Seat Inventory Control with Limited Demand Information

Yingjie Lan • Huina Gao

Robert H. Smith School of Business, University of Maryland, College Park, MD, USA

{ylan,hgao}@rhsmith.umd.edu

Faculty Advisors: Michael Ball, Itir Karaesmen ({mball,ikaraes}@rhsmith.umd.edu)

In this paper, we consider the classical *single resource (leg)* problem in revenue management for the case where demand information is limited. Our approach employs a competitive analysis, which guarantees a certain performance level under *all* possible input sequences where the only information available consists of lower and/or upper bounds on demand. We consider both competitive ratio and absolute regret performance criteria. We derive the best possible static policies whose booking limits remain constant throughout the booking horizon, under the various models we analyze. We prove that the optimal policies have the form of nested protection levels. We also show how dynamic policies, whose booking limits may be adjusted at any time based on the history of bookings, can be obtained. We provide extensive computational experiments and compare our methods to existing ones. The results of the experiments demonstrate the effectiveness of these new robust methods.

---

1 Introduction

Revenue management (RM) is one area of research and practice that has grown significantly in the past decade. To-date, the majority of the successful RM implementations rely heavily on use of demand information and experts advise against rushing to optimization in RM without first fully implementing accurate demand forecasting systems (Lahoti, 2002). Lennon (2004) mentions lack of data and naive forecasts based on inadequate data as two important factors limiting the realistic application of RM to a new industry. Even in industries such as airline 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, Weatherford and Pölt, 2002).

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 analyses 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). In this paper, we consider the classical *single resource (leg)* problem in RM and provide alternative forms of controlling the bookings when demand information is limited. We do not make any assumptions about the demand or the arrival process beyond upper and lower bounds, and do not need a risk neutrality assumption. Our approach relies on competitive analysis of online algorithms and we call the policies derived based on competitive analysis *robust* since they guarantee a certain performance level under *all* possible input sequences. We focus on the cases where the only information that is available is lower/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.

Our contributions can be summarized as follows. We propose new static and dynamic booking control policies for the single-leg, multiple-fare class problem when only upper/lower bounds on demand are available. We formulate two different problems for the decision-maker; these problems differ in their objectives where the first one *maximizes the competitive ratio* and the second *minimizes the maximum absolute regret.* We show that these two problems can be analyzed in a unified manner and we obtain *closed-form solutions* for both. The resulting optimal booking control policies preserve the *nesting* structure which is widely used in practice. We also show that the optimal booking limits obtained by the absolute regret criterion are no more than the ones obtained by competitive ratio. Using computational experiments, we show that the average revenues obtained by our policies in simulation studies are comparable to other well-known procedures even though our policies use less information. Further, our policies are robust and not prone to errors in modelling demand. Finally, our work builds on prior research by Ball and Queyranne (2006) who used competitive analysis of online algorithms to analyze the single-leg problem and assumed no demand information is available. Our policies, which assume demand bounds, generalize theirs, and we also provide alternative proofs for some of their results.

The rest of the paper is organized as follows: We provide a literature review in Section 2. Section 3 introduces the notation and describes our approach. In Section 4, we derive static nested booking control policies when the objective is to maximize the competitive ratio. Section 5 extends the analysis to the problem with the absolute regret criterion. We show how dynamic policies that improve the worst-case performance can be obtained in Section 6. Section 7 presents an extensive set of computational results and Section 8 provides conclusions and suggestions for further research.

2 Literature Review

The booking (or capacity) control problem for a single resource has long been studied in the RM literature. Littlewood (1972) was the first to analyze the stochastic, single-leg problem with two fare-classes and assumed product is sold in a low-before-high (LBH) manner; i.e., demand in lowest fare class arrives first. Belobaba (1987,1989) discussed heuristic extensions of Littlewood's 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. 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) discuss the static (i.e., arrivals are LBH) and dynamic problems (i.e., arrival sequence is not ordered by fare-class) 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. 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).

Until very recently, the adaptive method proposed by van Ryzin and McGill (2000) for the single-leg, multi-fare problem was the only work in RM that relied on no demand information. This approach updates booking limits iteratively - from flight to flight - based on past observations of *fill events*; i.e., whether a particular booking limit was reached or not. van Ryzin and McGill prove that the nested booking limits computed using this adaptive procedure converge to the optimal. More recently, Huh and Rusmevichientong (2006) suggested a more sophisticated adaptive method for the same problem. Both of these methods require LBH arrivals, suffer from poor performance during the transient period until the booking limits converge to the optimal, and the results can be sensitive to initial point and step-sizes.

The objective in all the above papers is to maximize revenues given the risk-neutrality of the decision-maker. Traditional assumptions on demand models, availability of demand information and also risk-neutrality have been relaxed in recent years, especially in the dynamic pricing context (e.g. Rusmevitchientong et al., 2006). In this new stream of research, ongoing work of Birbil et al. (2006), Eren and Maglaras (2006), Perakis and Roels (2006), and Ball and Queyranne (2006) are the only ones - to the best of our knowledge - that focus on seat inventory control in the single-leg RM problem. Birbil et al. (2006) assume there are inaccuracies associated with discrete probability distributions that characterize the demand in each fare class. They propose a robust framework where the probability value associated with the total demand of each fare-class lies in a polyhedral set (as opposed to the realized demand lying in a polyhedral set in our model and that of Perakis and Roels (2006)). Eren and Maglaras (2006) propose using the maximum entropy approach to update the booking limits while obtaining demand information. This approach makes use of past observations and accordingly adds fractile constraints to a convex optimization problem that is used to obtain an estimate of a discrete probability distribution. That estimate is then used to determine the booking limits, and this forecasting/optimization cycle is repeated throughout the booking horizon. Perakis and Roels (2006) assume limited demand information similar to ours and provide a general approach to both the single-leg and the network RM problems. They consider two objectives: Maximizing the minimum revenues and minimizing the maximum regret. They also restrict their analysis to particular policies (e.g., nested policies for single-leg, partitioned allocations for network). In contrast, we provide a unified framework for both absolute regret and competitive ratio criteria in the single-leg problem. Our analysis inspired by competitive analysis of online algorithms is significantly different than theirs. We provide *closed-form optimal solutions* and prove the optimality of the nested booking limits in our problem.

Ball and Queyranne (2006), which we refer as BQ, were the first to adapt the notion of *competitive analysis of online algorithms* in 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. Our approach is more general than theirs in that we assume availability of limited demand information and show the analysis extends to problems with *both* relative and absolute regret criteria. Our optimal policies reduce to theirs in special cases, hence we obtain alternative proofs for some of their results as well. Finally, our policies are practically more effective compared to BQ as we show in our computational experiments.

3 Problem Definition

We consider the single-leg RM problem from the perspective of *competitive analysis of online algorithms* (see the survey in Albers, 2003). This perspective evaluates the performance of a booking control policy relative to the performance of an offline algorithm that considers the entire input sequence simultaneously. An offline optimal solution is a solution obtained by an offline algorithm (with hindsight) that optimizes the objective function of interest. We have two performance metrics of interest: *absolute regret* (AR), which is the difference between the objective function values of the offline and online algorithms, and *relative regret*, which is the ratio of the absolute regret to the objective function value of the offline algorithm. Specifically, we are interested in having a guarantee on these measures under all input sequences. Based on the definition of robustness in Kouvelis and Yu (1997), optimal booking control policies obtained using these two performance metrics will yield *deviation robust* and *relative robust* decisions.

In competitive analysis, competitive ratio (CR) is commonly used as a measure of an algorithm's effectiveness. 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 $\tilde{v}_\Upsilon(I)$ be the objective value achieved by the online algorithm for input I and let $\tilde{v}^*(I)$ be the objective value achieved by an optimal offline algorithm. Then, CR is defined as:

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

Clearly, an algorithm/policy that maximizes the CR, minimizes the maximum relative regret. Likewise, the maximum absolute regret (MAR) of the online algorithm is defined as

$$\text{MAR of } \Upsilon \;\; = \;\; \sup_{I \in \Omega_\Upsilon} \{\tilde{v}^*(I) - \tilde{v}_\Upsilon(I)\}.$$

We use CR - as opposed to relative regret - and MAR in our analysis. Note that these two performance measures as defined above apply to "deterministic" algorithms, i.e., the algorithm applies the same decision rule and yields the same performance given the same input sequence as opposed to a "randomized" algorithm. While not as practical as deterministic ones, randomized algorithms are of special interest from a technical standpoint, and are analyzed in the doctoral dissertation of the first author (Lan, 2007). For now, our goal is to determine the best deterministic algorithm that maximizes the CR (or minimizes the MAR).

Throughout this paper, let n denote the total capacity of the resource (seats, rooms, etc.) available and $m \geq 2$ the number of fare-classes. Let f_i denote the fare for class i, where $f_1 > f_2 > \ldots > f_m \geq 0$. There are no restrictions on the demand or arrival process except that requests arrive one-by-one and cancellations are not allowed. In our analysis, we assume each request demands only one unit but this can be generalized as we discuss below. We are interested in determining nested booking limits $n \geq b_1 \geq b_2 \geq \cdot \geq b_m \geq 0$ (or nested protection levels $\theta_1 \leq \theta_2 \leq \cdot \leq \theta_m$) that maximize the CR or minimize the MAR. Note that *nesting by revenue order* remains the core of booking control in many traditional RM implementations. Although Talluri and van Ryzin (2004b) discuss that these types of nested policies may not be optimal in general, we show later in the text that they are optimal relative to the CR and MAR criteria.

The booking limit b_i defines the maximum number of booking requests to be accepted in classes k for $k \geq i$. Similarly, the protection level θ_i is the number of seats protected for classes $1, ..., i$ from classes k for $k \geq i+1$. By definition, $\theta_i + b_{i+1} = \theta_m = b_1$ for $i = 1, ..., m-1$. We use an additional variable called *bucket size* - denoted x_i - in our analysis; we define it as $x_i = b_i - b_{i+1}$. The notation b, θ and x are used for vectors of variables. Note any of these vectors is sufficient to characterize a nested booking control policy.

In simple terms, competitive analysis of online algorithms is based on worst-case analysis: One can think there is an adversary in charge of generating booking requests. The adversary is aware of the 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 RR or highest AR). Given the characterization of input sequences the adversary would choose, we determine the optimal parameters for the algorithm to maximize the CR (alternatively, to minimize the MAR).

By nature of the single-leg RM problem, the demand is discrete, although in many research papers, the analysis is done assuming continuity. In the *continuous problem*, the demands, and therefore the protection levels as well, may be any nonnegative real numbers. In this case, any request in a fare-class is *partially accepted* (split). Although the continuous problem is less realistic, its analysis is simpler. In the *discrete problem*, the demands are integral and the protection levels are restricted to being integer. Our analysis of the continuous problem constitutes a crucial stepping-stone for the discrete problem. By rounding the continuous solution and/or using randomized policies, effective integer protection levels can be obtained. Discussion and analysis of such discrete policies are discussed in the doctoral dissertation of the first author (see Lan, 2007). In the rest of the paper, our decision variables take continuous values, and we assume - without loss of generality - input sequences consist of requests that demand one unit of a fare-class each. Note that the continuous case easily generalizes to multi-unit and multi-fare requests (as in batch arrivals) by allowing splitting of requests. Our analysis carries through as long as each request demands a non-negative and finite amount. Note also that group bookings cannot be enforced/guaranteed in this case even though batch arrivals are permitted.

4 Optimal Static Policies for the Competitive Ratio Problem

In this section, we analyze the multi-fare single-leg problem and derive optimal booking control policies under the CR criterion assuming upper and lower bounds are available on demand in each fare class. That is, for each fare class i, we assume the input sequences, I, are restricted so that the total number of units demanded in class i falls between L_i and U_i where $0 \leq L_i \leq U_i$ for $i = 1, ..., m$.[1] We use the notation L and U to denote the respective vectors $(L_1, ..., L_m)$ and $(U_1, ..., U_m)$. Let $\Omega(L, U)$ be the set of all sequences where the total demand in each fare-class falls between the upper and lower bounds, $R^*(I)$ be the offline revenue obtained from sequence I, and $R(I; b)$ be the online revenue gained by a standard

[1] It is typical in the robust optimization framework to assume uncertain parameters belong to a polyhedral set. For instance, Bertsimas and Sims, 2004, assume parameters lie in a known interval centered at a nominal value. Similarly, we assume only "range forecasts" are given and demand lies in a given interval.

nested booking limit policy b. The following formulation, taking demand bounds into consideration, solves for the optimal booking control policy that maximizes the CR:

$$z^{CR} = \max_{b} z : z \leq \frac{R(I;b)}{R^*(I)}, \forall I \in \Omega(L,U) \tag{1}$$

where b is the vector of decision variables and z^{CR} is the optimal CR. There are some potential challenges in solving this maximization problem: 1) $R(I;b)$ might not have closed-form expression; 2) $|\Omega(L,U)|$ grows exponentially with n and m, prohibiting any serious direct attempt with even small problems. However, since there will be redundant constraints, which means the corresponding sequences are redundant as well, a sequence reduction approach seems natural, and, as we will show in the next two sections, this maximization problem can be rewritten in a very compact form. The solution to the above problem will result in a static control policy.

4.1 Sequence Reduction

A sequence consists of a finite stream of fare requests during the booking horizon. Since each request in a sequence demands one unit, we can characterize the sequences based on the fare-classes and units demanded: Let $|I|$ be the total number of requests in the sequence and $I[j]$ be the total number of requests in fare class j. A *profile* of I is an m-dimensional vector of $[I] = (I[j] : j = 1, \cdots, m)$. By definition, $|I| = \sum_{j=1}^{m} I[j]$.

Observe that offline optimal revenue realized in sequence I *only* depends on the profile, while online revenue (of a policy) also depends on the arrival order of requests. However, given a profile, we can ignore permutations of requests that do not yield the lowest online revenue. The proof of the next result is in Lan and Gao (2007), Appendix A.

Proposition 1 *(Continuous Problem) Given all input sequences with the same profile, a continuous, nested booking limit policy b generates the least revenue when applied to the unique LBH sequence with that profile.*

From the standpoint of the maximization problem in (1), all non-LBH sequences can be discarded based on Proposition 1. In the remaining sequences, there is a one-to-one relationship between sequences and profiles. This yields a substantial reduction in the size of the problem. Yet, the total number of profiles can be prohibitive. Therefore, we pursue further sequence reductions to define a tractable problem. We now introduce a categorization of input sequences, which is even a broader concept than profiles. The ultimate goal is to choose only one sequence in each category. For mathematical completeness, we now introduce a virtual fare-class $m+1$ where $b_{m+1} = 0$, $f_{m+1} = 0$, and $L_{m+1}, U_{m+1} \geq 0$ are arbitrary, non-negative numbers.

Definition 1 *(Input Sequence Category) An input sequence category A_j^b is the set of sequences such that an input sequence I belongs to category A_j^b under policy b, if and only if j is the lowest fare-class whose booking limit is reached after executing b on the sequence I, that is, $j = \min\{i : \hat{B}_i(I;b) = b_i\}$ where $\hat{B}_i(I;b)$ is the total number of requests accepted in fare classes i through m after policy b is applied to sequence I.*

Note that the virtual class has $b_{m+1} = 0$, so at least the booking limit of class $m+1$ is always reached, and j is well defined. By construction, A_{m+1}^b includes all sequences for which no request is rejected and $\frac{R(I;b)}{R^*(I)} = 1$ for all $I \in A_{m+1}^b$. Note also that some categories may be empty. For instance, in the trivial case of $b_i = n$ for all $i = 1, ..., m$ and $\sum_{i=1}^m U_i < n$, the only category that is non-empty is A_{m+1}^b. On the other hand, one can see that appending more requests in classes $k \geq j$ to a sequence in category j will not change the online revenue, since all such requests would be rejected. In addition, any permutation of the order of requests in a sequence will not change its category, which means the category is totally determined by the profile of the sequence. Finally, this categorization is complete, meaning $\cup_{j=1}^{m+1} A_j^b = \Omega(L, U)$. Note the notion of a category applies to both discrete and continuous problems, and so does our next result. Please see Appendix B in Lan and Gao (2007) for the proof.

Proposition 2 *(CAtegory-dominant-STreams, or CAST) Consider a nested booking limit policy b, all input sequences in some category A_j^b, and a sequence $CAST_k \in A_j^b$ such that $CAST_k$ is the LBH sequence with the profile of $CAST_k[j] = U_j$ for $j \geq k$ and $CAST_k[j] = L_j$ for $j < k$. $CAST_k$ dominates the other sequences in category A_j^b, i.e., $\frac{R(I;b)}{R^*(I)} \geq \frac{R(CAST_k;b)}{R^*(CAST_k)}$ for all $I \in A_j^b$.*

For any nested policy, Proposition 2 effectively reduces the number of sequences to be considered, down to m, disregarding the virtual class $m+1$ (which is not critical since $\frac{R(I;b)}{R^*(I)} = 1$ for all $I \in A_{m+1}^b$). With these reductions, problem (1) can be reformulated as

$$z^{CR} \quad = \quad \max_b \; z : z \leq \frac{R(CAST_k; b)}{R^*(CAST_k)}, \quad for \quad k = 1, ..., m. \qquad (2)$$

We next show how the optimal solution to the reduced formulation can be obtained by defining an appropriate linear programming model. We first focus on the continuous problem.

4.2 A Linear Programming Formulation for the Continuous Problem

In determining the optimal booking control policy, we use the bucket sizes $x_i = b_i - b_{i+1}$ ($i = 1, ..., m$) as the decision variables. First consider the case where $\sum_{i=1}^{m} U_i \leq n$. The optimal bucket sizes in this case are $x_i^* = U_i, \ \forall i$ and $z^{CR} = 1$. We assume $\sum_{i=1}^{m} U_i > n$ to exclude this trivial case from our analysis.

To express the term in the righthand side of the constraints in problem (2), we study the properties of nested policies under the CAST sequences. Note that any reasonable policy would have $x_1 \leq U_1$ so as not to protect any seats for class-1 that would definitely remain unsold under CAST sequences. Likewise, $x_i = b_i - b_{i+1} \leq U_i$ for all $i = 1, .., m$; otherwise there would be a class $k \geq 1$ whose bucket would have a "slack" (i.e., the booking limit of class k would not be reached). Naturally, we require that $\sum_{i=1}^{m} x_i \leq n$ so as not to exceed the capacity.

Let the number of class-j arrivals in sequence $CAST_k$ be denoted by T_j^k; we have $T_j^k = U_j$ for $j \geq k$ and $T_j^k = L_j$ for $j < k$ by definition. We further define $\rho(k) = \max\{j : \sum_{i=1}^{j} T_i^k \leq n\}$. Then, we can express $R^*(CAST_k)$, the offline optimal revenue for $CAST_k$, as R_k^* where

$$R_k^* = \left(n - \sum_{j=1}^{\rho(k)} T_j^k\right) f_{\rho(k)+1} + \sum_{j=1}^{\rho(k)} T_j^k f_j.$$

Next, we focus on the online revenue obtained from CAST sequences. For any non-trivial m-fare problem, the online revenue generated from $CAST_k$ can be considered in two parts: 1) from the L_i requests in classes $i < k$, with a subtotal revenue *at most* $R_k^+ = \sum_{i=1}^{k-1} f_i L_i$ (because not all L_i may be accepted based on the bucket sizes), and 2) from the U_i requests in classes $i \geq k$, from which revenue is $\sum_{i=k}^{m} x_i f_i$ (note $x_i \leq U_i$, from our discussion above). Then, $R(CAST_k; b)$, the online revenue obtained by policy b and corresponding bucket sizes x satisfy $R(CAST_k; b) \leq R_k^+ + \sum_{i=k}^{m} f_i x_i$. Therefore, we have an upper bound on the ratio of online revenues to that of offline, and this upper bound is linear in the bucket sizes x_j:

$$\frac{R(CAST_k; b)}{R^*(CAST_k)} \quad \leq \quad \frac{R_k^+ + \sum_{j=k}^{m} f_j x_j}{R_k^*}.$$

Combining these observations, we can formulate a linear program (LP) that finds a continuous, nested policy. We call this formulation the *General Bounded Model (GBM)*:

$$\textbf{GBM}: \qquad \bar{z}^{CR} \quad = \quad \max_{x} \quad z$$

$$\text{s.t.} \qquad R_j^* z \leq R_j^+ + \sum_{i=j}^{m} f_i x_i, \quad j = 1, ..., m \tag{3}$$

$$\sum_{i=1}^{m} x_i \leq n, \tag{4}$$

$$0 \leq x_j \leq U_j, \quad j = 1, ..., m. \tag{5}$$

A small numerical example to illustrate the GBM (and its solution) is provided in Appendix E of Lan and Gao (2007). Although GBM has been developed based on a necessary set of constraints and relations that constitute an *upper bound* on the optimal CR, the solution to GBM provides the optimal nested policy for the CR problem introduced in (1). We formalize this statement and show how a closed-form solution to GBM can be derived in the next section.

Note that the actual CR problem expressed in (1) is not only reduced to a problem with a small number of constraints based on CAST sequences, but its optimal policy in the continuous version is obtained by solving a LP with $2m + 1$ constraints. We can further make the following observations regarding the optimal solution of GBM (hence the optimal policy): For any $U_i > n$, the upper bound in constraint (5) is redundant. Therefore, demand information is most critical in determining the optimal policy when $U_i \leq n$ for a class i. Also, the optimal CR is non-decreasing in U_i and non-increasing in L_i for class i. Finally, the multi-fare continuous booking problem discussed in BQ can be represented using GBM. Theirs is a special case: They use no information about the demand, hence $L_i = 0$ and $U_i = \infty$ (or effectively $U_i \geq n$) for $i = 1, ..., m$ in their problem.

4.3 Optimal Seat Inventory Control Policy: Solution to GBM

Before we present the closed form solution to GBM and the CR problem, we make the following observation: The lower bounds imply that all sequences have at least L_i requests of class i. Let $N_j = [n - \sum_{i=1}^{j-1} L_i]^+$, $j = 2, ..., m$ denote the remaining number of seats after considering the requests for classes $j-1$ and lower. Any reasonable policy should have $\sum_{i=j}^{m} x_i = b_j \leq N_j$ for all $j = 1, ..., m$ so as never to accept a lower fare request at the expense of a guaranteed higher fare request.[2] Thus, when $N_k = 0$, for some class k, we must have $b_j = 0$ for all $j \geq k$ and classes $j \geq k$ can be simply ignored. We exclude such trivial cases from our analysis, by assuming $N_m = n - \sum_{j=1}^{m-1} L_j > 0$.

[2]These constraints can be added to GBM but are not binding in the optimal solution.

The analysis of GBM relies heavily on the investigation of the intrinsic structure of the problem and the relationship among the parameters $R_i^*, R_i^+, f_i, N_i, U_i$ and L_i; details are provided in Appendix C of Lan and Gao (2007). By definition of R_i^+, we can see that $L_i = (R_{i+1}^+ - R_i^+)/f_i$. The analysis in the Appendix shows that:

$$\sum_{i=1}^{m}(g_i + L_i) \geq n \qquad \text{and} \qquad g_i + L_i \leq U_i, \quad i = 2, \cdots, m \tag{6}$$

where

$$g_i \quad = \quad \frac{R_i^* - R_{i+1}^*}{f_i} \quad i = 1, \cdots, m \tag{7}$$

with $R_{m+1}^* = 0$. Recall that we have $\sum_{i=1}^{m} U_i > n$ to exclude trivial cases. Consequently, conditions in (6) indicate that each g_i is "locally" a feasible value for x_i but, in aggregate, they exceed total capacity. With these boundary conditions, we can derive a closed-form solution to GBM. The proof of the next result is in Lan and Gao (2007), Appendix D.

Proposition 3 *(a) The optimal solution of GBM is*

$$\bar{z}^{CR} \quad = \quad \frac{R_u^+/f_u + N_u}{R_u^*/f_u + \sum_{i=1}^{u-1} g_i} \tag{8}$$

$$x_j^{CR} \quad = \quad \begin{cases} g_j z^{CR} + L_j & j < u \\ (R_u^* z^{CR} - R_u^+)/f_u & j = u \\ 0 & j > u \end{cases} \tag{9}$$

$$u \quad = \quad \max\{j : R_j^+ \sum_{i=1}^{j-1} g_i < N_j R_j^*\} \tag{10}$$

where the index u denotes the critical fare-class such that all classes $k > u$ are closed.
(b) The nested protection levels defined by

$$\theta_i^{CR} \quad = \quad \sum_{j=1}^{i} x_i^{CR} \text{ for } i = 1, ..., m-1 \tag{11}$$

maximizes the CR in problem (1) and the optimal CR is $z^{CR} = \bar{z}^{CR}$.

Note that our analysis so far yields the optimal solution within the class of nested policies. We used reductions in the number of possible input sequences and formulated the GBM based on the properties of nesting. One question that remains to be answered is whether the nested policies are the best. The next result shows that, in fact, this is true. Please see Lan and Gao (2007), Appendix F for the proof.

Proposition 4 *For the continuous m-fare problem with demand bounds, no deterministic policy can achieve a CR greater than $\bar{z}^{CR}$.*

Our main result combines the analysis done so far:

Theorem 1 *For the continuous m-fare problem with demand bounds, the nested booking control policy with protection level vector θ^{CR} defined by (11) has a CR of $\bar{z}^{CR}$ given by (8) and this is the best possible among deterministic policies.*

Extension of our analysis to randomized policies reveals the following stronger result. Discussion and proof are available in Lan (2007).

Theorem 2 *For the continuous m-fare problem with demand bounds, no randomized booking policy has a CR larger than $\bar{z}^{CR}$ given in (8). Therefore, the deterministic nested booking control policy with protection levels defined by θ^{CR} in (11) is the best possible among all policies.*

In the case all lower bounds are zero, the above result can be simplified, and the optimal protection levels can be expressed in terms of n and g_j for $j = 1, ..., m$. See Appendix G of Lan and Gao (2007) for the derivation. For the special case of $L_j = 0, U_j \geq n$ for all $j = 1, ..., m$, that result reduces to $R_j^* = f_j n$, and $g_i = n(1 - f_{i+1}/f_i)$. This, in fact, is the optimal nested policy obtained by BQ.

5 Optimal Static Policies for the Absolute Regret Problem

The analysis in Section 4 easily extends to the problem with the absolute regret criterion where the objective is to minimize the MAR. Let us first express this problem in its general form:

$$z^{AR} \quad = \quad \min_b z : z \geq R^*(I) - R(I;b), \forall I \in \Omega(L,U) \qquad (12)$$

where b is the vector of decision variables and z^{AR} is the optimal MAR. We can reformulate this problem by reducing the number of input sequences using Propositions **??** to 2 and we only have to consider the CAST sequences. The following LP model (called GBM-AR) is designed to provide a *lower bound* on the minimum MAR in the continuous problem:

$$\mathbf{GBM-AR:} \quad \bar{z}^{AR} \; = \; \min \quad z$$

$$\text{s.t.} \qquad R_j^* - z \leq R_j^+ + \sum_{i=j}^{m} f_i x_i, \quad j = 1, ..., m \tag{13}$$
$$\text{and} \quad (4), (5).$$

This model is actually easier to analyze than the GBM model, and the closed-form solution is given below (proof is omitted):

$$\bar{z}^{AR} = R_{\tilde{u}}^* - R_{\tilde{u}}^+ - f_{\tilde{u}} x_{\tilde{u}}^{AR} \tag{14}$$
$$x_j^{AR} = \begin{cases} g_j + L_j & j < \tilde{u} \\ N_{\tilde{u}} - \sum_{i<\tilde{u}} g_i & j = \tilde{u} \\ 0 & j > \tilde{u} \end{cases} \tag{15}$$
$$\tilde{u} = \max\{j : \sum_{i<j} g_i < N_j\} \tag{16}$$

where $\tilde{u}$ is the critical fare-class such that classes $k > \tilde{u}$ are closed. A small numerical example to illustrate GBM-AR and its solution is provided in Appendix E in Lan and Gao (2007).

Following a similar argument for the CR problem, we can show that the optimal solution of GBM-AR provides the optimal MAR of the problem in (12)[3]. In addition, the optimal nested policy obtained by GBM-AR is the best possible under the MAR criteria (proof is omitted). Hence, nesting leads to both *deviation robust* and *absolute robust* decisions in the single-leg RM problem with limited demand information.

Theorem 3 *For the continuous m-fare problem with demand bounds, the nested booking control policy defined by the protection levels*

$$\theta_i^{AR} = \sum_{j=1}^{i} x_i^{AR} \text{ for } i = 1, ..., m-1 \tag{17}$$

has minimum MAR of $\bar{z}^{AR}$ (where x_i^{AR} and $\bar{z}^{AR}$ are given in (15) and (14)). No other policy, deterministic or randomized, has a lower MAR.

A close look at the optimal policies of CR and MAR problem reveals that $u \geq \tilde{u}$. Therefore, MAR tends to deny bookings to a higher number of fare-classes. Furthermore, $x_i^{CR} \leq x_i^{AR}$ for $i < \tilde{u} \leq u$ and $\sum_{i=1}^{u} x_i^{CR} = \sum_{i=1}^{\tilde{u}} x_i^{AR} = n$ at optimality. Our next result combines these observations:

[3]Perakis and Roels (2006) develop a LP model, which is slightly different than ours, and obtain the same closed-form solution for the MAR problem. Their analysis of the MAR problem is significantly different, and they do not present a formal proof of optimality of nesting.

Proposition 5 *In the continuous problem with demand bounds, the optimal nested protection levels obtained by (11) and (17) satisfy:* $\theta_i^{AR} \geq \theta_i^{CR}$ *for* $i = 1, ..., m$.

While competitive analysis of algorithms using CR and MAR criteria provides conservative solutions to guarantee worst-case performance, the above result shows that MAR criteria is more aggressive in protecting seats for classes with higher fares (i.e., optimal booking limits of MAR are lower). Note that there is no easy comparison of the revenues obtained by CR or MAR solutions in general: The revenues obtained by the optimal policy of MAR in the two-fare problem can be lower (higher) than that of the optimal policy of CR under LBH input sequences when the demand for class-1 is low (high).

6 Dynamic Policies

In this section, we show how dynamically adjusting a static policy can improve the CR (or MAR) in our problem. Our analysis extends the discussion in BQ (who provide a direction for the CR problem only for $m = 2$) to multiple fare classes with demand bounds. Using the imaginary adversary paradigm, the CR (or MAR) of a policy effectively assumes that the adversary always sticks to an "optimal" strategy (i.e., adversary sends LBH inputs). On the other hand, if the adversary "makes a mistake", we can consider a new problem scenario based on the remaining capacity and requests accepted so far, create a new policy for the remaining requests in order to guarantee better overall performance.

Consider a dynamic scenario where a dynamic policy has been executing so that a partial input sequence has already been processed. Suppose h_i bookings have been accepted for fare class i for $i = 1, \cdots, m$ by processing the partial sequence I_0. This accumulates revenue of $R_0 = \sum_{i=1}^m h_i f_i$ from the $\sum_{i=1}^m h_i$ sold seats. The question is whether the booking limits can be adjusted to improve the CR (MAR) achievable under a future input sequence $\hat{I}$, which yields the complete input sequence of $I = I_0\hat{I}$. The concatenation of I_0 and $\hat{I}$ into the complete input sequence I produces the following profile: $L_j \leq I[j] = I_0[j] + \hat{I}[j] \leq U_j, \;\; j = 1, \cdots, m$. Since $I_0[j] \geq h_j, j = 1, \cdots, m$, it follows that $I[j] \geq h_j$ must hold, so that we can update lower bounds for $\hat{I}$ as follows:

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

Let $\hat{n} = n - \sum_{i=1}^m h_i$ denote the remaining number of available seats, $\hat{b}$ the nested booking limits for allocating the remaining $\hat{n}$ seats, and $\hat{x}$ the corresponding bucket sizes. That is,

$0 = \hat{b}_{m+1} \leq \hat{b}_m \leq \cdots \leq \hat{b}_1 \leq \hat{n}$, $\hat{x}_j = \hat{b}_j - \hat{b}_{j+1}$, $j = 1, \cdots, m$.

Given a partial policy $\hat{b}$ and partial input sequences $I_0, \hat{I}$, the overall CR can be improved by solving

$$\hat{z}^{CR} = \max_{\hat{b}} z : z \leq \frac{R_0 + R(\hat{b}; \hat{I})}{R^*(I_0 \hat{I})}, \quad \forall \hat{I} \in \Omega(\hat{L}, \hat{U}). \tag{18}$$

where $\hat{z}^{CR}$ is the new guarantee on the worst-case CR performance. Similarly, the improved minimum MAR of $\hat{z}^{AR}$ is determined by:

$$\hat{z}^{AR} = \max_{\hat{b}} z : R^*(I_0 \hat{I}) - z \leq R_0 + R(\hat{b}; \hat{I}), \quad \forall \hat{I} \in \Omega(\hat{L}, \hat{U}). \tag{19}$$

Applying proposition 2 to the partial $\hat{I}$ sequences, we only have to focus on m sequences that dominate the others in $\Omega(\hat{L}, \hat{U})$. Details are omitted. We denote these new sequences $CAST'_k$ for $k = 1, ..m$. Following exactly the same steps in Sections 4.2 and 5, the optimal CR (MAR) can be computed solving a LP. This LP has the same structure as GBM (GMB-AR) except for its parameters: The new policies for the CR and MAR problems are obtained by equations (8)-(10) and (14)-(16), respectively, by replacing L_i by $\hat{L}_i$ and defining $\hat{R}^*_j = R^*(I_0 CAST'_j)$, $\hat{R}^+_j = R_0 + \sum_{i=1}^{j-1} \hat{L}_i f_i$, $\hat{N}_j = \hat{n} - \sum_{i=1}^{j-1} \hat{L}_i$, $\hat{g}_j = (\hat{R}^*_j - \hat{R}^*_{j+1})/f_j$.

One key observation pertaining our dynamic policies is the following: Static and dynamic policies are identical under LBH input sequences when $m = 2$ but not when $m > 2$. This is because the only 'mistake' the adversary can make in LBH sequences is in the total amount of requests of each class. When $m = 2$, there is no room for improvement by adjusting the booking limit of class-2 in case of a 'mistake' because the input is LBH.

These dynamic policies are easy to implement; specifically, the h_j should be updated each time a request is accepted, LPs are re-solved and the booking limits adjusted accordingly. These tasks require minimal computational resources as we have closed-form optimal solutions of the LPs.

7 Computational Results

While the policies obtained above come with performance guarantees, such worst-case guarantees are no indication of their effectiveness in practical situations. We have designed computational experiments (i) to quantify the performance of policies obtained by CR and MAR criteria, (ii) to illustrate the differences between static and dynamic policies, (iii) to compare our policies to other well-known procedures in single-leg RM, and (iv) to show the robustness of our policies under various demand scenarios and parameter settings.

7.1 Performance of our policies: CR vs. MAR, Static vs. Dynamic

In this section, our goal is to test the effectiveness of CR vs. MAR and static vs. dynamic policies. The policies that use demand bounds are denoted BSTAT-CR, BSTAT-AR, BDYN-CR and BDYN-AR, and are obtained by solving GBM, GBM-AR, dynamically adjusted GBM and dynamically adjusted GBM-AR, respectively. The policies STAT-CR, STAT-AR, DYN-CR, DYN-AR are special cases of these policies where no demand information is used (effectively, $L_i = 0, U_i \geq n$ in optimization). We only implement the deterministic versions of our algorithms, and use continuous booking limits; randomized algorithms are excluded from our computational experiments. We use the standard-nesting implementation of the nested policies. [4]

The demand is uniform distributed, $m = 2$, and $n = 100$ in all the examples in this section. We use simulation to study the (experimental) performance of the policies. We vary the fare values, demand parameters, arrival regimes with respect to fare classes (LBH vs. time-homogeneous arrivals[5]), and the demand-mix (mean demand of a fare class relative to the demand of other classes) in our experiments. Requests arrive one-by-one in all the experiments. We also use FCFS and OFFLINE policies as benchmarks; FCFS accepts all arrivals up to the capacity and OFFLINE computes the hindsight optimal solution and revenue at the end of each simulation run.

Example-1a. The fares are $(f_1, f_2) = (500, 100)$, and demand in each fare class is discrete uniform distributed 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 (computed by taking the ratio of policy revenues to OFFLINE in each simulation run, and averaging this ratio across all simulation runs) of policies, and the average number of seats sold are displayed in Table 1. The static and

[4]Nested policies can be implemented in two different ways: Standard- and theft-nesting (see Talluri and van Ryzin, 2004a). We only use standard nesting in our experiments. Discussion of theft-nesting and its implications for our problem are available in the doctoral dissertation of the second author (Gao, 2007).

[5]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 to distribute the arrivals in each fare class uniformly over the entire booking horizon. The corresponding arrival times determine the sequence of arrivals in the time-homogeneous case.

Policy	*Protection level of class-1* θ_1	*Theoretical CR (%)*	*Average revenues*	*Average of ratio of policy revenues to OFFLINE (%)*	*Average number of seats sold*
BSTAT-CR, BDYN-CR	68.49	89.04	129,235	95.37	89.3
BSTAT-AR, BDYN-AR	72	87.69	129,501	95.28	86.93
STAT-CR, DYN-CR	44.5	66.19	114,302	85.84	98.79
STAT-AR, 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 1: The protection level, theoretical CR, and average performance of policies in Example-1a.

dynamic policies are equivalent in this example because of LBH arrivals. Note that the differences between the protection levels of CR and MAR 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 to 30 samples of size 200 each, we compute the percentiles in each sample, and we take the averages of the percentiles of the samples to obtain the estimate. This information is provided in Figure 1. 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 MAR policies and also the effect of demand information. Because MAR 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.

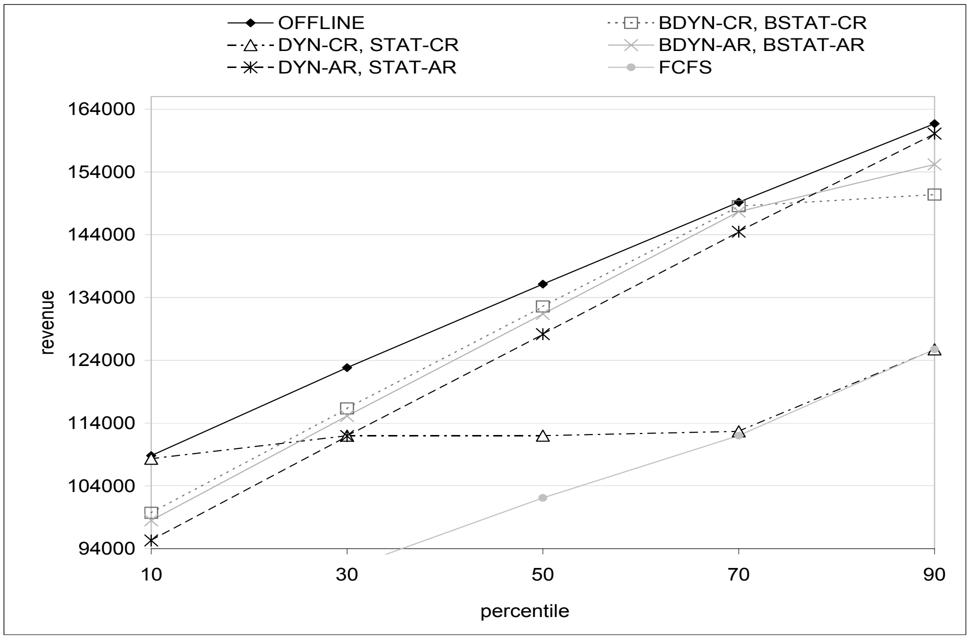

Figure 1: Distribution of revenues for each policy in Example-1a

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).

Example-2a. We next investigate the effect of the fare ratio $r = f_2/f_1$. The demand has the same distribution as in Example-1a. We vary the fare ratio from 0.1 to 1. The arrivals are LBH. Figure 2 shows how the average of ratio of policy revenue to OFFLINE revenue change with r. Notice that in this case the CR and MAR policies result in similar average relative performance when demand information is used. Use of demand information not only improves average relative performance but it also makes the performance of policies less sensitive to changes in the fare ratio r. Lower fare ratio favors MAR policies over CR when there is no demand information: STAT-AR is significantly better than STAT-CR for $r = 0.1$; this is expected because STAT-AR accepts more class-1 requests than STAT-CR.

We repeated Example-1a and Example-2a with different demand parameters and arrival regimes; see Example-1b through Example-1f and Example-2b in Appendix H of Lan and Gao (2007). Our observations so far can be summarized as follows: (1) CR and static policies are better in terms of the guaranteed minimum revenue, except when class-2 demand is low:

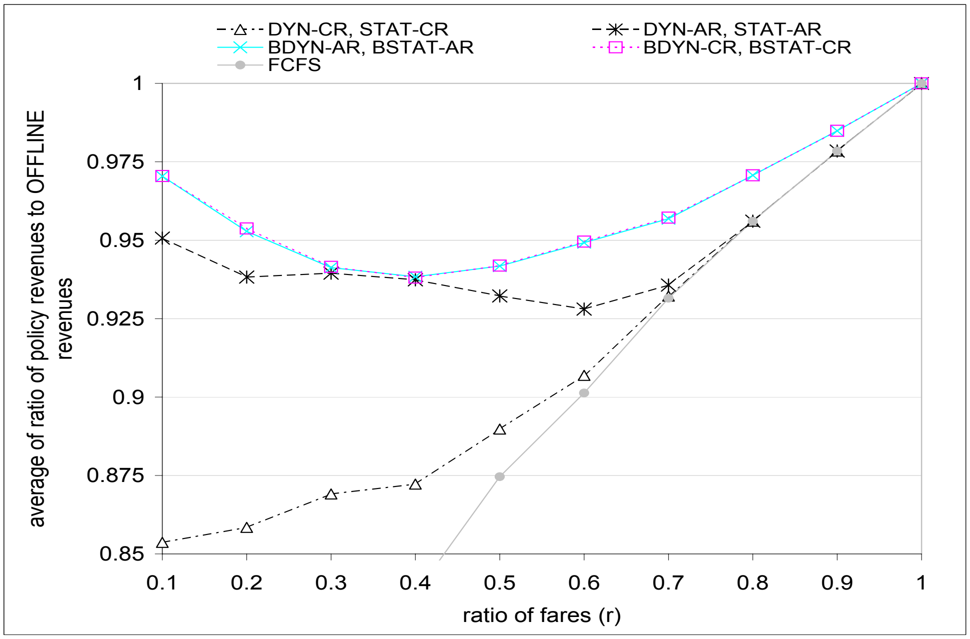

Figure 2: Performance of each policy relative to OFFLINE in Example-2a

These policies are more conservative and tend to protect less seats for class-1 compared to other policies. (2) MAR and dynamic policies are better in terms of guaranteed maximum revenue, 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 MAR policies when class-1 demand is low. (4) The use of correct demand information decreases the variance in revenues of MAR 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 MAR and CR policies is negligible when correct demand information is used and class-1 demand is low.

We also conducted experiments to see the effect of 'quality of' demand information on our policies, please see the doctoral dissertation of the second author (Gao, 2007) for details. Clearly, the performance of our policies can be very poor if the estimate of the range of demand distribution is very poor, 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.

7.2 Comparison to other well-known policies

In this section, we compare our policies to other well-known methods. We use (i) Littlewood's rule (we call it EMSR as in Belobaba, 1989) for $m = 2$, (ii) EMSR-b heuristic for $m > 2$ (see Talluri and van Ryzin, 2004a), (iii) the dynamic programming (DP) model of Brumelle and McGill (1993) when $m > 2$ and arrivals are LBH, and (iv) the Markov Decision Process (MDP) model of Lee and Hersh (1993) for Markovian arrivals. These policies are benchmarks for our methods and represent the *ideal* situations where one has complete information about arrival processes or probability distribution of demand. In this section, we only report the performance of the policies relative to OFFLINE. Each instance in each experiment involves 300 simulation runs. The capacity is $n = 100$ unless noted otherwise. The parameters of our policies are computed by setting the lower and upper bounds of demand to two standard deviations away from the (estimated) mean of each fare class. Our policies that use no demand information are excluded from the discussion in this section unless noted otherwise, because the focus is on the methods that use some form of demand information.

Example-3a. In this example, the fares are $(f_1, f_2) = (500, 100)$. The demand for each fare-class is independent and Poisson distributed. We denote the mean demand of class i as λ_i, $i = 1, 2$. The demand factor (ratio of total mean demand to capacity) is 1.2 while the mean demand of the fare classes $[\lambda_2, \lambda_1]$ range from [120,0] to [0,120]. When arrivals are time-homogeneous (see Figure 3), MDP is the optimal policy and it performs remarkably well compared to OFFLINE. Note that CR policies perform comparable to 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). There is less than 1% difference in the relative performance of any of our policies. We repeated this example with LBH arrivals. The arrival regime does not affect the relative performance of the policies, and their performance relative to EMSR does not change (EMSR is optimal when arrivals are LBH). The results with LBH arrivals and also the performance of policies that do not use demand information are presented in Example-3b in Appendix H of Lan and Gao (2007).

We utilized the same example keeping the mean demand at [60,60] but varying the fare ratio. The main observations remain the same: MDP (as the ideal solution) is close to OFFLINE when arrivals are time-homogeneous. Our policies are as good as EMSR at all fare ratios regardless of the arrival regime (except when class-2 demand is negligible). The relative performance of our policies is no worse than 95% in any of the experiments. In

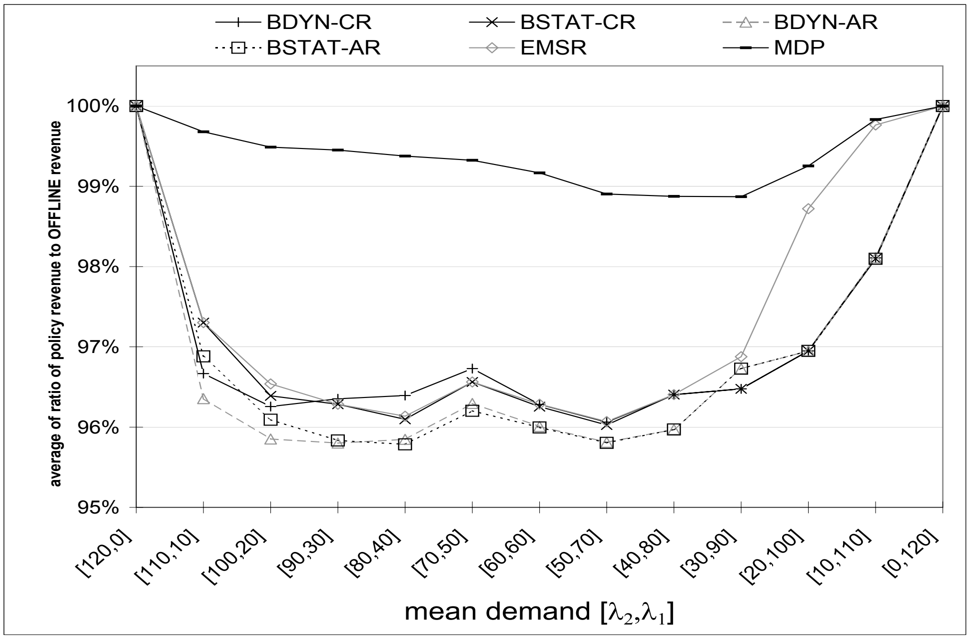

Figure 3: Average relative performance of policies in Example-3a

these experiments, the additional benefit of using the dynamic versions of our policies is not significant. This is because the demand is stationary and demand-mix is balanced. Next, we look at an example with more than two fare-classes.

Example-4a. This is adapted from the example used in Talluri and van Ryzin (2004a), Section 2.2.3.4. We have $m = 4$, fares are $f = (1050, 567, 527, 350)$, $n = 124$, the demand is Normally 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 1 through 4. All fare-classes have the same coefficient of variation (CoV) and we vary the CoV in this experiment. Note that static and dynamic policies are not equivalent in this example. The average relative performance is reported in Figure 4. 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 of our policies dominate EMSR-b, and the performance of EMSR-b degrades as the CoV increases. The performance of our policies with no demand information are given in Example-4b in Appendix H of Lan and Gao (2007).

The results of the last two experiments are very encouraging: Our methods with limited demand information are practically as good as EMSR and even better than EMSR-b which is commonly used in airline RM practice. But these examples assumed we have accurate

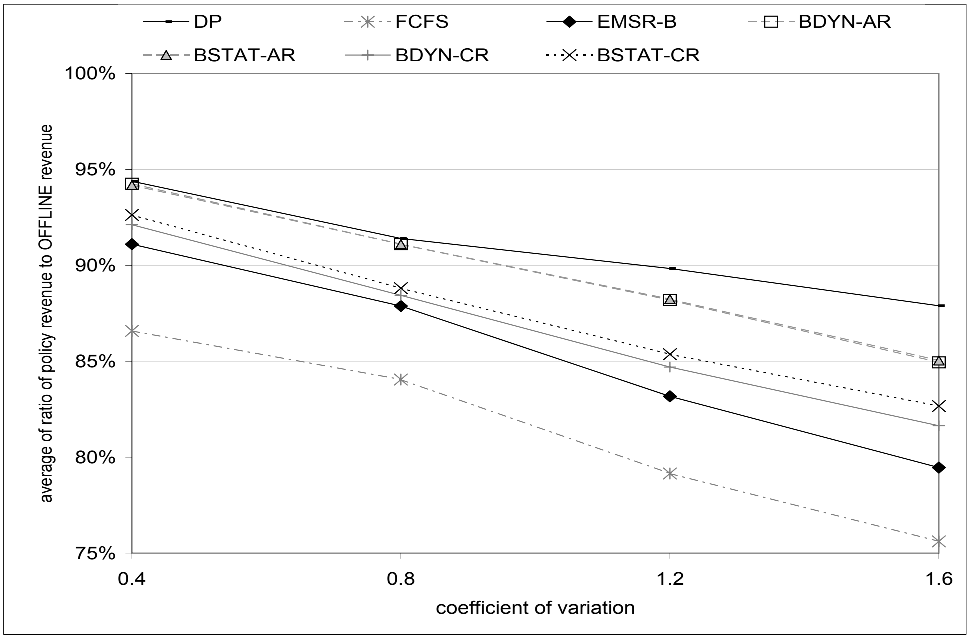

Figure 4: Average relative performance of policies in Example-4a

information about the demand. However, the main motivation for our work on robust methods is lack of data and accurate forecasts. In real life, not only will the demand forecasts be wrong, but also some of the assumptions in models such as EMSR, DP or MDP may fail to hold. We conducted further experiments to see the effect of inaccurate forecasts and/or demand information on these policies, as well as ours. See Examples-5a,5b,6a,6b,7 in Appendix H of Lan and Gao (2007). Those experiments reveal that our policies with limited demand information perform as good as EMSR under various scenarios. While the performance of MDP is close to OFFLINE in all the examples provided in this paper, its performance is highly sensitivity to the accuracy of rate of arrivals, especially when the demand rate is time-variant. In contrast, our policies are only sensitive to aggregate demand information, and do not require a complete characterization of the arrival process.

8 Conclusions

In this paper, we have analyzed the traditional single-leg RM problem from the perspective of competitive analysis of online algorithms. Our models make use of limited demand information and we derive static and dynamic policies for both competitive ratio and absolute regret criterion. Our optimal booking control policies have significant practical advantages:

the nesting property of booking limits is preserved while the need for information is reduced, and the optimal policy parameters are obtained in closed-form, hence the computational burden is minimal. Our computational experiments show that, when coupled with good estimates of the bounds on demand, our policies have average performance that is similar to EMSR and EMSR-b, which are two widely used methods. At the same time, our policies also perform well when subjected to demand fluctuation or errors.

From a research perspective, competitive analysis of online algorithms approach is very promising in RM. In this paper, we took a first step in using limited demand information to increase practical effectiveness of online algorithms in RM. Adding more demand information to the single-leg problem (e.g., use of time-varying bounds on demand in each fare class as opposed to static, aggregate bounds in our model) and application of the competitive analysis ideas to the network RM problem remain as challenging future research topics. In this paper, analysis was carried out assuming the demand information is given and static. However, it is worth investigating how demand bounds can be estimated, what types of estimation procedures and/or choice of demand bounds make our policies most effective, and how robust methods can be adapted to changes in demand information.

References

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

Ball, M. O. and M. Queyranne. 2006. Toward robust revenue management: competitive analysis of online booking. Working paper, available at SSRN: http://ssrn.com/abstract=896547.

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

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

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

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.

Bitran, G. and R. Caldentey. 2003. An Overview of Pricing Models for Revenue Management. *M&SOM,* 5, 203-229.

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

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

Curry, R.E. 1990. Optimal airline seat allocation with multiple nested fare classes. *Trans. Sci.* 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.

Feng, Y. and B. Xiao. 1999. Maximizing revenues of perishable assets with a risk factor. *Oper. Res.* 47, 337-341.

Gao, H. (2007) *Robust Revenue Management and Data-Driven Pricing.* On-going doctoral dissertation, Robert H. Smith School of Business, University of Maryland, College Park, MD.

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.

Kouvelis, P. and G. Yu. 1997. *Robust Discrete Optimization and Its Applications.* Kluwer Academic Publishers, Nonconvex Optimization and Its Applications, V.14, Boston.

Lan, Y. (2007). *Robust Revenue Management: Applications and Theory.* On-going doctoral dissertation, Robert H. Smith School of Business, University of Maryland, College Park, MD.

Lan, Y. and H. Gao (2007). *Technical Addendum to 'Seat Inventory Control with Limited Demand Information,'* Technical report, Robert H. Smith School of Business, University of Maryland, College Park, MD. URL: http://wam.umd.edu/~ylan/agifors-add.pdf.

Lautenbacher, C.J. and S. Stidham, Jr. 1999. The underlying Markov Decision Process in the single-leg airline yield management problem, *Trans. Sci.* 33, 136-146.

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.

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.

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

McGill, J. and G. van Ryzin. 1999. Revenue Management: Research Overview and Prospects. *Trans. Sci.* 33, 233-256.

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

Robinson,L. W. 1995. Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes. *Oper. Res.* 43, 252-263.

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

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

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

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.

Weatherford, L.R. and S. Pölt. 2002. Better unconstraining of airline demand data in revenue management systems for improved forecast accuracy and greater revenues. *J. Revenue and Pricing Manag.* 1, 234-254.

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