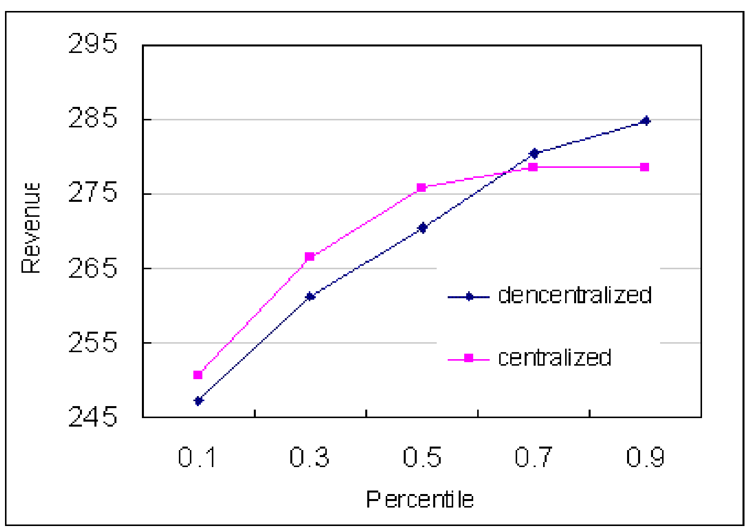

Figure 4.8: Distribution of Revenues in Example 4.3 with fare ratio 1:2

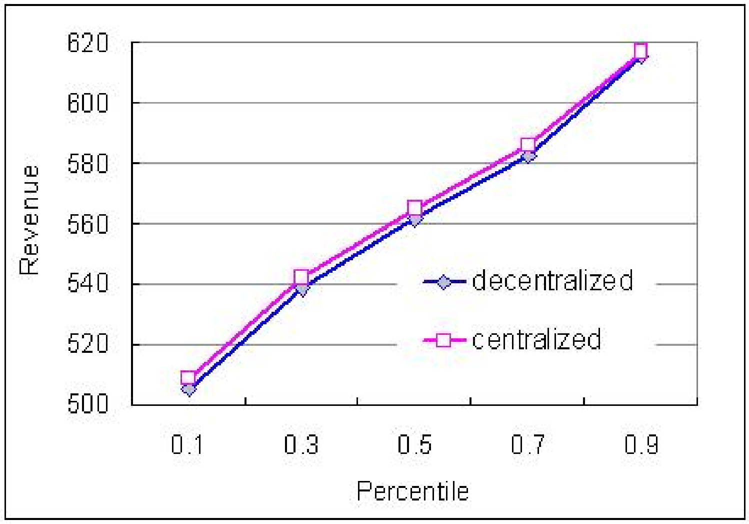

Figure 4.9: Distribution of Revenues in Example 4.3 with fare ratio 1:5

decentralized solution are close. However, when fare ratio is moderate, centralized performs better than decentralized solution in the 10^{th} percentile, while there's the opposite relationship in the 90^{th} percentile. Competition certainly makes airlines more aggressive, regardless of its size and market power, in protecting seats for class-1 requests.

4.8 Extensions

The immediate modeling extensions of our work include increasing the number of flights (airlines) and increasing the number of fare classes.

4.8.1 Multi-fare problem

Up to now, we provide analysis for both centralized and decentralized decision-making with $m = 2$. We provide both solutions and show the existence and uniqueness of Nash Equilibrium for two competing airlines(flights). In this section, we will consider the possible extension to the multi-fare problem with two airlines.

Because the strategy space of the airlines $[0, n^i]$ and $[0, n^j]$ are compact and objective function is continuous, a Nash Equilibrium in mixed strategy always exists. Supermodularity is the sufficient condition to show the existence of a pure-strategy Nash Equilibrium for a multi-fare problem. This means the objective function should have increasing differences with respect to the nested booking limits. We will use b_3^i, b_3^j to denote the nested booking limits for class-3 requests and b_2^i, b_2^j to represent the nested booking limits for class-3 and class-2.

However, the next two examples show "increasing difference" assumption may be violated when $m = 3$. Therefore, supermodularity does not necessarily hold when $m > 2$.

Example 4.4: This is a symmetric case when two competing airlines(flights) have the same capacity 9 and demand bounds: $U_1 = 4, L_1 = 1$, $U_2 = 5, L_2 = 2$, $U_3 = 6, L_3 = 3$. Fares are $f_1 = 3, f_2 = 2, f_3 = 1$. We can enumerate airline j's choices b_3^j, b_2^j to compute the objective of airline i's best response. We use function $F(\cdot)$ to denote the optimal objective function value in airline i's optimization problem given b_2^j, b_3^j. We compute the following differences:

$$F(b_3^j + \epsilon, b_2^j) - F(b_3^j, b_2^j) = F(1, 3) - F(0, 3) = 0.764706 - 0.772277 = -0.00757$$

$$F(b_3^j + \epsilon, b_2^j + \epsilon) - F(b_3^j, b_2^j + \epsilon) = F(1, 4) - F(0, 4) = 0.772277 - 0.78 = -0.00772$$

In this example,

$$F(b_3^j+\epsilon, b_2^j+\epsilon) - F(b_3^j, b_2^j+\epsilon) < F(b_3^j+\epsilon, b_2^j) - F(b_3^j, b_2^j)$$

that is, the function does not have increasing difference.

Example 4.5: We change the capacity in Example 4.3, $n^i = 8$, $n^j = 10$ which gives us an asymmetric problem. We calculate the differences in $F(\cdot)$ for the same booking limit parameters:

$$F(b_3^j+\epsilon, b_2^j) - F(b_3^j, b_2^j) = F(1,3) - F(0,3) = 0.76364 - 0.77064 = -0.00700$$

$$F(b_3^j+\epsilon, b_2^j+\epsilon) - F(b_3^j, b_2^j+\epsilon) = F(1,4) - F(0,4) = 0.77064 - 0.77777 = -0.00713$$

In this example :

$$F(b_3^j+\epsilon, b_2^j+\epsilon) - F(b_3^j, b_2^j+\epsilon) < F(b_3^j+\epsilon, b_2^j) - F(b_3^j, b_2^j),$$

again violating increasing differences.

Therefore, we know that multi-fare problem is not a supermodular game. Note that supermodularity is the sufficient but not necessary condition for the existence of a pure-strategy Nash Equilibrium. We leave further analysis of this problem to future research. However, our computational experiments with $m > 2$ have some interesting observations. In search of the pure-strategy equilibrium, we found that the best response vectors usually converge to the same mutual best response point even when our iterative procedure starts with different initial points. So we feel that because of nesting property, the multi-fare problem might still have unique pure-strategy equilibrium even though it is not supermodular. These interesting findings are going to be addressed and further explored in the future.

4.8.2 Multi-flight problem with m=2

While our analysis in Section 4.3-4.7 was focused on two airlines, our results can easily be extended to the case where there are more than two flights (airlines), each offering only two fares. If we assume a_k^{ij} is the fraction of overflow in class k from flight i to flight j; with $\sum_{j:j\neq i} a_k^{ij} \leq 1 fork = 1, 2$, the structure of the overflow function is no different then the one we have in our two-flight model. Our analysis carries through and all results are valid if the problem is extended to multiple flights under this overflow assumption.

4.9 Summary of Chapter 4

In this chapter we examined how booking limit decisions are affected by competition when sellers use distribution-free methods and focus on worst-case performance. We have shown that there is a unique pure strategy Nash Equilibrium under this competing game with two classes. In general, we also find that the equilibrium decision of competing airlines for the decentralized model can be very different from the centralized solution. These results can be useful for decision-makers who plan expansion into new markets, who try to defend an entry by a rival, who have problem of data scarcity or inaccuracy, or who just want to remain solvent in the highly competitive market. We also provide additional numerical experiments to show the effect of centralized versus decentralized solutions. Analytically, we showed class-1 (high-fare) customers have more availability of their service under competition than under a monopoly.

Chapter 5

Conclusions and Future Research

In this thesis, we have analyzed the traditional single-leg seat allocation problem from the perspective of competitive analysis of online algorithms. We compare robust policies using both competitive ratio and absolute regret criterion. We also compare robust methods with other bench-mark policies used in practice. Extensive computational examples have shown that robust booking control have significant practical advantages. They provide effective results in many cases; performance is impressive while the need for information is reduced. Furthermore, they are able to hedge against inaccuracies in information and are not prone to forecast error as other well-known policies are. From a research perspective, the competitive analysis of online algorithms approach is very promising in RM.

In this thesis, we provided multiple approaches to including partial demand information within the competitive analysis framework. Included was a multi-period model that considered demand information at a more disaggregated level that is typical in RM. We also extended the basic model with a single decision-maker into a competitive framework with two competing airlines and/or flights. We have demonstrated analytically and evaluated numerically the difference between centralized versus decentralized decision-making.

We have focused on single-leg seat allocation problem from the perspective of competitive analysis of online algorithms so far. However, there can be significant benefits in considering a network approach within this framework. In the airline case, the general/typical problem involves managing the capacities of a set of connecting flights, where the only information available consists of demand bounds on each flight. In the hospitality industry, the problem is managing room capacity on consecutive days, where the arrival distribution is unknown, and a mix of customers with different lengths of stay share the capacity on any given day. Network RM dos creates methodological and operational challenges, but the potential improvements should be sufficient to justify the investment in both research and practice.

Independent demand is another basic assumption in all of our models. This would appear to e a rather simplistic assumption, but it is widely used in practice. A natural extension would be to consider correlated demand because it would be common that if class-1 demand is up then demand for class-2 also rise. Considering demand correlations should provide a better representation of the real business environment. One way to model this is to define a parameter (e.g. demand intensity) so that the demand bounds depend on the level of this parameter.

The analysis in this thesis also ignores one important feature: consumer choice. However, consumers are becoming increasingly aware of the existence of pricing strategies. In addition, the availability of pricing information on the Internet affords consumers the opportunity to behave more strategically when making purchase decisions. There are studies showing that the use of standard yield management approaches to pricing by airlines can result in significantly reduced revenues when buyers are using an informed and strategic approach in purchasing. Therefore, distribution-free models that endogenize strategic behavior will not only fill a void in the academic literature, but would have significant practical value.

Appendix A
Appendices for Chapter 2
A.1 Standard vs. Theft-nesting Implementation of Static Policies

As discussed in Talluri and van Ryzin (2004a), a static nested booking control policy can be implemented in two different ways: The standard-nesting and theft-nesting. 2 In case of m = 2, standard nesting accepts class-2 requests as long as the booking limit b2 is not reached and there is capacity. In contrast, theft-nesting stops accepting class-2 after the first b_2 requests (regardless of their class). Clearly, both implementations yield the same booking curve when the arrivals are LBH. In a way, theft-nesting implementation of a static policy reduces the booking limit of class-2 upon observing a class-1 request, hence it is aggressive in protecting seats for class-1. This is similar to our dynamic policies except that DYN-CR and DYN-AR need not reduce booking limit of class-2 by one unit at each update.

Although we have implemented theft-nesting and tested its performance in our computational experiments, it is not the focus of our attention and we prefer to exclude it from the paper. Our dynamic policies suggest the potential for theft-nesting when demand for higher fare classes is more. On the other hand, theft-nesting is not desirable when the demand is known to be limited: Theft-nesting will close a low fare class early because of high-fare requests that arrive early in the booking horizon. If the total number of high-fare requests is limited, a significant portion of seats protected for future high-fare demand goes unsold and revenues obtained by theft-nesting are lower.

In the Appendix, we provides numerical results on theft-nesting implemented with the robust static policies. For theft-nesting, a booking in class j not only reduces the allocation for class j but also ”steals” from the allocation of all lower classes. So when we accept a request for class j, not only is the class j allocation reduced by one but so are the allocations for classes j+1, j+2,...n. Note that we apply similar idea in implementing theft-nesting Van Ryzin's adaptive algorithm as the booking limits are first determined by previous fill-event.

In practice, demand rarely arrives in low-to-high order and the choice of standard versus theft nesting matters. As we can see from the Example-A in Figure A.1 and Example-B in Figure A.2, with mixed order of arrivals, theft nesting protects more capacity for higher classes (equivalently, allocates less capacity to lower classes). Hence, achieve a higher revenue when there are more high-fare class requests or the fare ratio is small.

Example-A.1: In this two-fare example, fares are f_1=200, f_2=100. Demand are stationary and normally distributed with mean range from [120,0] to [0,120] and time-homogeneous arrivals. Knowing the arrival process.

Example-A.2: In this two-fare example, demand are stationary and normally distributed with mean fixed at [60,60] fare ratio range from 0.1 to 1.0. Arrivals are time-homogeneous.

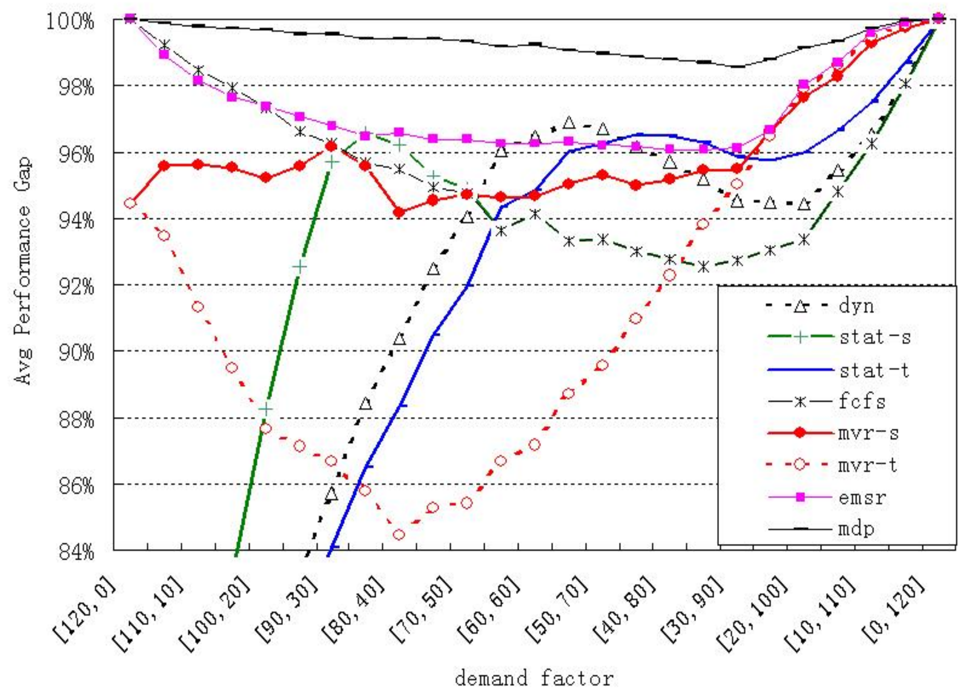

Figure A.1: Average performance gap in Example-A

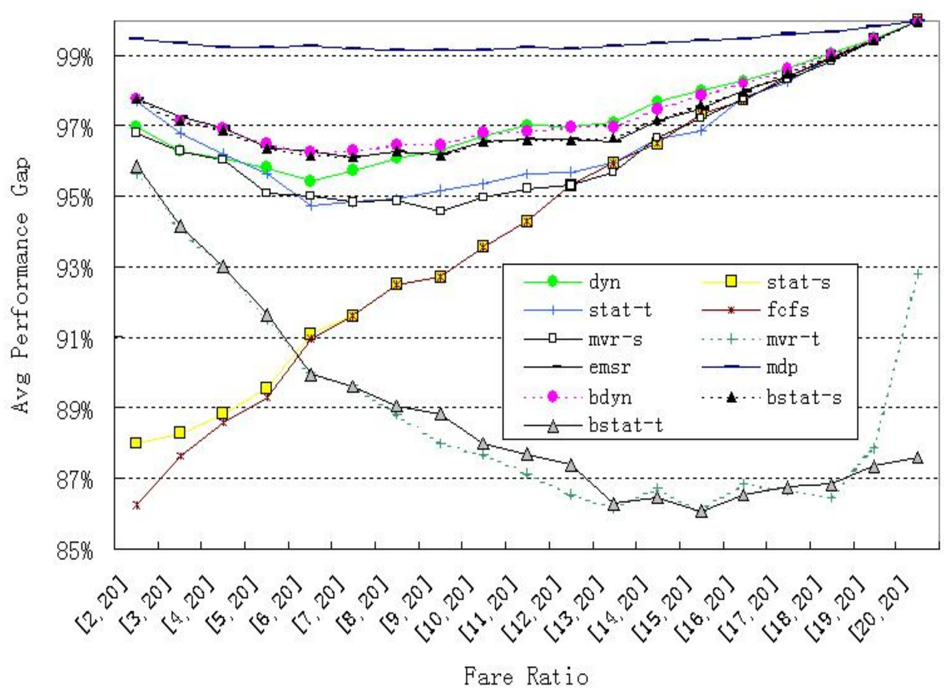

Figure A.2: Average performance gap in Example-B

A.2 Numerical Examples

Example-A.3: In this example, fare ratio range from 0.1 to 1.0 with mean demand fixed at $(\lambda_1, \lambda_2) = (60, 60)$. Low fare class requests arrivals earlier than high fare requests. The average performance gap for the policies are presented in Figure A.3 .

Example-A.4: In this example, fare ratio range from 0.1 to 1.0 with mean demand fixed at $(\lambda_1, \lambda_2) = (40, 80)$. Low fare class requests arrivals earlier than high fare requests. The average performance gap for the policies are presented in Figure A.4 .

Example-A.5: This example is the same as Example-3 in Chapter 2 except the demand-mix: The lower and upper bounds of demand are $L_1 = 20$, $U_1 = 60$, $L_2 = 60$ and $U_2 = 100$, i.e., the mean demand for class-1 is lower. The revenue percentiles for the policies are presented in Figure A.5. In this case, no policy (except OFFLINE) is stochastically dominant. This particular demand-mix affects STAT-AR most: It has a significantly lower revenue at the 10th percentile, and provides the second highest revenue value at the 90th percentile, following BSTAT-AR. This is expected because STAT-AR protects the highest number of seats for class-1, for which the demand is lower in this case. The observations regarding the relative performances of the other policies remain the same.

Example-A.6: This example is the same as Example-3 in Chapter 2 except the demand-mix: The lower and upper bounds of demand are $L_1 = 60$, $U_1 = 100$, $L_2 = 20$ and $U_2 = 60$, i.e., the mean demand for class- 1 is higher. The revenue percentiles for the policies are presented in Figure A.6. In this case, STAT-CR is stochastically dominated because its booking limit for class-2 is significantly higher than the other policies. When the arrivals occur homogeneously over time in both fare classes, dynamic and static policies are not equivalent. If demand-mix is balanced as in Example-3, then the average revenues obtained by each of the policies except DYN-AR are higher when arrivals are time-homogeneous; DYN-AR sets a low booking limit for class-2 initially, and this booking limit is updated with each class-1 request leading to rejecting far too many class-2 requests and having too many idle seats at the end of the booking horizon. However, DYN-AR is not stochastically dominated, i.e., has a higher 90th percentile than, for e.g., BSTAT-AR. These experiments are repeated by changing the arrival regime. Our previous observations regarding stochastic dominance relations among the policies do not change with the arrival regime. Next three examples provide the revenue percentiles for these experiments when arrivals are time-homogeneous.

Example-A.7: This example is the same as Example-3 in Chapter 2 except that the requests in each fare class arrive homogeneously over time. The revenue percentiles for the policies are presented in Figure A.7.

Example-A.8: This example is the same as Example-3b above except that the requests in each fare class arrive homogeneously over time. The revenue percentiles for the policies are presented in Figure A.8.

Example-A.9: This example is the same as Example-3c above except that the requests in each fare class arrive homogeneously over time. The revenue per-

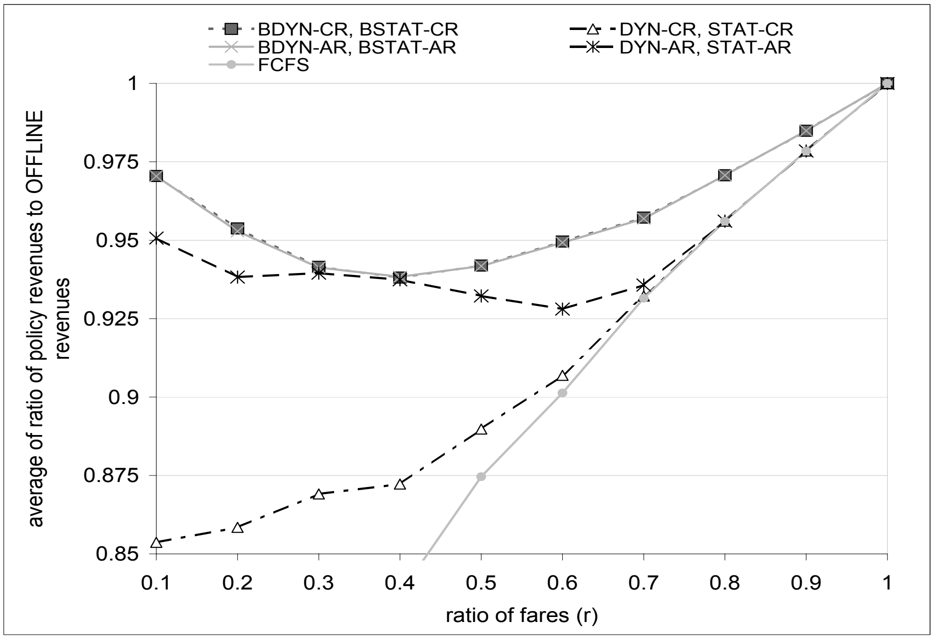

Figure A.3: Average performance gap in Example-2b

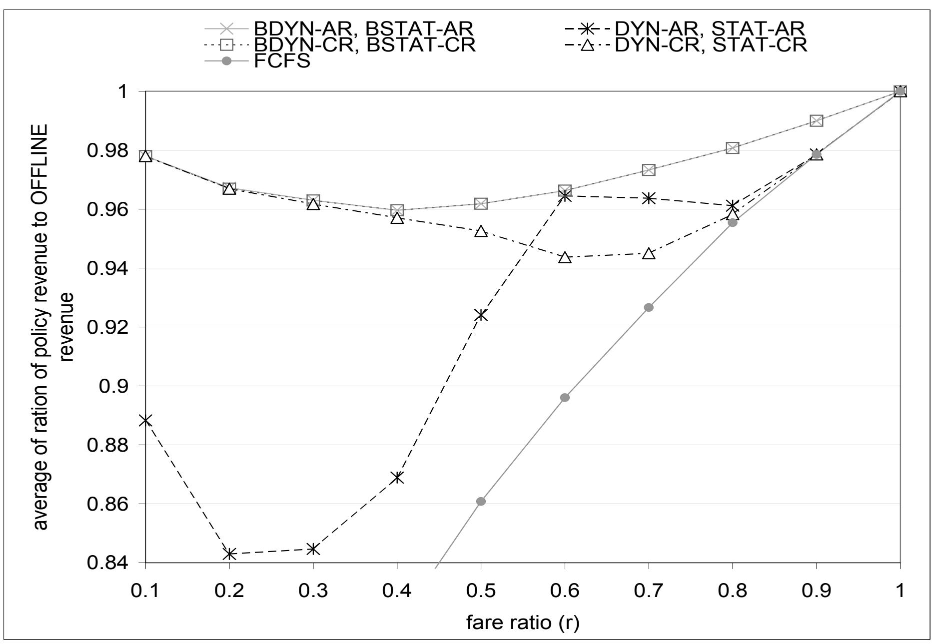

Figure A.4: Average performance gap in Example-2c

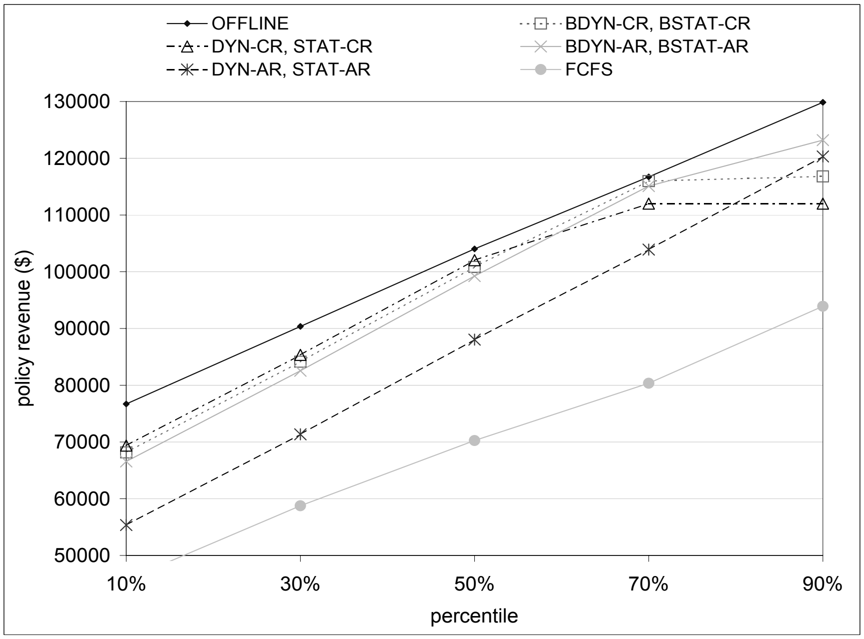

Figure A.5: Average performance gap in Example-3b

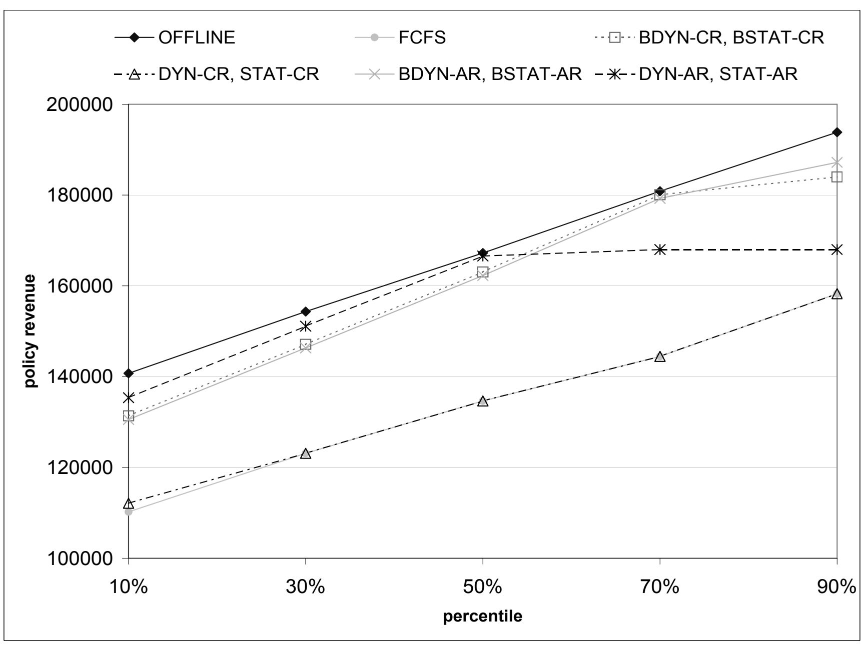

Figure A.6: Average performance gap in Example-3c

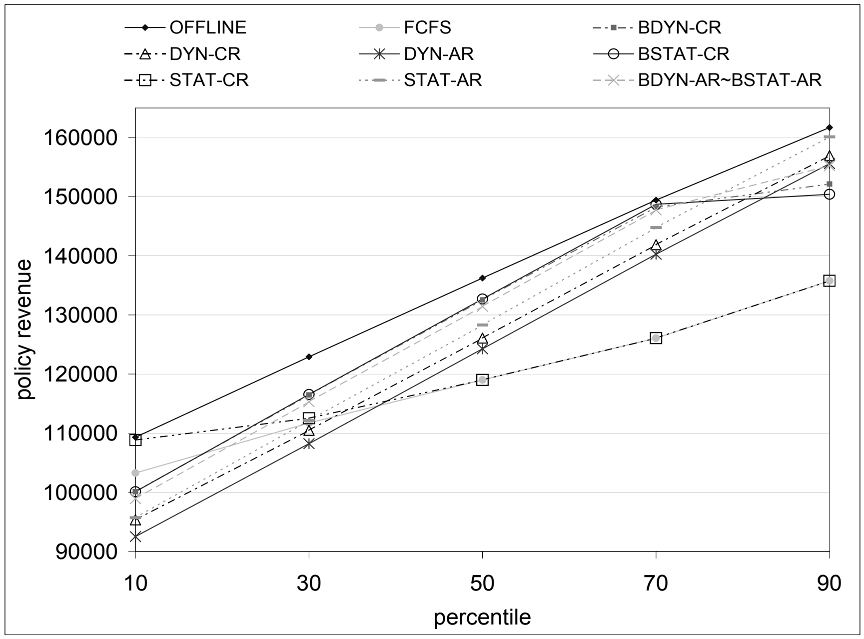

Figure A.7: Average performance gap in Example-3d

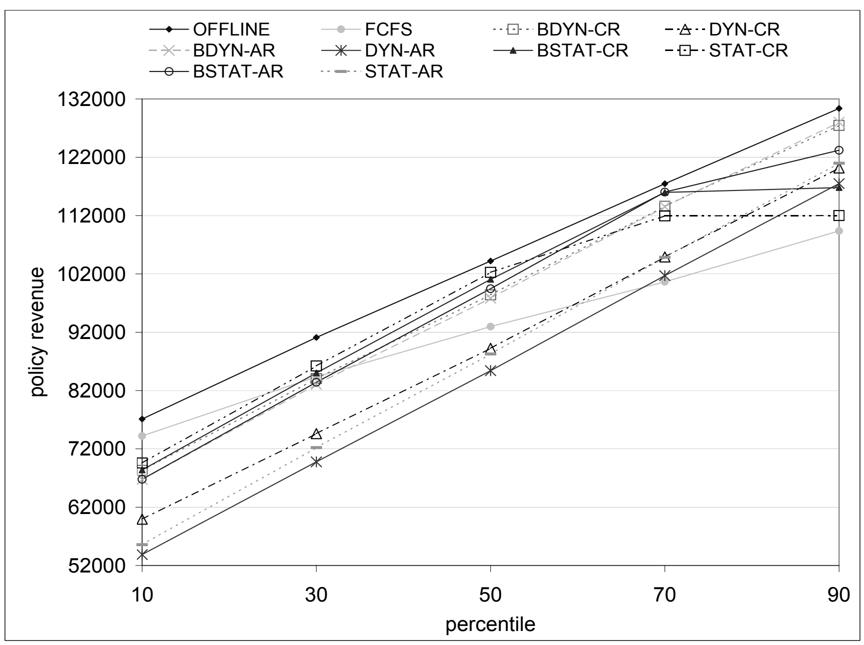

Figure A.8: Average performance gap in Example-3e centiles for the policies are presented in Figure A.9.

Example-A.10: This example is the same as Example-7 with performance of robust policies require no demand information. The average performance gaps are specified in Figure A.10

Example-A.11: This example is the same as Example-9 except that all the demand arrivals later in the second half of the booking horizon. The average performance gaps are specified in Figure A.11

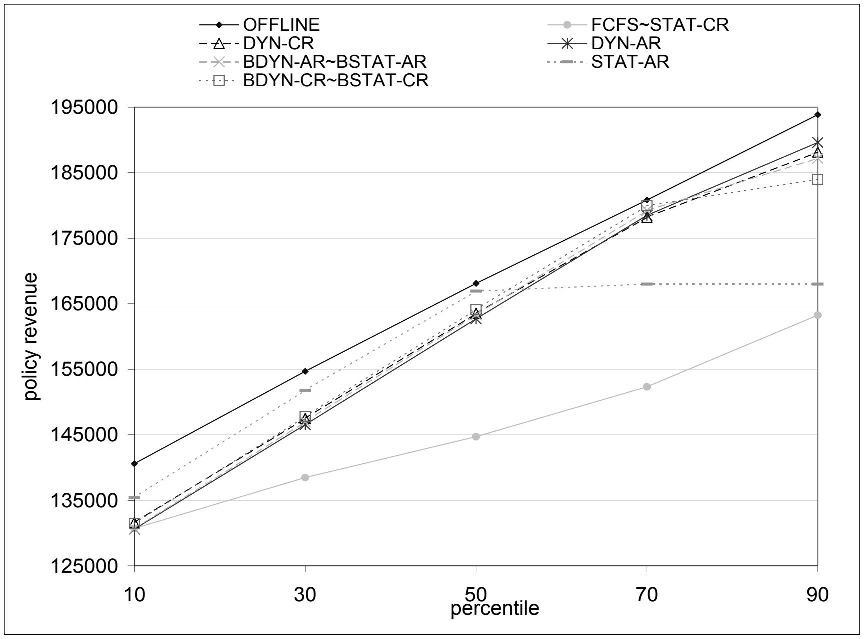

Figure A.9: Average performance gap in Example-3f

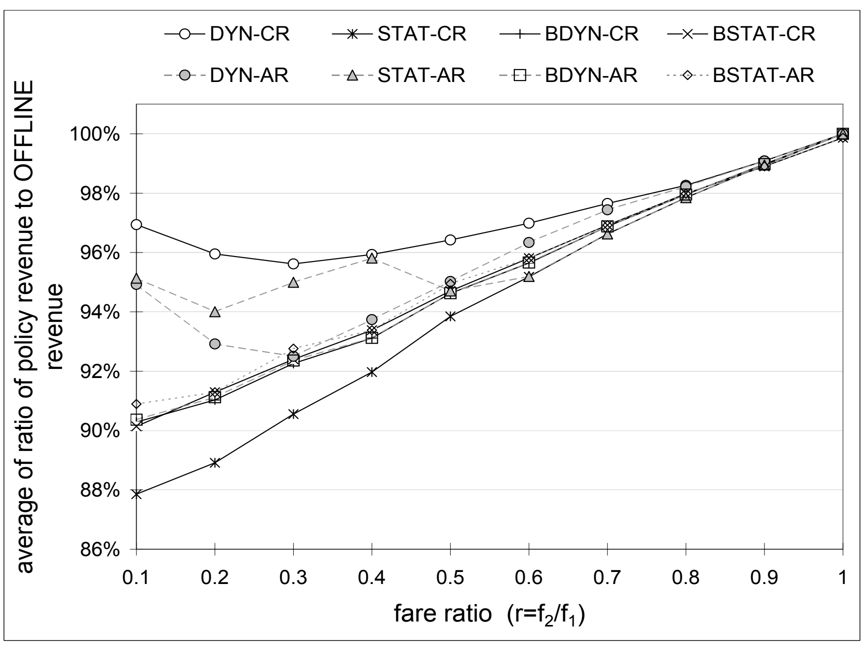

Figure A.10: Average performance gap in Example-7b

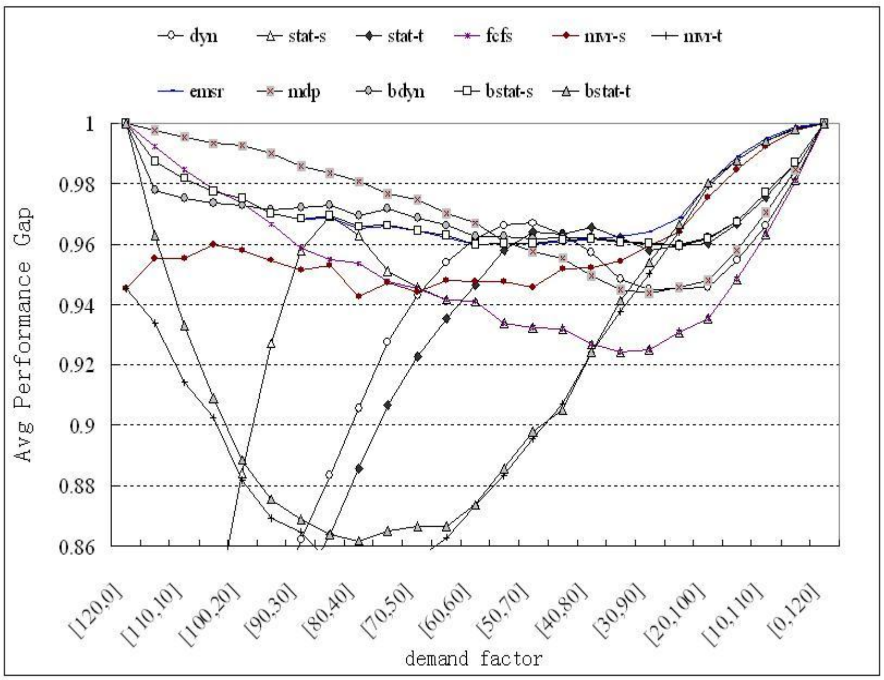

Figure A.11: Average performance gap in Example-9b

Appendix B
Appendices for Chapter 3
B.1 Proof of Proposition 2

To simplify the notation and to suppress the time index, we refer to the seller's static decision vector as β, as opposed to b^1, in the remainder of this section. We first provide additional notation and state an important observation. Combining the notion of effective booking limits (as expressed by algorithm $PNEST(t)$ in Section 3.2.1) with Proposition 1, one can iteratively compute the number of requests accepted in each period and the effective booking limit at the start of each period:

Observation 1 *Consider a nested booking limit policy β and input $\vec{I} = [I^1, .., I^T]$ where inputs arrive in LBH order in each period. Let the number of class j requests accepted in period t after processing I^t be $h_j^{I^t}$. Let the effective nested booking limit vector at the beginning of period t be $\beta^{I^t} = (\beta_1^{I^t}, ..., \beta_m^{I^t})$. Parameters $h_j^{I^t}$ and $\beta_j^{I^t}$ are computed iteratively as follows:*

Step 0: Set $t = 1$. Set $\beta^{I^1} = \beta$, i.e., $(\beta_1^{I^1}, \beta_2^{I^1}, ..., \beta_m^{I^1}) = (\beta_1, ..., \beta_m)$.

Step 1: Compute the number of reservations accepted in each fare class in period t, starting with class m and working backwards from class m to class 1:

$$h_m^{I^t} = min(\beta_m^{I^t}, I^t[m]), \tag{B.1}$$

$$h_j^{I^t} = min(\beta_j^{I^t} - \sum_{k=j+1}^{m} h_k^{I^t}, I^t[j]) \quad for \quad j = m-1, m-2, ..., 1 \tag{B.2}$$

Step 2: Compute the effective booking limit of each fare class at the beginning of period $t+1$, starting with class 1 and working forwards from class 1 to class m:

$$\beta_1^{I^{t+1}} = \beta_1^{I^t} - \sum_{k=1}^{m} h_k^{I^t}, \tag{B.3}$$

$$\beta_j^{I^{t+1}} = min(\beta_{j-1}^{I^{t+1}}, \beta_j^{I^t} - \sum_{k=j}^{m} h_k^{I^t}) \quad for \quad j = 2, ..., m. \tag{B.4}$$

Step 3: If $t = T$, STOP. Otherwise, set $t \leftarrow t+1$ and go to Step 1.

The vector $\beta^{I^{T+1}}$ provides information (in a nested fashion) on the number of unused seats at the end of the booking horizon after processing input $\vec{I}$ with policy β.

We use these relations and the corresponding notation in the proof below.

Proposition 2: *Given a static, nested booking limit policy β and the set Q of extreme input sequences, the lowest CR is achieved by one of the inputs in set Q. That is, for any feasible input sequence $\vec{I}$, we have*

$$\frac{R(\vec{I},\beta)}{R^*(\vec{I})} \geq min\{\frac{R(\vec{I^*},\beta)}{R^*(\vec{I^*})}, \quad \forall \vec{I^*} \in Q\}.$$

Proof Let us consider a sequence $\vec{I} = [I^1, ..., I^T]$ which is period-wise LBH (without loss of generality). We compute the effective booking limits $\beta^{I^1}, ..., \beta^{I^{T+1}}$ for input $\vec{I}$ and policy β as described in Observation 1 above. For all $t = 1, ..., T$, let $j^*(t)$ be the highest-fare (lowest-index) fare class whose effective booking limit is zero at the end of period t, i.e., $\beta_k^{I^{t+1}} = 0$ for all $k \geq j^*(t)$ and $\beta_k^{I^{t+1}} > 0$ for all $k < j^*(t)$, $k = 1, ..., m$. If $\beta_k^{I^{t+1}} > 0$ for all $k = 1, .., m$, let $j^*(t) = m + 1$. Note that if $j^*(t) = m + 1$, then all requests that arrive in period t are accepted (because no effective booking limit is reached in that period). Note also $j^*(1) \geq j^*(2) \geq \cdots \geq j^*(T)$ by definition of standard nesting.

Starting with $t = T$, we show that the adversary can get a CR lower than $R(\vec{I},\beta)/R^*(\vec{I})$ by sending one of the m extreme input sequences in period t instead of sending I^t. For $t = T$, the proof in Lan (2008) follows immediately: When $j^*(T) < m+1$, then the adversary can not increase the CR by sending an input $I^{T\prime}$ such that $I^{T\prime}[k] = L_k^T$ $k < j^*(T)$ and $I^{T\prime}[k] = U_k^T$ for $k \geq j^*(T)$, $k = 1, ..., m$. When $j^*(T) = m + 1$, all requests have been accepted throughout the booking horizon and the CR is $R(\vec{I},\beta)/R^*(\vec{I}) = 1$. In this latter case, adversary can send any of the j-extreme input sequences for period T without increasing the CR.

Consider any period $t < T$.

Case 1: For all $k \geq j^*(t)$, the adversary can lower the CR by sending a LBH input in period t with $I^{t\prime}[k] = U_k^t \geq I^t[k]$, using the same argument in Lan (2008): since effective booking limits of $k \geq j^*(t)$ are all zero, additional requests in these classes are all rejected. Online revenues are unaffected by the change, while offline optimal can only increase or remain the same. Therefore CR is improved by this change.

Case 2: For all $k < j^*(t)$, the adversary can lower the CR by sending a LBH input in period t with the profile (i) $I^{t\prime}[k] = L_k^t$ if $b_k^{T+1} > 0$ (in this case, a class k request in period t does not displace an equal or higher fare request in period t or any time in the remainder of the booking horizon) or (ii) $I^{t\prime}[k] = U_k^t$ if $\beta_k^{T+1} = 0$ (in this case, a class k request displaces equal or higher fare requests either in period t or in the remainder of the booking horizon). The former follows from Lan (2008). In the latter, while the effective booking limit of class k is positive at the end of period t, it will become zero by the end of the planning horizon. This means that the argument for Case 1 will be valid for class k some time before the end of the booking horizon. Consequently, it is to the adversary's advantage to send the maximum number of requests for classes that will be rejected either in this period or in the future. •

B.2 Determining the Optimal Booking Limits in the Static Problem

In this section, we show that (3.4) can be expressed as a mixed integer programming (MIP) model to determine the optimal static booking limits. The binary variables are used in determining the correct number of fare-class requests accepted in processing an input sequence, which is then used in calculating the online revenues. Binary variables are also used to ensure effective booking limits are calculated correctly, preserving the nesting property. To simplify the notation in the problem formulation, we present the model for all extreme sequences belonging to the set Q, without referring to the property in Corollary 1.

B.2.1 The MIP Model

In MIP, the input parameters are f_k for $k = 1, ..., m$, the capacity n, the period-wise profile $I^t[j]$, $t = 1, .., T$, for each input $\vec{I} = [I^1, ..., I^T] \in Q$, and the corresponding offline optimal revenue $R^*(\vec{I})$ for $\vec{I} \in Q$. $R^*(\vec{I})$ is easily computed a priori given the profile of the extreme sequences. The decision variable for the seller is the vector $\beta = (\beta_1, ..., \beta_m)$ and the scalar z represents the CR. All other variables are auxiliary and ensure the effective booking limits follow the relations provided in Observation 1 of Appendix B.1.

MIP:

$$\max \quad z$$

$$\text{s.t.} \quad R^*(\vec{I})z \leq \sum_{t=1}^{T}\sum_{k=1}^{m} f_k h_k^{I^t}, \quad \forall\ t = 1, ...T,\ \forall \vec{I} \in Q \tag{B.5}$$

$$h_j^{I^t} \leq I^t[j], \quad j = 1, ..., m, \quad t = 1, ..., T, \forall \vec{I} \in Q \tag{B.6}$$

$$\sum_{k=j}^{m} h_k^{I^t} \leq \beta_j^{I^t}, j = 1, ..., m, t = 1, ...T, \forall \vec{I} \in Q \tag{B.7}$$

$$I^t[j] - h_j^{I^t} - nv_j^{I^t} \leq 0, j = 1, ...m, t = 1, ..., T, \forall \vec{I} \in Q \tag{B.8}$$

$$\beta_j^{I^t} - \sum_{k=j}^{m} h_k^{I^t} - n(1 - v_j^{I^t}) \leq 0, j = 1, ..., m, t = 1, ..., T, \forall \vec{I} \in Q \tag{B.9}$$

$$\beta_j^{I_{t+1}} \leq \beta_{j-1}^{I_{t+1}}, j = 2, ..., m, t = 1, ..., (T-1), \forall \vec{I} \in Q \tag{B.10}$$

$$\beta_j^{I_{t+1}} \leq \beta_j^{I^t} - \sum_{k=j}^{m} h_k^{I^t}, j = 2, ..., m, t = 1, ..., (T-1), \forall \vec{I} \in Q \tag{B.11}$$

$$\beta_j^{I_{t+1}} - \beta_{j-1}^{I_{t+1}} - nw_j^{I_{t+1}} \le 0, j = 2, ..., m, t = 1, ..., (T-1), \forall \vec{I} \in Q \tag{B.12}$$

$$\beta_j^{I^t} - \sum_{k=j}^{m} h_k^{I^t} - \beta_j^{I_{t+1}} - n(1 - w_j^{I_{t+1}}) \le 0,$$

$$j = 2, ..., m, t = 1, ..., (T-1), \forall \vec{I} \in Q \tag{B.13}$$

$$\beta_1^{I_{t+1}} = \beta_1^{I^t} - \sum_{k=1}^{m} h_k^{I^t}, t = 1, ..., (T-1), \forall \vec{I} \in Q \tag{B.14}$$

$$\beta_j^{I^1} = \beta_j, \quad j = 1, ..., m, \quad \forall \vec{I} \in Q \tag{B.15}$$

$$\beta_j^{I^t}, h_j^{I^t} \ge 0; v_j^{I^t}, w_j^{I^t} \in \{0, 1\}, j = 1, ..., m, t = 1, ..., T, \forall \vec{I} \in Q. \tag{B.16}$$

$$0 \le \beta_j \le n, \quad j = 1, ..., m. \tag{B.17}$$

The size of MIP is O(m^{T+1}) as the number of constraints and binary variables are proportional to m^{T+1}, which is polynomial in m for fixed T. The optimal objective value in MIP is at least as much as the objective value derived in the static model of Lan (2008) when the model in Lan (2008) uses the aggregate lower and upper bounds $\sum_{t=1}^{T} L_j^t$ and $\sum_{t=1}^{T} U_j^t$, respectively, for a single-period problem. This is because there are period-wise restrictions on the input sequence. We provide an example to show the differences in the solutions of our static model and that of Lan (2008) below.

B.2.2 Comparison of our static model to that of Lan (2008)

Here is an example to show the differences in solutions between our static model and that of Lan (2008). Both the theoretical CR and the average performance of the policies, computed using simulations, are considered: Suppose $n = 10$,$f_1 = 5$, $f_2 = 1$, $T = 2$, $L_1^1 = 0$, $L_2^1 = 0$, $L_1^2 = 2$, $L_2^2 = 0$, $U_1^1 = 8$, $U_2^1 = 3$, $U_1^2 = 6$, and $U_2^2 = 12$. Demand for class 2 is assumed to be uniform distributed between L_2^t and U_2^t in period $t = 1, 2$. In two different experiments, the demand for class 1 is first uniform distributed, and then triangular distributed, each with the range $[L_1^t, U_1^t]$ for $t = 1, 2$. The theoretical CR and the average gap (computed as the ratio of policy revenues to offline optimal revenues, averaged across all simulation runs), as an indication of the actual CR, are computed using simulation for policies obtained by MIP and Lan (2008). The results are presented in Table B.1. The optimal booking limit for class 2 derived using the static CR model in Lan (2008) provide $b_2 = 3.28$, and the theoretical CR of this policy is 0.73. However, the static seller's optimal strategy from the MIP model suggests $b_2 = 3.68$ and the theoretical competitive ratio is 0.76. The average performance of the policies differ depending on the demand distribution used in this experiment. When class 1 demand is uniform distributed, the average performance gap is 0.878 for Lan (2008) and 0.84 for MIP. If class 1 demand in the first period is triangular distributed with a mean of 2, then the average performance gap is 0.9788 for Lan (2008) and 0.99 for MIP. Although MIP guarantees a better worst-case performance, its actual average performance need not always be superior compared to Lan (2008).

Table B.1: Booking Limits, Theoretical CR and Average Performance Gap

Model	Booking Limit b_2	Theoretical CR	Average Gap D_1 is Uniform	Average Gap D_1 is Triangular
Lan (2008)	3.28	0.73	0.878	0.9788
MIP	3.68	0.76	0.84	0.99

B.3 Proof of Proposition 4

Proof Through induction, we will show that for a m-fare, T-period problem, the lowest competitive ratio is achieved by one of the $m+1$ extreme input sequences. We use the notation $I \prec\prec I'$ to denote that given booking limit vector b^T of the seller and two inputs I and I', I results in a CR that is no-higher than I' (i.e., I' is dominated by I). We focus only on LBH inputs.

STEP 1: We will prove that in the last period and for the lowest fare, adversary will send in the amount of $L_T^m + \Delta^m$, $\Delta^m \in \{0, U_T^m - L_T^m\}$. Consider three LBH input sequences, S_1, S_2, S_3 with the following profiles: $S_1^T[m] = U_m^T$, $S_2^T[m] = L_T^m + \Theta^m$ for $0 < \Theta^m < U_T^m - L_T^m$, $S_3^T[m] = L_T^m$, and $S_1^T[j] = S_2^T[j] = S_3^T[j]$ for all $j = 1, .., m-1$. S_2 is dominated by either S_1 or S_3:

Case 1.1: If $0 \leq b_m^T \leq L_T^m + \Theta^m$, $S_1 \prec\prec S_2$ because the online revenues are equal with inputs S_1 and S_2 and the offline optimal revenue $R^*(S_1)$ is the highest with S_1.

Case 1.2: Consider $b_T^m > L_T^m + \Theta^m$. (i) If processing S_2 with b^T and accepting $L_T^m + \Theta^m$ requests in class m does not displace any higher fare requests, then S_3 also does not displace any higher fare requests when processed by b^T. In this situation, online revenues satisfy $R(S_2, b^T) \geq R(S_3, b^T)$ and offline revenues satisfy $R^*(S_2) \geq R^*(S_3)$. Notice that $R^*(S_2) - R^*(S_3) \leq f_m * \Theta^m$ and $R(S_2, b^T) = f_m * \Theta^m + R(S_3, b^T)$. This implies that $\frac{R^T + R(S_2, b^T)}{R^*(S_2)} \geq \frac{R^T + R(S_3, b^T)}{R^*(S_3)}$. Therefore, $S_3 \prec\prec S_2$. (ii) If processing S_2 with b^T and accepting $L_T^m + \Theta^m$ requests in class m displaces higher-fare requests, then S_1 also leads to displacement of higher-fare requests. In this situation, $R^T + R(S_1, b^T) \leq R^T + R(S_2, b^T)$ because S_1 accepts a higher number of class m requests, and $R^*(S_1) \geq R^*(S_2)$, which leads to $\frac{R^T + R(S_1, b^T)}{R^*(S_1)} \leq \frac{R^T + R(S_2, b^T)}{R^*(S_2)}$. Thus,$S_1 \prec\prec S_2$.

Therefore, a LBH sequence I^T with profile $L_m^T < I^T[m] < U_m^T$ is dominated.

STEP 2: Through backward induction on fare classes, we will show that a LBH sequence I^T with profile $L_j^T < I^T[j] < U_j^T$ is dominated. The statement is true for class m based on STEP 1. Suppose it is true for classes $m, m-1, ..., k$. We will show this is true for class $k-1$. Consider six LBH input sequences $S_1', ..., S_6'$ with profiles $S_1'[k] = L_k^T$, $S_1'[k-1] = L_{k-1}^T$, $S_2'[k] = L_k^T$, $S_2'[k-1] = L_{k-1}^T + \theta$ for $0 < \theta < U_{k-1}^T - L_{k-1}^T$, $S_3'[k] = L_k^T$, $S_3'[k-1] = U_{k-1}^T$, $S_4'[k] = U_k^T$, $S_4'[k-1] = L_{k-1}^T$,

$S'_5[k] = U_k^T$, $S'_5[k-1] = L_{k-1}^T + \theta$ for $0 < \theta < U_{k-1}^T - L_{k-1}^T$, and $S'_6[k] = U_k^T$, $S'_6[k-1] = U_{k-1}^T$.

Case 2.1: Consider S'_1, S'_2, S'_3 where L_k^T requests of class k arrive. (i) If all of the $S'_2[k-1]$ requests are accepted by b^T and none of the requests in classes $k-2$ to 1 are displaced, then all of the $S'_1[k-1]$ requests are also accepted by b^T without displacing higher-fare requests. This is similar to Case 1.2(i) above. Therefore, $S'_1 \prec\prec S'_2$. (ii) If all of the $S'_2[k-1]$ requests are accepted by b^T, leading to displacement of requests in any of classes $k-2$ to 1, then at least $S'_2[k-1]$ requests of class $k-1$ will be accepted when S'_3 is processed by b^T, leading to at least as many displacements of higher-fare. This is similar to Case 1.2(ii) above. Thus, $S'_3 \prec\prec S'_2$.

Case 2.2: Consider S'_4, S'_5, S'_6 where U_k^T requests of class k arrive. Using the same argument as in Case 2.1 above, one can show that either $S'_4 \prec\prec S'_5$ or $S'_6 \prec\prec S'_5$.

Therefore, an input I^T with $L_j^T < I^T[j] < U_j^T$ for any j is dominated. Next, we show there are only $m+1$ of the non-dominated sequences.

STEP 3. We know any extreme LBH input I^T has profile $I^T[m] = L_m^T$ or $I^T[m] = U_m^T$. Consider LBH sequences S''_1, S''_2, and S''_3 with profiles $S''_1[m] = U_m^T$, $S''_1[m-1] = U_{m-1}^T$, $S''_2[m] = L_m^T$, $S''_2[m-1] = U_{m-1}^T$ and $S''_3[m] = L_m^T$, $S''_3[m-1] = L_{m-1}^T$.

Case 3.1: If processing S''_2 with b^T displaces any request in classes $m-1$ to 1, then S''_1 leads to at least the same number of displaced requests in these classes. Then, using the logic in Case 1.2(ii) above, we have $S''_1 \prec\prec S''_2$.

Case 3.2: If processing S''_2 with b^T does not result in displacement of any request in classes $m-1$ to 1, $S''_3 \prec\prec S''_2$ by Case 1.2(i) above.

Combining Cases 3.1 and 3.2, we see that S''_2 is dominated. Therefore a non-dominating sequence I^T has a profile such that if $I^T[m] = L_m^T$, then $I^T[m-1] = L_{m-1}^T$. Note that, this observation can be extended to any two consecutive class j and j-1.

Therefore, the total number of non-dominated input streams in period T is only $m+1$ and Q^{*T} is the set of all non-dominated sequences for period T. •

B.4 Proof of Proposition 5

B.4.1 Preliminaries

In this section, we present elementary observations and results that are used in the proof of Proposition 5. First of all, we make use of duality in proving the optimality of the suggested solution for LPM. The dual formulation of LPM, called

DLPM, is given below.

$$DLPM: \quad \min \ n^T\nu + \sum_{j=1}^{m} U_j\varpi_j + \sum_{j=1}^{m}(R^T + \sum_{i=1}^{j-1} f_i L_i^T)\lambda_j + (R^T + \sum_{i=1}^{m+1} f_i L_i^T)\lambda_{m+1} \tag{B.18}$$

$$s.t. \quad \sum_{j=1}^{m+1} R_j^* \lambda_j \geq 1, \tag{B.19}$$

$$\varpi_j + \nu - f_j * \sum_{i=1}^{j} \lambda_i \geq 0, \quad j = 1,..,m, \tag{B.20}$$

$$\nu \geq 0, \quad \lambda_j, \varpi_j \geq 0, j = 1,..,m. \tag{B.21}$$

Second, the solution to LPM is closely related to the solution of the following linear program, which disregards the $(m+1)^{st}$ extreme sequence:

$$GBM: \max \quad z$$

s.t.

$$R^*(I^1 \cdots I^{T-1} I^{*T,k})z \quad \leq \quad R^T + \sum_{i=1}^{k-1} f_i L_i^T + \sum_{i=k}^{m} f_i x_i^T, \quad k = 1,...,m, \tag{B.22}$$

$$\sum_{j=1}^{m} x_j^T \quad \leq \quad n^T \tag{B.23}$$

$$0 \leq \quad x_j^T \quad \leq U_j^T, \quad j = 1,...,m \tag{B.24}$$

When $R^T = 0$, this is equivalent to the formulation in Lan (2008) that determines the static nested booking limits in a single-period problem. The closed-form solution to GBM is given by

$$z^{*GBM} \quad = \quad \theta_u \tag{B.25}$$

$$x_k^{*GBM} \quad = \quad \begin{cases} g_k z^* + L_k^T & 1 \leq k < u \\ (R_k^* z^* - \sum_{i=1}^{k-1} f_i L_i^T - R^T)/f_k & k = u \\ 0 & k > u \end{cases} \tag{B.26}$$

where the parameters are defined as

$$g_m = R_m^*, \quad g_k = (R_k^* - R_{k+1}^*)/f_k \geq 0, \quad k = 1,...,m-1$$

$$u = \max\{j : (R^T + \sum_{i=1}^{j-1} f_i L_i^T) \sum_{i=1}^{j-1} g_i < N_j^T R_j^*\}, \tag{B.27}$$

$$\theta_u = \frac{(R^T + \sum_{i=1}^{u-1} f_i L_i^T)/f_u + N_u^T}{R_u^*/f_u + \sum_{i=1}^{u-1} g_i}.$$

We omit the proof of optimality of the solution given in (B.25) and (B.26) for GBM here; that result follows from the proof in Lan (2008) with the addition of the scalar

$R^T > 0$ to the online policy revenues for each extreme input. Notice that the solution of GBM only partially characterizes the optimal solution of LPM.

Finally, we use one additional parameter in the solution of LPM:

$$\pi = \frac{R^T + \sum_{i=1}^{m-1} f_i L_i^T + f_m * min(L_m^T, N_m^T)}{R_{m+1}^*}.$$

Note that π is a constant and it represents the highest competitive ratio possible for any booking limit policy when $(m+1)^{st}$ extreme sequence is processed. This is what is represented in constraint (3.9) of LPM. We have the following observation on parameter π.

Lemma 1 *If $u = m$ and $\theta_m \geq \pi$, then $\pi = \frac{R^T + \sum_{i=1}^{m} f_i L_i^T}{R_{m+1}^*}$.*

<u>Proof</u> We show this by proving that when $u = m$ and $\theta_m \geq \pi$, then it is impossible to have $N_m^T < L_m^T$. Suppose $N_m^T < L_m^T$, $u = m$, and $\theta_m \geq \pi$. Then, by definition,

$$\theta_m = \frac{R^T + \sum_{i=1}^{m-1} f_i L_i^T + N_m^T * f_m}{R_m^* + \sum_{i=1}^{m-1} g_i * f_m} = \frac{R^T + \sum_{i=1}^{m-1} f_i L_i^T + N_m^T * f_m}{R_{m+1}^* + \sum_{i=1}^{m} g_i * f_m},$$

$$\pi = \frac{R^T + \sum_{i=1}^{m-1} f_i L_i^T + N_m^T * f_m}{R_{m+1}^*},$$

and $\theta_m < \pi$, which is contradictory to the condition that $\theta_m \geq \pi$. Thus,in this case, $N_m^T \geq L_m^T$ and the resulting π is $\frac{R^T + \sum_{i=1}^{m} f_i L_i^T}{R_{m+1}^*}$. •

In addition to the solution to GBM, consider the alternative

$$z^{*ALT} = \pi \tag{B.28}$$

$$x_k^{*ALT} = \begin{cases} g_k z^{*ALT} + L_k^T & 1 < k \leq m \\ \min(U_T^1, n^T - \sum_{i=2}^{m}(L_i^T + g_i \pi)) & k = 1 \end{cases} \tag{B.29}$$

which if feasible for LPM under certain conditions:

Lemma 2 *When $u = m$ and $\theta_m \geq \pi$, then x^{*ALT} of equation (B.29) and $z^{*ALT} = \pi$ is feasible for LPM.*

<u>Proof</u> To prove this, we first have to prove that $x_1^{*ALT} \geq L_T^1 + g_1 \pi$: When $u = m$, we have $N_m^T > 0$, that is, $n^T > \sum_{i=1}^{m-1} L_T^i$. Let's define $R_m^+ = \sum_{i=1}^{m-1} f_i L_T^i$. In this case, we know that $\theta_m \geq \pi$, so based on Lemma 1,

$$\theta_m = \frac{R^T + R_m^+ + (n^T - \sum_{i=1}^{m-1} L_i^T) * f_m}{R_m^* + \sum_{i=1}^{m-1} g_i * f_m} \geq \frac{R^T + \sum_{i=1}^{m} f_i L_i^T}{R_{m+1}^*} = \pi.$$

Then, we get the following series of relations by algebraic manipulations:

$$
\begin{aligned}
(n^T - \sum_{i=1}^{m-1} L_i^T) f_m R_{m+1}^* &\geq (R^T + R_m^+)(R_m^* - R_{m+1}^*) + R_m^* f_m L_m^T \\
&\qquad + f_m \sum_{i=1}^{m-1} g_i (R^T + R_m^+) + \sum_{i=1}^{m-1} g_i f_m L_m^T \\
(n^T - \sum_{i=1}^{m-1} L_i^T) R_{m+1}^* &\geq (R^T + R_m^+) \sum_{i=1}^{m} g_i + R_m^* L_m^T + \sum_{i=1}^{m-1} g_i f_m L_m^T \\
f_m g_m + R_{m+1}^* &= R_m^* \\
(n^T - \sum_{i=1}^{m-1} L_i^T) &\geq \frac{(R^T + R_m^+) \sum_{i=1}^{m} g_i}{R_{m+1}^*} + \frac{\sum_{i=1}^{m} g_i f_m L_m^T}{R_{m+1}^*} + L_m^T \\
n^T &\geq \sum_{i=1}^{m} L_i^T + \sum_{i=1}^{m} g_i * \pi.
\end{aligned}
$$

Then

$$
x_1^{*ALT} = n^T - \sum_{i=2}^{m} L_i^T + \sum_{i=2}^{m} g_i * \pi \geq L_1^T + g_1 \pi. \tag{B.30}
$$

Next we prove feasibility of x^{*ALT}. Here are the observations:

- Constraint (3.11) is satisfied: We have $g_i \leq U_T^i - L_T^i$ for $i = 1, ..., m$; this follows from the proof in Lan (2008). By construction, $x_1^{*ALT} \leq U_T^1$. We also have $z^* = \pi \leq 1$ by design. This implies $x_k^{*ALT} = \pi * g_i + L_T^i \leq U_T^i$ for $1 < i \leq m$.

- Constraint (3.8) is satisfied: We have $u = m, \pi \leq \theta_m = (R^T + \sum_{i=1}^{m} f_i L_i^T)/R_{m+1}^*$. So for any k, $2 \leq k \leq m$, we have

$$
R_k^* \pi \leq R^T + \sum_{i=1}^{m} f_i L_i^T + R_k^* \pi - R_{m+1}^* \pi.
$$

Further algebra yields

$$
R_k^* \pi \leq R^T + \sum_{i=1}^{m} f_i L_i^T + \sum_{i=k}^{m} (R_i^* - R_{i+1}^*) \pi
$$

$$
R_k^* \pi \leq R^T + \sum_{i=1}^{k-1} f_i L_i^T + \sum_{i=k}^{m} f_i (L_i^T + g_i \pi)
$$

$$
R_k^* \pi \leq R^T + \sum_{i=1}^{k-1} f_i L_i^T + \sum_{i=k}^{m} f_i x_i^{*ALT}.
$$

For $k = 1$, we have $L_T^1 + g_1 \pi \leq x_1^{*ALT} \leq U_T^1$ which follows from equation (B.30) above. This yields $R_1^* \pi \leq R^T + \sum_{i=1}^{m} f_i x_i^{*ALT}$.

- Constraint (3.9) is satisfied: This follows from the definition of π:

$$R^*_{m+1}\pi \leq R^T + \sum_{i=1}^{m-1} f_i L_i^T + f_m\ min(L_m^T, N_m^T).$$

- The remaining constraint $\sum_{i=1}^{m} x_i^{*ALT} \leq n^T$ is satisfied.

•

B.4.2 Proof of Proposition 5

Proof We first prove the feasibility of the solution suggested, then prove its optimality. (a) Proof of feasibility :

Case 1: If the optimal solution of GBM has $u < m$, then constraint (3.8) for $k = m$ is not binding because only the first u out of those m constraints in the GBM are binding. This follows from the proof in Lan (2008). Thus, in this case, given x^{*GBM}, the corresponding $z^* = \theta_u$ satisfies constraint (3.8) when $k = m$, and also satisfies constraint (3.9) of LPM because constraint (3.9) yields the same online revenue but lower offline optimal revenue given the $(m+1)^{st}$ extreme sequence. So $z^* = \theta_u$ and x^{*GBM} constitute a feasible solution to LPM.

Case 2: If the optimal solution of GBM has $u = m$ and $\theta_m < \pi$, then constraint (3.9) in LPM is not binding for the solution x^{*GBM} and $z^* = \theta_m$. Therefore the solution x^{*GBM} and $z^* = \theta_m$ of GBM is again feasible for LPM.

Case 3: If the optimal solution of GBM has $u = m$ and $\theta_m \geq \pi$, then Lemma 2 proves x^{*ALT} and z^{*ALT} are feasible for LPM.

(b)Proof of optimality:

Cases 1 and 2: Consider the following solution for DLPM:

$$\varpi_j^* = 0, \quad j = 1, ..., m, \quad \lambda_j^* = 0 \quad j = u+1, ..., m+1 \tag{B.31}$$

$$\nu^* = 1/(R_u^*/f_u + \sum_{i=1}^{u-1} g_i) \tag{B.32}$$

$$\lambda_j^* = \nu^*(1/f_j - 1/f_{j-1}) \quad j = 1, ..., u. \tag{B.33}$$

It can be showed that this solution is feasible for DLPM and its objective function value is equal to θ_u. Therefore the solution x^{*GBM} and z^{*GBM} is the optimal solution for LPM. Case 3: Consider the following feasible solution to DLPM:

$$\begin{aligned}
\lambda_j &= \begin{cases} 0 & j \leq m \\ 1/R^*_{m+1} & j = m+1 \end{cases} \\
\nu &= 0, \\
\varpi_j &= 0 \qquad\qquad\qquad \forall j = 1, ..., m.
\end{aligned}$$

The objective function value of DLPM is equal to π for this trivial solution. Therefore, the solution x^{*ALT} and z^{*ALT} is the optimal solution to LPM. •

B.5 Proof of Proposition 6

Proof We consider the same cases introduced in the proof of Proposition 5.

Cases 1 and 2: If $u < m$ or $u = m$ and $\theta_m < \pi$ in the solution of GBM, then constraint (3.9) is not binding in LPM. In these cases, GBM solution provides the optimal booking limit and its CR is correctly calculated in GBM as shown in Lan et al.(2007).

Case 3: If solution to GBM has $u = m$ and $\theta_m \geq \pi$, then x^{*ALT} is the optimal booking limit policy. In this case, we know $u = m$, $z^* = \pi = \frac{R^T + \sum_{i=1}^{m} f_i L_i^T}{R^*_{m+1}}$ is the optimal objective value of LPM, and constraint (3.9) is binding in LPM. Given the k^{th} extreme input sequence $I^{*T,k}$ and a nested booking limit policy b^{*T} obtained from x^{*ALT} of equation (B.29), the online policy revenue satisfies $R^T + R(I^{*T,k}, b^{*T}) = R^T + \sum_{i=1}^{k-1} f_i min(x_i^{*ALT}, I^{*T,k}[i])$ by construction of x^{*ALT}. Note that $L_i^T \leq x_i^{*ALT} \leq U_i^T$ and $R^T + R(I^{*T,k}, b^{*T}) = R^T + \sum_{i=1}^{k-1} f_i L_i^T + \sum_{i=k}^{m} f_i x_i^{*ALT}$ for extreme sequence k. Therefore, the constraints of LPM correctly calculate the online revenues of policy b^{*T} for inputs $k = 1, .., m+1$. The minimum CR is achieved by input $m + 1$ in this case, and the corresponding CR is equal to π. •

Albers, S. 2003. Online algorithms: A Survey. *Math. Programming, Ser. B,* 97, 3-26.

Alstrup, J., S. Boas, O.B.G. Madsen and R.V.V. Vidal. 1986. Booking Policy for Flights with Two Types of Passengers. *European Journal of Operational Research* **27** 274-288.

Alstrup, J., S. Andersson, S. Boas, O.B.G. Madsen and R.V.V. Vidal. 1989. Booking Control Increases Profit at Scandinavian Airlines. *Interfaces* **19** 10-19.

Bailey, J. 2007. Bumped Fliers and No Plan B. *The New York Times,* May 30, 2007.

Ball, M. O. and M. Queyranne. 2006. Toward Robust Revenue Management: Competitive Analysis of Online Booking. Working paper, University of Maryland, College Park, MD; available at SSRN: http://ssrn.com/abstract=896547.

Beckmann, M.J. 1958. Decision and Team Problems in Airline Reservations. *Econometrica,* **26** 134-145.

Belobaba, P.P. 1987. Airline Yield Management: An Overview of Seat Inventory Control. *Transportation Science* 21, 63-73.

Belobaba, P.P. 1989. Application of a Probabilistic Decision Model to Airline Seat Inventory Control. *Operations Research,* 37, 183-197.

Belobaba, P. 2006. Flight Overbooking: Models and Practice. Lecture notes. Massachusettes Institute of Technology, Boston, MA, $http : //ocw.mit.edu/NR/rdonlyres/Aeronautics-and-Astronautics/16-75JSpring-2006/ E2A09BF6-E275 - 431B - 9C07 - 2C98BEA50541/0/lect19.pdf$

Birbil, S. I., J.B.G. Frenk, J. A.S. Gromicho and S. Zhang. 2006. An Integrated Approach to Single-Leg Airline Revenue Management: The Role of Robust Optimization. Working paper, Erasmus University, Erasmus Research University Institute of Management (ERIM), Rotterdam, The Netherlands.

Bertsimas, D. and M. Sims. 2004. The price of robustness. *Oper. Res.* 52, 35-53.

Besbes, O. and A. Zeevi. 2006. Dynamic pricing without knowing the demand function: Risk bounds and near optimal pricing algorithms. Working paper, Columbia University, Graduate School of Business.

Bitran, G.R. and S. Gilbert. 1996. Managing Hotel Reservations with Uncertain Arrivals. *Operations Research* **44** 35-49.

Boyd,D. and Phillips,R. 1998, "Revenue Management",

$$http : //www.talus.net/news/boyd_phillips_article.htm.$$

Brumelle, S.L. and J.I. McGill. 1993. Airline Seat Allocation with Multiple Nested Fare Classes. *Operations Research* **41** 127-137.

Brumelle, S. and D. Walczak. 2003. Dynamic Airline Revenue Management with Multipe Semi-Markov Demand. *Oper. Res.* 51, 137-148.

Chatwin, R.E. 1992. Optimal Airline Overbooking. Ph.D. Thesis, Stanford University, Palo Alto.

Chatwin, R.E. 1999. Continuous-Time Airline Overbooking with Time Dependent Fares and Refunds. *Transportation Science* **33** 182-191.

Couglan, J. 1999. Airline overbooking in the multi-class case. *Journal of the Operational Research Society,* 50, 1098-1103.

Curry, R.E. 1990. Optimal Airline Seat Allocation with Fare Classes Nested by Origins and Destinations. *Transportation Science* **24** 193-204.

Eren, S. and C. Maglaras. 2006. Revenue management heuristics under limited market information: A maximum entropy approach. presented at The 6th Annual INFORMS Revenue Management Conference, June 5-6, Columbia University.

Fox, S.1992 *Pricing and Rate Forecasting Using Broadcast Yield Management,* National Association of Broadcasters, Item 3555, Washington DC.

Geraghty, M.K. and E. Johnson. 1997. Revenue Management Saves National Car Rental. *Interfaces,* v.27, pp.107.

hotel-online.com. 2005. Choice Hotels International Signed 276 New Franchise Contracts in the First Six Months of 2005, Up 19% Over the Same Period in 2004. URL:
http://www.hotel-online.com/News/PR2005_3rd/July05_ChoiceMidYear.html

Huh, T. and Rusmevichientong, P. 2006. An Adaptive Algorithm for Multiple-Fare-Class Capacity Control Problems. Presented at the 6th Annual INFORMS Revenue Management Conference, June 5-6, Columbia University.

Hoseason, J. 2000 *"Capacity management in the cruise industry"* , In A. Ingold, I. Yeoman and U. McMahon-Beattie (Eds), Yield Management: Strategies for the Service Industries (2nd ed.) (pp.289C302). London: Thomson.

Jiang, Netessine and Savin, 2007, Robust Newsvendor Competition, Working paper.

Kasilingam,R.G. 1996, *"Air Cargo Revenue Management: Charactersitics and Complexities,"* EJOR,96,36-44

Kleywegt,Cooper.B 2006, *Pricing Dynamics of Competitors Whose Models Ignore Competition* 6th Annual INFORMS Revenue Management and Pricing Conference, New York.

Kunnumkal, S. and H. Topaloglu. 2007. *A stochastic approximation method for the revenue management problem on a single flight leg with discrete demand distributions.* Working paper, School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY.

Lan (2008), Robust Revenue Management: Theory and Experiments. PhD thesis, University of Maryland, Robert H. Smith School of Business, College Park, MD.

Lawrence, R.D., S.J. Hong and J. Cherrier. 2003. Passenger-based predictive modeling of airline no-show rates. *KDD '03: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining,* pp. 397-406.

Lahoti, A. 2002. Why CEOs Should Care About Revenue Management. *OR/MS Today*, February 2002, http://www.lionhrtpub.com/orms-2-02/rm.html.

Lee, T.C. and M. Hersh. 1993. A Model for Dynamic Airline Seat Inventory Control with Multiple Seat Bookings. *Trans. Sci.* 27, 252-265.

Leibs.S. 2001, Ford heeds the profits. *CFO Magazine* Aug.2001.

Lennon, J.J. 2004. Revenue management and customer forecasts: A bridge too far for the UK visitor attractions sector *J. Revenue and Pricing Manag.* 2, 338-352.

Lim, A.E.B. and J.G. Shanthikumar. 2006. Relative entropy, exponential utility, and robust dynamic pricing. forthcoming in *Operations Research.*

Li,M.Z.F and Qum, T.H.(1998). Seat allocation game on flights with two fares", Working Paper, Nangyang Business School, Nangyang Technological University, Singapore.

Littlewood, K. 1972. Forecasting and Control of Passengers. *12th AGIFORS Symposium Proceedings,* 95-128.

Liu, Smith, Orkin, Garey, 2002, Estimating unconstrained hotel demand based on censored booking data.*Journal of Revenue and pricing Management,* Vol.1, No.2, 1,pp. 121-1 38(18)

Oberwetter.R. 2001, Building blockbuster business: Can revenue managment land a starring role in the movie theater industry? *OR/MS Today, June 2001*

Pölt S 1998. "Forecasting is difficult-especially if it refers to the future" *In Reservations and Yield mangement Study Group Annual Meeting Proceddings* Melbourne, Australia.

Pölt S 1999. "Back to the roots: New results on leg optimization" *In Reservations and Yield mangement Study Group Symposium* London, UK.

Netessine,S. and Shumsky, R.A.(2005)."Revenue management games: horizontal and vertical competiton", *Management Science* **51** 813-831

Perakis, G. and G. Roels. 2006. Robust Controls for Network Revenue Management. Working paper, UCLA, Anderson School of Management, Los Angeles, CA.

Phillips, R. L. 2005. *Pricing and Revenue Optimization.* Stanford University Press, Stanford.

Ratliff, R.M. 1998. Ideas on Overbooking. Presentation at *AGIFORS Reservation and Yield Management Study Group Meeting,* Melbourne, Australia.

Robinson, L.W. 1995. Optimal and Approximate Control Policies for Airline Booking with Sequential Nonmonotonic Fare Classes. *Operations Research* **43** 252-263.

Rothstein, M. and A.W. Stone. 1967. Passenger Booking Levels. *Proceedings of the Seventh AGIFORS Symposium*, American Airlines, New York.

Rothstein, M. 1971. An Airline Overbooking Model. *Transportation Science* **5** 180-192.

Rothstein, M. 1974. Hotel Overbooking as a Markovian Sequential Decision Process. *Decision Sciences* **5** 389-404.

Rothstein, M. 1985. OR and the Airline Overbooking Problem. *Operations Research* **33** 237-248.

Rusmevichientong, P., B. Van Roy, and P. W. Glynn. 2006. A Non-Parametric Approach to Multi-Product Pricing. *Oper. Res.* 54, 82-98.

Shlifer, E. and Y. Vardi. 1975. An Airline Overbooking Policy. *Transportation Science* **9** 101-114.

Simon, J.L. 1968. An Almost-Practical Solution to Airline Overbooking. *Journal of Transport Economics and Policy* **II** 201-202.

Smith,B., J. Leimkuhler, R. Darrow and J. Samuels. 1992. Yield Management at American Airlines. *Interfaces* **22** 8-31.

Subramanian, J., S. Stidham Jr.and C. Lautenbacher. 1999. Airline Yield Management with Overbooking, Cancellations and No-Shows. *Transportation Science* **33** 147-167.

Talluri, K. and G. van Ryzin. 2004a. *The Theory and Practice of Revenue Management*, Kluwer Academic Publishers, Boston, 2004.

Talluri, K. and G. van Ryzin. 2004b. Revenue Management under a general discrete choice model of customer behavior. *Manag. Sci.* 50, 15-33.

Taylor, C.J. 1962. The Determination of Passenger Booking Levels. *Proceedings of the Second AGIFORS Symposium*, American Airlines, New York.

Thompson, H.R. 1961. Statistical Problems in Airline Reservation Control. *Operations Research Quarterly* **12** 167-185.

van Ryzin, G. J. and J. McGill. 2000. Revenue Management without Forecasting or Optimization: An Adaptive Algorithm for Determining Airline Seat Protection Levels. *Manag. Sci.* 46, 760-775.

Vives Xavier, 1999. Oligopoly Pricing Old Ideas and New Tools, MIT Press, Cambridge, Massachusetts

Wade, B. 1996. PRACTICAL TRAVELER;No-Show Fees: Well, Yes and No. *The New York Times,* March 17, 1996.

Weatherford, L.R. and S.E. Bodily. 1992. A Taxonomy and Research Overview of Perishable-Asset Revenue Management: Yield Management, Overbooking and Pricing. *Operations Research* **40** 831-844.

Yeoman, I. and A. Ingold. 1997. *Yield Management Strategies for the Service Industries.* Cassell, London.

Wollmer, R.D. 1992. An airline seat management model for single-leg route when low fare classes book first. *Oper. Res.* 40, 26-37.