PHBS WORKING PAPER SERIES

Adaptive Robust Capacity Control with Adjustable Regret

Yingjie Lan
Peking University

January 2023

Working Paper 20230103

Abstract

A robust single-leg capacity control method based on adjustable regret is proposed to address the practical needs largely overlooked by previous studies for both smooth control of conservatism and adaptability to changing environments. Only the lower and upper bounds on demand are needed, and a performance guarantee is provided. Joint reduction finds one extreme scenario for each fare class, and a linear program of modest size is formulated, providing nested booking limit policies that are provably optimal among all online policies. Closed-form solutions are found for continuous problems, and their effectiveness is shown by simulation study.

Keywords: Revenue management, Adjustable regret, Robust optimization, Competitive analysis, Game theory, Competitive analysis
JEL Classification: E12, E44, G28, G32, G33

Peking University HSBC Business School
University Town, Nanshan District
Shenzhen 518055, China

Highlights

Adaptive Robust Capacity Control with Adjustable Regret

Yingjie Lan

- Fine-tuned policy that offers smooth adaptability and control of conservatism
- Extreme scenarios by joint reduction to formulate a linear programming model
- Closed-form robust revenue management policies with adjustable regret
- Extensive simulations to demonstrate effectiveness of new approach

Adaptive Robust Capacity Control with Adjustable Regret

Yingjie Lan[a]

[a]*Peking University, HSBC Business School, Shenzhen, 518005, Guangdong, China*

Abstract

A robust single-leg capacity control method based on adjustable regret is proposed to address the practical needs largely overlooked by previous studies for both smooth control of conservatism and adaptability to changing environments. Only the lower and upper bounds on demand are needed, and a performance guarantee is provided. Joint reduction finds one extreme scenario for each fare class, and a linear program of modest size is formulated, providing nested booking limit policies that are provably optimal among all online policies. Closed-form solutions are found for continuous problems, and their effectiveness is shown by simulation study.
Keywords: revenue management, adjustable regret, robust optimization, game theory, competitive analysis

1. Introduction

Revenue management (RM) is a highly successful practice of management science, with wide applications in industries such as air travel, hospitality, and car rentals. To maximize revenues from heterogeneous customers arriving stochastically over a time period, sales of products provided from unreplenishable resources must be strategically controlled. Uncertainty in revenues due to demand fluctuations require certain criteria and assumptions for revenue comparison and optimization. Classical RM models assume risk neutrality and exact demand distributions, and maximize expected revenues, on which Talluri et al. (2004) provides a comprehensive overview.

As application of RM deepens, it is gradually realized that the assumptions in classical models can be unrealistic in practice. The risk neutrality assumption is questioned by Feng and Xiao (1999) and Lancaster (2003), and Barz (2007) observes risk reduction efforts from managers for infrequent decisions. Meanwhile, accurate demand distributions may be hard to obtain. Vinod (2021b) points out that even in a stable business environment with plenty of historical data, demand forecasts can still be quite inaccurate at the itinerary level most crucial for airline RM. Things get worse in a new or unstable business environment (Lan et al. 2011) when historic data are scarce or do not have much relevance or predictive power.

To apply RM without such assumptions, a rich set of *robust RM models* have been developed recently under alternative specifications of demand uncertainty. Distributional vagueness within a continuum is employed in Birbil et al. (2009), Rusmevichientong and Topaloglu (2012), and Sierag and van der Mei (2016), while Lan et al. (2008), Ball and Queyranne (2009), Perakis and Roels (2010), Lan et al. (2011), and Ma et al. (2021) are *distribution-free* in that no distributions or probabilities are required at all, instead, an *uncertainty set* is used to specify the possible realizations of demand, e.g., within some lower and upper bounds. Such uncertainty sets can be estimated by expert judgments (Perakis and Roels 2010), or by numerical methods (Sierag and van der Mei 2016) from historical data.

These robust models cover a rich set of RM tasks under various settings. Some study single-leg capacity control (e.g. Lan et al. 2008 and Ball and Queyranne 2009), some study joint decisions of capacity control and overbooking (Lan et al. 2011), and others study network RM problems (Perakis and Roels 2010). They require less input information than their classical counterpart, and provide conservative solutions that sacrifice some revenues to guard against downside risks with a performance guarantee.

The reader is referred to Gönsch (2017) for a systemic survey on robust RM models, and to Perakis and Roels (2008) and Kouvelis and Yu (2013) for general literature on robust decision-making and decision criteria.

However, as a seasoned practitioner observes in Vinod (2021b), conservatism and inflexibility are the two main reasons preventing robust RM models from being widely adopted by airlines. The conservatism raises concern about missed revenue opportunities, resulting in significant revenue gaps between robust and classical methods. In a tumultuous time of the global COVID-19 pandemic, when classical methods do not work due to the volatile nature and unpredictability of future outbreaks, it calls for an adaptive robust RM approach that takes corrective actions on key performance indicators (KPIs) monitored in real-time. This adaptive approach to master uncertainty has similarities to the sales and operations workflow (Palmatier and Crum 2003) with continuous demand management for peak performance.

Yet the robust RM models in the literature only recommend stable controls that can not readily provide adaptability. Admittedly, there are some interesting results on dynamical recalibration of the policy on the fly in Lan et al. (2008), but it does not allow for controlled response to real time KPIs or to a more optimistic outlook from textual data analysis by artificial intelligence. Therefore, those robust RM methods have never been a real priority for airlines and RM software providers (Vinod 2021b).

This paper sets out to develop an adaptive robust RM method by working with a new decision criterion. Three robust decision criteria are commonly seen in the RM literature so for: the maximin, the absolute regret (or Savage), and the relative regret (or competitive ratio), which do not have any adaptivity built in them and their level of conservatism can not be fine-tuned. The maximin criteria are considered the most conservative for revenue maximization, and Perakis and Roels (2010) further observes that the absolute regret criteria are less conservative than the relative regret criteria.

A crude way to be adaptive is to switch between these criteria, which is very coarse-grained, providing only three distinct levels of conservatism and revenue, and quite limited in its adaptive range. Fortunately, the adjustable regret minimization (ARM) criterion proposed in Lan (2021) presents a viable alternative for the purpose at hand, which has a control parameter to interpolate between these criteria for fine-tuned control, and extrapolate beyond them to provide a much wider adaptive range.

An adaptive robust RM method for the classical single-leg capacity control problem can address these practical issues: (i) accurate data may not be available for new or unstable business environments; (ii) decision-makers are not necessarily risk-neutral, e.g., they may be concerned about downside risks or worst-case results in the short run; (iii) the ability to fine-tune the level of conservatism can help reduce the revenue gap between robust and classical methods; (vi) adaptive robust RM methods are needed to respond in real-time to an ever-changing environment. The new method is analytically appealing in many ways. The extreme scenarios are identified by a new technique of *joint reduction* to greatly simplify the problem. Then a simple linear program is formulated, with the optimal solution in closed-form, providing a nested booking limit policy that is provably optimal among all online policies. The policy gets more aggressive as the ARM control parameter increases, giving a clear direction of change. The more general analysis here provides alternative proofs for some of the results in Lan et al. (2008).

The rest of the paper is organized as follows. Section 2 introduces the problem and outlines the new adaptive approach. Reduction of the scenarios under nested booking policies are included in section 3. Derivation of the optimal solution are presented in 4. Section 5 provides computational results. Conclusions with future research suggestions are given in section 6.

2. Problem Specification

As a cornerstone problem in RM, single-leg capacity control has its origin in airlines but quickly spreads to other industries. It involves selling products that require a unit of a single resource, such as the limited capacity in the main cabin of a flight leg. The products or fares are offered at different prices with different features or restrictions, such as cancellation options and luggage limits, to cater to different customer segments. Ideally, price-insensitive customers would be dissuaded to buy lower fares by the restrictions, so that customers are well segmented by the design. This study limits itself to independent demand, with each product corresponding to exactly one disjoint customer segment. The capacity, products, and prices are determined long in advance, from factors such as business strategy, market segments, and competitor price points.

Capacity control in airlines takes place as fare requests arrive one by one at a booking process, which opens a certain length of time before and closes at the flight's departure, when an empty seat on the flight has no more value. The challenge to maximize revenue lies in making an irrevocable decision upon the arrival of a request whether to satisfy it with the capacity on hand or to reject it and preserve that capacity for uncertain future requests. Unlike classical models, there is no distributional information on demand, instead, it is only assumed to know the lower and upper bounds on the number of units demanded in each fare, which defines an *uncertainty set* Ω with all possible request sequences satisfying these bounds.

To develop an adaptive robust RM method that works with such uncertainty sets, the ARM criterion is adopted along the competitive analysis of online algorithms (see Albers 2003 for a survey). An online algorithm in this context implements a booking policy, which tells what to do under all possible situations in a booking process. In competitive analysis, an online algorithm π is evaluated relative to an offline optimal algorithm Υ that knows in advance (with hindsight) the entire sequence of requests $I \in \Omega$ to its advantage. Let $R(I; \dot{})$ denote the objective function value of an algorithm on input I, and let $R^*(I) = R(I; \Upsilon)$ for the offline optimal revenue from I. There are two common metrics of relative performance: absolute regret, which is the difference $R^*(I) - R(I; \pi)$, and relative regret, which is $1 - R(I; \pi)/R^*(I)$. As these metrics do not allow adaptability, the ARM criterion introduces an *adjustable regret metric* $\mathcal{M}(I; \pi) = \beta R^*(I) - R(I; \pi)$ that is parameterized by $\beta \geq 0$. The *guarantee of performance* by an algorithm π is the worst-case metric $\max_{I \in \Omega} \mathcal{M}(I; \pi)$. Like the other relative criteria, it chooses a π from a feasible set Π to minimize the regret guarantee:

$$\min_{\pi \in \Pi} \max_{I \in \Omega} \mathcal{M}(I; \pi), \tag{1}$$

The β parameter in (1) plays a crucial role in moderating aggressiveness and making it adaptive. As β increases, the benchmark becomes more demanding in a more favorable scenario I indicated by a bigger $R^*(I)$, and an algorithm that best conforms to such requirement will be recommended. This intuition is reflected in Lan (2021) with an observation: the ARM criterion degenerates into the most conservative maximin criterion with $\beta = 0$; as *beta* increases to the competitive ratio (between 0 and 1), it is equivalent to the more aggressive relative regret criterion; and when β reaches 1, the even more aggressive absolute regret criterion is recovered; as β approaches ∞, it becomes the most aggressive maximax criterion. That observation also shows that β may be seen as an interpolation parameter, with a bigger β for a more aggressive criterion. The interplay between β, expected revenue, and overall risk has been studied for one-way trading in Lan (2021) by simulation, and section 5 provides some simulation results in RM.

This study focuses on determining optimal algorithms for both continuous and discrete problems. In continuous (discrete) problems, the capacity, the demands, and the accepted requests can be any positive real (integral) numbers, while a request in a fare class may be split and partially accepted as necessary. Although the continuous case is less realistic, it admits closed-form solutions, while the discrete case requires solving a mixed integer linear program.

Allowing request splitting simplifies the analysis with multiunit requests (as in batch arrivals), but group bookings cannot be guaranteed by splitting. The analysis of the more intuitive discrete problems is fully presented first, while the results carry over to continuous problems via a limit process. Splitting of requests in discrete problems breaks up a multiunit request into multiple requests that demand one unit of capacity in a fare class. The analysis with unit requests carries through as long as each multiunit request demands a nonnegative and finite amount. In the rest of the paper, it is assumed, without loss of generality, that input sequences consist of unit requests.

3. Scenario Reduction

The scenario reduction is inspired by the idea of elimination of dominated strategies in game theory, but it is based on a new concept of *joint dominance* between sets of strategies. The problem in (1) may be viewed as a sequential zero-sum game between a manager and an imaginary adversary who chooses a scenario I from Ω after observing the algorithm chosen by the manager. In such a game, a scenario I dominates another I' if there is $\mathcal{M}(I;\pi) \geq \mathcal{M}(I';\pi)$ for all $\pi \in \Pi$. It can be generalized to two sets of strategies:

Definition 1. *(Joint Dominance and Reduction) A set* Ω_2 *jointly dominates* Ω_1

if for any $\pi \in \Pi$, a scenario $I_2 \in \Omega_2$ can be found for any $I_1 \in \Omega_1$, such that $\mathcal{M}(I_1; \pi) \leq \mathcal{M}(I_2; \pi)$ (or I_2 dominates I_1 under π). If additionally there is $\Omega_2 \subset \Omega_1$, then Ω_2 is a joint reduction of Ω_1.

Some comments may be helpful. At the core of joint dominance is a mapping $F_\pi : \Omega_1 \to \Omega_2$ for each $\pi \in \Pi$ such that $\forall I \in \Omega_1 : \mathcal{M}(I; \pi) \leq \mathcal{M}(F_\pi(I); \pi)$. The dominance of one strategy over another is recovered with singleton sets. Joint reduction ensures that

$$\forall \pi \in \Pi : \max_{I \in \Omega_1} \mathcal{M}(I; \pi) = \max_{I \in \Omega_2} \mathcal{M}(I; \pi), \tag{2}$$

while for joint dominance there is $\forall \pi \in \Pi : \max_{I \in \Omega_1} \mathcal{M}(I; \pi) \leq \max_{I \in \Omega_2} \mathcal{M}(I; \pi)$. Under joint dominance, all strategies in $\Omega_1 \setminus \Omega_2$ can be eliminated from Ω, without changing the game value in (1), since $\Omega \setminus (\Omega_1 \setminus \Omega_2)$ is a joint reduction of Ω. Both joint dominance and reduction are transitive relationships. One can start with Ω and use joint reduction recursively to simplify the problem.

A demand scenario is a sequence of booking requests for a unit of a product. The uncertainty set Ω contains a combinatorial number of all the sequences that observe the demand bounds. Let $L = (L_1, \cdots, L_m) \geq 0$ and $U = (U_1, \cdots, U_m) \geq L$ for the lower and upper bounds on requests for the fare products. The *profile* of a sequence I is the m-dimensional vector $(I[j], j = 1, \cdots, m)$, where $I[j]$ is the total number of class j requests in sequence I. The notation $I(t)$ represents the class of t^{th} request in sequence I, for $t = 1, \cdots, |I|$, with $I(1 : t)$ for the subsequence of the first t requests. The uncertainty set is explicitly defined as $\Omega = \{I : L_i \leq I[i] \leq U_i, i = 1, \cdots, m\}$.

The size of Ω can be reduced under *standard nested booking limit* (SNBL) policies, which are proved later to be optimal among all online policies. Let n denote the total capacity of the resource (seats, rooms, etc.) and let m denote the number of fare classes (products) with fare prices $f_1 > f_2 > \cdots > f_m > 0$. A record $a(t) = (a_1(t), \cdots, a_m(t))$ is maintained to track the number of accepted requests for each fare class *after* the first t requests are processed. A SNBL policy $b = (b_1, \cdots, b_m)$ has m nested booking limits $b_1 \geq b_2 \geq \cdots \geq b_m \geq 0$ such that b_i limits the total number of requests accepted for classes i to m: $\forall \tau : b_i \geq \sum_{j=i}^{m} a_j(\tau), i = 1, \cdots, m$. The policy is open to fare class k at $(\tau+1)^{th}$ request if the affected limits have room:

$$\sum_{j=i}^{m} a_j(\tau) < b_i, i = 1, \cdots, k, \tag{3}$$

and closed to it otherwise. Due to the nested nature in (3), if a SNBL policy is closed to a fare class, then it is closed to all lower fare classes. The critical class $C(I; b)$ is the highest fare class that is closed after processing all requests in I by b:

$$C(I; b) = \max\{c \leq m + 1 : \sum_{i=j}^{m} a_i(|I|) < b_j, j = 1, \cdots, c - 1\}, \tag{4}$$

where $|I|$ is the total number of requests in I. If the policy is open to all fare classes, then $C(I; b) = m + 1$ (the virtual class with $b_{m+1} = 0, f_{m+1} = 0$). The decision for $I(t + 1)$ is to accept if $I(t + 1) > C(I(1 : t), b)$ and reject otherwise, with the record update given by $a_i(t+1) = a_i(t) + \mathbf{1}\{i = I(t+1) > C(I(1 : t), b)\}$, where $\mathbf{1}\{\cdot\}$ is the indicator function.

An alternative way to describe a SNBL policy is to use *bucket sizes*, denoted by a vector x defined by $x_i = b_i - b_{i+1}$ for $i = 1..m$, where a virtual fare class $m + 1$ is introduced with $b_{m+1} = 0, f_{m+1} = 0$ for convenience. Note that either of these vectors is sufficient to characterize a SNBL policy. The protection levels commonly used in single-leg RM are easily derived from the booking limits (see Talluri et al. 2004). The number of seats protected for classes 1 to k $(k = 1, \cdots, m - 1)$ is equal to $n - b_{k+1}$, which must never be sold to the lower fares classes $k+1$ to m. Let $a(\cdot)$ be the booking record for policy b on input I, then there is $R(I;b) = \sum_{i=1}^{m} f_i a_i(|I|)$. A knapsack problem finds $R^*(I) = \max\{\sum_{i=1}^{m} f_i y_i : \sum_{i=1}^{m} y_i \leq n, I[i] \geq y_i \geq 0, i = 1, \cdots, m\}$.

The set of SNBL policies can be slightly reduced without any side effects. Note if $x_i > U_i$ (or $b_i > b_{i+1} + U_i$), then b_i can never be reached in (3) since $a_i(\tau) \leq U_i$ as long as $I \in \Omega$. A new policy b' can be defined by $b'_i = b_{i+1} + U_i$ and $b'_j = b_j$ for $j \neq i$ so that $x'_i = U_i$, and (3) makes it clear that b' always make the same booking decisions as b. Therefore, the restriction $x \leq U$ is enforced on policies for consideration:

$$B(n) = \left\{ x : \sum_{i=1}^{m} x_i \leq n, 0 \leq x \leq U \right\}. \tag{5}$$

The first step of reduction is to consider the low-before-high (LBH) sequences, where a lower fare request always arrives before a higher one. For a sequence I, let LBH(I) denote the rearrangement of all requests in I into LBH order. There is $R^*(I) = R^*(\text{LBH}(I))$ as the profiles of I and LBH(I) are the same (to define the same knapsack problem), and Lan et al. (2008) prove for any SNBL policy b there is always $R(I;b) \geq R(\text{LBH}(I);b)$. Therefore, there is $\mathcal{M}(\text{LBH}(I);b) \geq \mathcal{M}(I;b)$ for any $b \in B(n)$, which leads to

Proposition 1 (LBH Reduction)**.** *The LBH subset defined by LBH(Ω) = \{LBH(I) : $I \in \Omega$\} is a joint reduction of Ω with SNBL policies.*

From the standpoint of ARM, all non-LBH sequences can be ignored due to Proposition 1. This yields a substantial scenario reduction, yet the total number of sequences can still be prohibitive, which can be further reduced by considering critical classes.

Proposition 2 (Permutation Invariance). *The critical class is permutation invariant, so that $C(I;b) = C(LBH(I);b)$ for any input sequence I and SNBL policy b.*

Proof. Let I' only differ from I by a swap of two requests in I at position s and $s+1$: $I'(s) = I(s+1)$ and $I'(s+1) = I(s)$. Assume $I(s) < I(s+1)$, otherwise, simply switch I' and I to satisfy it. It suffices to show $C(I;b) = C(I';b)$, as any permutation can be generated by many such swaps at different positions.

Let $a_j(t)$ ($a'_j(t)$) be the booking records for processing sequence I (I'). Note that as I and I' are the same before request s, there is $a_j(s-1) = a'_j(s-1), j = 1, \cdots, m$. It is clear from (3) that it is impossible to reject $I(s)$ and then accept $I(s+1)$ as $I(s) < I(s+1)$, so there are only three possible cases of outcomes to consider.

Case 1: both $I(s)$ and $I(s+1)$ are accepted. The record update is

$$a_j(s+1) = a_j(s-1) + \mathbf{1}\{j \in \{I(s), I(s+1)\}\}, j = 1, \cdots, m,$$

with $a(s+1)$ satisfying (3). Accepting $I'(s)$ and $I'(s+1)$ will result in $a'(s) \leq a'(s+1) = a(s+1)$. Therefore both $a'(s)$ and $a'(s+1)$ satisfy (3), which means both $I'(s)$ and $I'(s+1)$ will be accepted. As $a'(s+1) = a(s+1)$ and $I(t) = I'(t), t > s+1$, the same decisions will be made for both sequences afterward.

Case 2: $I(s)$ is accepted but $I(s+1)$ is rejected. The record updated is

$$a_j(s+1) = a_j(s) = a_j(s-1) + \mathbf{1}\{j = I(s)\}, j = 1, \cdots, m.$$

Consider two subcases. **Subcase 2.1**: $I'(s)$ is rejected. Then $a'_j(s) = a'_j(s-1) = a_j(s-1), j = 1, \cdots, m$, and it is clear that $I'(s+1)$ will be accepted just like $I(s)$, since (3) is identical for both. The new records will be $a'(s+1) = a(s+1)$, and the same decisions will be made for both sequences afterward. **Subcase 2.2**: $I'(s)$

is accepted. Note that Case 1 does not require the assumption of $I(s) < I(s+1)$, thus the reasoning applies backward, which means the acceptance of both $I'(s)$ and $I'(s+1)$ implies the acceptance of both $I(s)$ and $I(s+1)$ as well. Therefore, $I'(s+1)$ must be rejected in this subcase, and the record update is

$$a'_j(s+1) = a'_j(s) = a'_j(s-1) + \mathbf{1}\{j = I'(s)\}, j = 1, \cdots, m.$$

Note the differences between $a'(s+1)$ and $a(s+1)$ are

$$a'_j(s+1) = \begin{cases} a_j(s+1) + 1 & \text{if } j = I(s+1), \\ a_j(s+1) - 1 & \text{if } j = I(s), \\ a_j(s+1) & \text{otherwise,} \end{cases}$$

which immediately implies

$$\sum_{j=i}^{m} a'_j(s+1) = \sum_{j=i}^{m} a_j(s+1), i = 1, \cdots, I(s).$$

The rejection of $I'(s+1)$ by (3) requires that there exists $i \in \{1, \cdots, I'(s+1) = I(s)\}$ such that

$$b_i = \sum_{j=i}^{m} a'_j(s) = \sum_{j=i}^{m} a'_j(s+1) = \sum_{j=i}^{m} a_j(s+1).$$

Therefore, all future requests with fare lower than $I(s)$ are rejected, and the acceptance condition (3) for requests with higher fares in the future will always be the same for both $a'(t)$ and $a(t)$. In other words, the future decisions must be identical, ending up with $C(I; b) = C(I'; b) \geq I(S)$.

Case 3: both $I(s)$ and $I(s+1)$ are rejected. The updated records are $a_j(s+1) =$

$a_j(s-1), j = 1, \cdots, m$, and it is clear that both $I'(s)$ and $I'(s+1)$ must also be rejected. Since $a'(s+1) = a(s+1)$ and $I(t) = I'(t), t > s+1$, the same decisions will be made for both sequences afterward.

In all cases, there is $C(I; b) = C(I; b')$. $\square$

The step sets are defined for further reduction on the input sequences. Let LBH$[P]$ denote the LBH sequence with profile P.

Definition 2 (Step Set)**.** *The step set $D^j, j = 1, \cdots, m+1$ contains a series of LBH sequences $D^j_d = LBH[P^j_d], d = 0, \cdots, \sum_{k=1}^{j-1}(U[k] - L[k])$ whose profiles P^j_d are stepwise incremental: $P^j_0 = (L[1], \cdots, L[j-1], U[j], \cdots, U[m])$, and $P^j_{d+1}[i] = P^j_d[i] + \mathbf{1}\{i = k(d+1)\}, i = 1, \cdots, m$, where the increment class function*

$$k(d) = \min\{k : \sum_{i=1}^{k}(U[i] - L[i]) \geq d\}. \tag{6}$$

Intuitively, the step set starts from D^j_0, then repeatedly add a request of the highest fare class allowed by the upper bounds to have the next sequence in the series, until reaching LBH$[U]$. The set size $|D^j| = 1 + \sum_{i=1}^{j-1}(U[i] - L[i])$ is linear in the demand spread $U[i] - L[i]$. In the special case of $j = 1$, there is $|D^1| = 1$ as $\sum_{k=1}^{0}(U[k] - L[k])$ is an empty sum that evaluates to 0, with $D^1_0 =$ LBH$[U]$. It turns out that any sequence is dominated by a step sequence under a given policy.

Proposition 3 (Step Map)**.** *Any sequence $I \in \Omega$ is dominated by a step sequence $D^{C(I;b)}_{d(I;b)}$ under a policy $b \in B(n)$, with both sequences having the same critical class under b, where $d(I; b) = \mathbf{1}\{\beta > 1\} \sum_{i=1}^{C(I;b)-1}(I[i] - L[i])$.*

Proof. The proof constructs a finite series of sequences sharing the same critical class and increasing in metric $\mathcal{M}(\cdot; b)$, with the series starting from I and ending with $D^{C(I;b)}_{d(I;b)}$. Let $c = C(I; b)$ for short.

If $I[k] < U[k]$ for some $k \geq c$, let I' be obtained by appending a request of class k to the end of I. By definition of $C(I;b)$, the appended request will be rejected by b, resulting in: 1. online revenue will not change, 2. offline revenue may increase, and 3. the critical class will not change. Therefore, the metric increases and I' dominates I under policy b. Append enough fare class k requests until there are $U[k]$ requests for all $k = c, \cdots, m$, then rearrange it in LBH order without affecting the critical class (Proposition 2). Let $\hat{I}$ denote the final sequence rearranged in LBH order.

Note that for any sequence J with $C(J;b) = c$, the requests of classes 1 to $c-1$ in J must all be taken by the optimal knapsack problem for $R^*(J)$, for the following reasons. As no requests of classes 1 to $c-1$ are rejected, the booking record satisfies $a_i(|J|) = J[i]$ for all $i < c$. Meanwhile, the definition of critical class requires by (3) that $\sum_{j=i}^{m} a_j(|J|) < b_i$ for all $i < c$, which implies $\sum_{j=1}^{c-1} J[j] < n$, therefore all requests in $J[i], i < c$ must be accepted to have the maximal $R^*(J)$. With this, consider two cases to further work on $\hat{I}$.

Case 1: $\beta \in [0,1]$. If $\hat{I} \neq D_0^c$, then there is a $k < c$ such that $\hat{I}[k] > L[k]$. Take one class k request away from $\hat{I}$ to get $\hat{I}'$, which maintains $C(\hat{I}';b) = c$. There is $R(\hat{I}';b) = R(\hat{I};b) - f_k$ as the removed request was originally accepted. And $R^*(\hat{I}') = R^*(\hat{I}) - f_k$ must hold as $C(\hat{I}';b) = C(\hat{I};b)$. Therefore, the metrics satisfy $\mathcal{M}(\hat{I}';b) \geq \mathcal{M}(\hat{I};b) + f_k(1-\beta)$, which means $\hat{I}'$ dominates $\hat{I}$ under policy b. Keep doing this until D_0^c is obtained. Therefore, I is dominated by D_0^c under policy b, while maintaining $C(D_0^c;b) = c$.

Case 2: $\beta > 1$. If $\hat{I} \notin D^c$, then there exist classes i and k that satisfy $i < k < c$, $\hat{I}[k] > L[k]$, and $\hat{I}[i] < U[k]$. Take one class k request away and add one class i request to $\hat{I}$ while maintaining the LBH order to get a new $\hat{I}'$. There is $R(\hat{I}';b) = R(\hat{I};b) + f_i - f_k$ as all requests of any class $j < c$ are accepted. And $R^*(\hat{I}') = R^*(\hat{I}) + f_i - f_k$ must also hold, since there is $C(\hat{I}';b) = C(\hat{I};b)$ with $i < k$.

Therefore, the metrics satisfy $\mathcal{M}(\hat{I}'; b) = \mathcal{M}(\hat{I}; b) + (\beta - 1)(f_i - f_k)$, which means $\hat{I}'$ dominates $\hat{I}$ under policy b. Keep doing this until having $\hat{I}' \in D^c$, and since $\sum_{i=1}^{c-1} \hat{I}'[i] = d(I; b)$ is maintained, the final sequence must be $D^c_{d(I;b)}$. □

Note that when $\beta \leq 1$, Proposition 3 implies that the set $D_0^* = \{D_0^i : i = 1, \cdots, m+1$ is a joint reduction of Ω. For $\beta > 1$, however, it can only ensure that the set $D = \bigcup_{i=1}^{m+1} D^i$ is a joint reduction of Ω, for which a final reduction is possible. As $C(D^c_d; b) \geq c$ for any policy $b \in B(n)$ and $C(D^c_{d+1}; b) \leq C(D^c_d; b)$, let $\check{D}^c(b) = \{I \in D^c : C(I; b) = c\}$ and $\check{d}^c(b) = |\check{D}^c(b)|$. Clearly there is $\check{D}^c(b) = \{D^c_d : 0 \leq d < \check{d}^c(b)\}$.

Proposition 4. *For any policy $b \in B(n)$, the critical class $C(D^j_d; b) = 1$ for $d \geq \check{d}^j(b)$ and the step change of online revenues $\Delta R(D^j_d; b) = R(D^j_d; b) - R(D^j_{d-1}; b)$ satisfies*

$$\Delta R(D^j_d; b) = \mathbf{1}\{d \leq \check{d}^j(b)\} f_{k(d)}. \tag{7}$$

Proof. For $d \leq \check{d}^j(b)$, as $k(d) < j = C(D^c_{d-1}; b)$, the added request of class $k(d)$ is accepted by b, giving $\Delta R(D^c_d; b) = f_{k(d)}$. Note that at $d = \check{d}^j(b)$, there is $C(D^c_d; b) < j$ as a result of accepting the added request of class $k(d)$, which means $k(d) \geq C(D^c_d; b)$ from (4). But since $D^d_j[k] = U_k \geq x_k = b_k - b_{k+1}, k = 1..k(d) - 1$, all booking limits from $k(d)$ up to 1 must be reached, which gives $C(D^c_d; b) = 1$ at $d = \check{d}^j(b)$ and beyond. Thus all added requests for $d > \check{d}^c(b)$ are rejected, resulting in $\Delta R(D^c_d; b) = 0$. □

Let $\bar{G}^j_d \equiv \beta R^*(D^j_d) - R^+_j(D^j_d)$ be a surrogate independent of b for $\mathcal{M}(D^j_d, b)$, where $R^+_j(D^j_d) = \sum_{i=1}^{j-1} f_i D^j_d[i]$. A joint reduction by critical classes to $m+1$ scenarios is always possible.

Theorem 1 (Reduction by Critical Classes)**.** *Given a policy $b \in B(n)$, any $I \in \Omega$ can be mapped to a step sequence $D^c_{d(c,\beta)}$ that dominates I under b, where $c = C(I; b)$*

is the critical class and

$$d(c,\beta) \equiv \min \operatorname*{arg\,max}_{d\in 0..|D_j|-1} \bar{G}_d^c$$
$$= \max\{d : \bar{G}_d^c > \bar{G}_{d-1}^c\}, \tag{8}$$

is non-decreasing in β *and does not depend on* b*. The mapping establishes set* $D(\beta) = \{D^c_{d(c,\beta)} : c = 1, \cdots, m+1\}$ *as a joint reduction of* Ω*.*

Proof. It can be easily verified for $\beta \in [0,1]$ that $d(c,\beta) = 0$ and $\mathcal{D}(\beta) = D_0^*$, so the proof only considers $\beta > 1$. The case of $|D^c| = 1$ is also simple, so suppose $|D^c| > 1$, implying $c > 1$ as $|D^1| = 1$.

Let $\Delta\bar{G}_d^c \equiv \bar{G}_d^c - \bar{G}_{d-1}^c$, and $\Delta R^*(D_d^c) \equiv R^*(D_d^c) - R^*(D_{d-1}^c)$. Then with (6) there is $\Delta R_j^+(D_d^c) \equiv R_j^+(D_d^c) - R_c^+(D_{d-1}^c) = f_{k(d)}$, and

$$\Delta\bar{G}_d^c = \beta\Delta R^*(D_d^c) - f_{k(d)}.$$

Let $D_c^d(-t)$ denote the t^{th} last request in D_c^d, then there is

$$R^*(D_d^c) = \sum_{t=1}^{n} f_{D_d^c(-t)},$$

where $D_d^c(-t) = m+1$ is assumed for convenience if $t > |D_d^c|$. If $k(d) \geq D_c^d(-n)$, the added request of class $k(d)$ will not increase revenue, resulting in $\Delta R^*(D_c^d) = 0$ and $\Delta\bar{G}_d^c = -f_{k(d+1)} < 0$ from then on, which means $k(d) < D_c^d(-n)$ at $d = d(c,\beta)$. Therefore, $d(c,\beta)$ satisfies $k(d) < D_d^j(-n)$, which means the added request of class $k(d)$ is accepted by the offline optimal policy and the request $D_d^j(-n)$ is "crammed out", giving $\Delta R^*(D_d^j) = f_{k(d)} - f_{D_d^j(-n)} > 0$. Clearly there is

$$\Delta \bar{G}_d^j = (\beta - 1) f_{k(d)} - \beta f_{D_d^j(-n)},$$

which decreases in d, as $k(d)$ increases in d, while $D_d^j(-n)$ decreases in d. Therefore, the solution in (8) is correct. As $\Delta \bar{G}_d^j$ increases in β, it is clear that $d(c, \beta)$ increases in β.

Let $\Delta \mathcal{M}(D_d^c; b) \equiv \mathcal{M}(D_d^c; b) - \mathcal{M}(D_{d-1}^c; b) = \beta \Delta R^*(D_d^c) - \Delta R(D_d^c; b) = \Delta \bar{G}_d^c + \Delta R_c^+(D_d^c) - \Delta R(D_c^d; b)$, Since $\Delta R_c^+(D_d^c) = f_{k(d)}$, Proposition 4 provides

$$\Delta \mathcal{M}(D_d^c; b) = \Delta \bar{G}_d^c + \mathbf{1}\{d > \check{d}^c(b)\} f_{k(d)}$$

which means that $\mathcal{M}(D_d^c; b) - \bar{G}_d^c$ is a constant if $d \leq \check{d}^c(b)$.

As $D_{d(I;b)}^c$ dominates I under b with $C(D_{d(I;b)}^c; b) = C(I; b) = c$ by Proposition 3, there is $d(I; b) \leq \check{d}^c(b)$. Consider two cases. **Case 1**: $d(c, \beta) \leq \check{d}^c(b)$. As $\mathcal{M}(D_d^c; b) - \bar{G}_d^j$ is constant for $d = 0, \cdots, \check{d}^c(b)$, there is $d(c, \beta) = \min \arg\max_{d \in 0..\check{d}^c(b)} \mathcal{M}(D_d^c; b)$, implying that $D_{d(c,\beta)}^c$ dominates $D_{d(I;b)}^c$ under b. **Case 2**: $d(c, \beta) > \check{d}^c(b)$. There is $\Delta \mathcal{M}(D_d^c; b) \geq \Delta \bar{G}_d^j > 0$ for $d \leq d(c, \beta)$, and again $D_{d(c,\beta)}^c$ dominates $D_{d(I;b)}^c$ under b. Therefore, in both cases $D_{d(c,\beta)}^c$ dominates I under b. $\square$

All the results on scenario reduction so far are for discrete problems that permit multiunit requests with splitting, which can be readily extended to continuous problems. The key is to approximate a continuous problem with a discrete problem with multiunit requests, where a discrete unit is redefined to be a fractional amount of $1/N$ (N is a big integer). A continuous amount a in the continuous problem is converted to $\lceil aN \rceil$ units in the discrete problem. Then a limiting process with $N \to \infty$ will bring the approximation to perfection, with all results carried over to continuous problems that allow request splitting.

4. Optimal Policy

The problem (1) is first considered with $\Pi = B(n)$ to formulate a linear program that enjoys closed-form solutions. Then the SNBL policies in $B(n)$ are proved optimal among online algorithms. The effect of β on the solution is also studied.

According to (2), Theorem 1 greatly reduces the complexity of regret guarantee

$$\forall b \in B(n) : Z(b) \equiv \max_{I \in \Omega} \mathcal{M}(I; b) = \max_{j \in 1..m+1} \mathcal{M}(D^j_{d(j,\beta)}; b). \tag{9}$$

A major difficulty to formulate a linear program out of (9) lies in the online revenue $R(D^j_{d(j,\beta)}, b)$ being a nonlinear function of b or x, which is involved in $\mathcal{M}(D^j_{d(j,\beta)}; b)$. To circumvent, a linear approximation of the online revenue is defined as

$$\bar{R}_j(d, x) \equiv R^+_j(D^j_d) + \sum_{i=j}^{m} f_i x_i, \tag{10}$$

which is an upper bound for $R(D^j_d, b)$, but can replace it in (9).

Proposition 5 (Linearized Regret)**.** *The regret guarantee $Z(b)$ for $b \in B(n)$ or its equivalent x can be computed with $\bar{R}_j(d, x)$ by*

$$Z(b) = \max_{j \in 1..m+1} \beta R^*(D^j_{d(j,\beta)}) - \bar{R}_j(d(j, \beta), x). \tag{11}$$

Proof. Consider two cases: **Case 1**: $\check{d}^j(b) > d(j, \beta)$. There is $C(D^j_{d(j,\beta)}; b) = j$, which implies $\bar{R}_j(d(j, \beta), x) = R(D^j_{d(j,\beta)}, b)$. **Case 2**: $d(j, \beta) \geq \check{d}^j(b)$. Then $C(D^j_{d(j,\beta)}; b) = 1$ by Proposition 4 so that $D^j_{d(j,\beta)}$ is dominated by D^1_0 under b by Proposition 3. This dominance means that replacing $R(D^j_{d(j,\beta)}, b)$ by its upper bound $\bar{R}_j(d(j, \beta), x)$ will not affect $Z(b)$, since D_0^1 always shows up in (11) by reason of $d(1,\beta)=0$, and there is always $R(D_0^1, b) = \bar{R}_1(0,x)$ from Case 1 with $1 = \check{d}^1(b) > d(1,\beta) = 0$. $\square$

Based on Proposition 5, a linear program for the ARM (LPARM) problem (1) with $\Pi = B(n)$ is formulated with bucket sizes x, and its regret guarantee z.

$$\text{LPARM} : \min \quad z$$

$$s.t. \quad z + \sum_{i=j}^{m} f_i x_i \geq \bar{G}^j_{d(j,\beta)} \quad j = 1, \cdots, m+1 \tag{12}$$

$$\sum_{i=1}^{m} x_i \leq n \tag{13}$$

$$0 \leq x_j \leq U_j \quad j = 1, \cdots, m \tag{14}$$

whereas for discrete problems, an integral constraint on x must be included. Note that constraint (12) degenerates to $z \geq \bar{G}^j_{d(j,\beta)}$ for $j = m+1$. It is assumed that $\sum L \leq n$ (if $\sum L \geq n$ then it is never optimal to accept more than $n - \sum L + L_m < L_m$ requests in fare class m, thus reducing L_m to the level $n - \sum L + L_m$ will not affect the problem.

Proposition 6 (Auxiliary LP)**.** *The value of* $\bar{G}^j_{d(j,\beta)}$ *for* $j = 1, \cdots, m+1$ *can be* *found by an auxiliary LP as follows:*

$$LPAUX : \bar{G}^{j}_{d(j,\beta)} = \max \beta \sum_{i=1}^{m} f_i x_i - \sum_{i=1}^{m} f_i y_i \tag{15}$$

$$s.t. \quad \sum_{i=1}^{m} x_i \le n \tag{16}$$

$$x_i \le y_i \ge L_i \quad i = 1, \cdots, j-1 \tag{17}$$

$$0 \le x \le U, \quad 0 \le y \le U, \tag{18}$$

and $\bar{G}^{j}_{d(j,\beta)}$ decreases in j but the decrement is bounded above by $f_j U_j$:

$$0 \le \bar{G}^{j}_{d(j,\beta)} - \bar{G}^{j+1}_{d(j+1,\beta)} \le f_j U_j \quad j = 1, \cdots, m. \tag{19}$$

Proof. Only discrete problems are dealt with here, as the results can be carried over to continuous problems by the limit process discussed earlier.

Since any step sequence D^j_d defines a feasible solution to the auxiliary LP by $\bar{x} = \arg\max\{\sum_{i=1}^{m} x_i f_i : \sum_i x_i \le n, x_i \le D^j_d[i], i = 1, \cdots, m\}, \bar{y} = \mathbf{1}\{i < j\} \max(L_i, \bar{x}_i), i = 1, \cdots, m)$ with $\bar{G}^j_d$ being its objective value, it suffices to show that there exists an optimal solution $x^*_i, y^*_i, i = 1, \cdots, m$ such that $P^* = (y^*_1, \cdots, y^*_{j-1}, U_j, \cdots, U_m)$ is a profile for a step sequence $D^j_{d^*}$ with $\bar{G}^j_{d^*}$ being the objective value. From the objective (15), y^*_i must be as small as possible given x^*_i (so $y^*_i = \mathbf{1}\{i < j\} \max(L_i, x^*_i), i = 1, \cdots, m$), and x^* must be the offline optimal booking decision given y^* (so $x^*_k = P^*[k]$ if $\sum_{i=1}^{k} P^*[i] \le n$ and $x^*_k = (n - \sum_{i=1}^{k} P^*[i])^+$ otherwise).

Consider two cases: $\beta \le 1$ and $\beta > 1$. Case 1: $\beta \le 1$. If there is $y^*_i > L_i$ for some $i < j$ (implying $y^*_i = x^*_i$), then the modified solution $y'_k = \mathbf{1}\{k \ne i\} y^*_k + \mathbf{1}\{k = i\} L_k, x'_k = \mathbf{1}\{k \ne i\} x^*_k + \mathbf{1}\{k = i\} L_k, k = 1, \cdots, m$ will not deteriorate the objective value. Let $y^* = y', x^* = x'$ and keep modifying until all $y^*_k = L_k, k < j$ and its P^*

becomes the profile of D_0^j. Case 2: $\beta > 1$. Let $\bar{k} = \min\{k \leq j : y_k^* < U_k\}$. If $\bar{k} \geq j-1$, then $y_k^* = U_k, k < j-1$ and clearly LBH(P^*) belongs to D^j. If $\bar{k} < j-1$, and suppose there is a $y_{\dot{k}}^* > L_i$ for some $\dot{k} > \bar{k}$ (implying $x_k^* = y_k^*, k \in \{\bar{k}, \dot{k}\}$), then let $y'_{\bar{k}} = y_{\bar{k}}^* + 1$, $y'_{\dot{k}} = y_{\dot{k}}^* - 1$, and $y'_k = y_k^*$ for $k \notin \{\bar{k}, \dot{k}\}$, and let $x' = x^* + (y' - y^*)$. Then x', y' is feasible and improves the objective value by $(\beta - 1)(f_{\bar{k}} - f_{\dot{k}}$, which contradicts the optimality of x^*, y^*, therefore, no such $\dot{k}$ can exist, and $y_k^* = L_k$ for all $k > \bar{k}$, which means LBH(P^*) belongs to D^j.

In both cases there is a $D_{b^*}^j = \text{LBH}(P^*) \in D^j$, and the corresponding objective value is clearly $\bar{G}_{d^*}^j$, with $R^*(D_{d^*}^j) = \sum_{i=1}^m f_i x_i^*$ and $R^+(D_{d^*}^j) = \sum_{i=1}^m f_i y_i^*$ (note that $y_i^* = 0, i \geq j$). Therefore, it must be true that $\bar{G}_{d^*}^j = \bar{G}_{d(j,\beta)}^j$.

To show that $\bar{G}_{d(j,\beta)}^j$ is non-increasing in j, note that the auxiliary LP model has more constraints as j increases, therefore $\bar{G}_{d(j,\beta)}^j$ decreases as j increases. Let x^*, y^* be an optimal solution to the model for j, let $y'_j = U_j$, and $y'_k = y_k^*$ for $k \neq j$, then x^*, y' is a feasible solution to the model for $j+1$, which gives $\bar{G}_{d(j,\beta)}^j - f_j U_j \leq \bar{G}_{d(j+1,\beta)}^{j+1}$. $\square$

The LPAUX model and the bounds on $\bar{G}_{d(j,\beta)}^j - \bar{G}_{d(j+1,\beta)}^{j+1}$ are crucial for solving the continuous problems in closed-form and studying the properties of the solutions.

Theorem 2 (Closed-form Optimal Policy)**.** *The LPARM model for continuous problems can be solved in closed-form as follows:*

$$x_i^* = \begin{cases} g_i^+, & i < \bar{u} \\ n - \sum_{j=1}^{\bar{u}-1} g_j^+, & i = \bar{u} \\ 0, & i > \bar{u} \end{cases} \tag{20}$$

$$z^* = \bar{G}_{d(\bar{u},\beta)}^{\bar{u}} - x_{\bar{u}} f_{\bar{u}} \tag{21}$$

where $g_j^+ \equiv (\bar{G}^j_{d(j,\beta)} - \bar{G}^{j+1}_{d(j+1,\beta)})/f_j, j = 1..m$ *and*

$$\bar{u} \equiv \max\{i \in \{1, \cdots, m\} : \sum_{j=1}^{i-1} g_j^+ < n\}. \tag{22}$$

Proof. Clearly x^*, z^* is a feasible solution. Consider the dual problem with y_j for (12), v for (13), and w_j for (14):

$$\max \quad \sum_{j=1}^{m+1} \bar{G}^j_{d(j,\beta)} y_j - nv - \sum_{i=1}^{m} U_i w_i \tag{23}$$

$$s.t. \quad \sum_{i=1}^{m+1} y_i = 1 \tag{24}$$

$$f_j \sum_{i=1}^{j} y_i \leq v + w_j \quad j = 1, \cdots, m \tag{25}$$

$$y, v, w \geq 0 \tag{26}$$

Let $v^* = f_{\bar{u}}$, $w^* = 0$, $y_i^* = (f_i^{-1} - f_{i-1}^{-1})v^*, i = 1..\bar{u}$ (let $f_0 = \infty$ for convenience), and $y_i^* = 0, i > \bar{u}$. Clearly v^*, w^*, y^* is a dual feasible solution with a dual objective value of $\bar{G}^{\bar{u}}_{d(\bar{u},\beta)} - (n - \sum_{j=1}^{\bar{u}-1} g_j^+)v^* = z^*$, which is identical to the objective of the primal objective for the solution x^*, z^*. Therefore, x^*, z^* must be an optimal solution. $\square$

Closed-form solutions make it possible to study the effect of β on optimal policies, which can help with conservatism control and being adaptive to real-time KPIs or environment changes.

Theorem 3 (Increasing Aggressiveness)**.** *For continuous problems, as β increases, the optimal policy becomes more aggressive by (1) increasing reserved seats $\sum_{j=1}^{i} x_j^*$ for higher fares for $i = 1, \cdots, m$ if $\beta \leq 1$, or (2) increasing reserved revenue potential $\sum_{j=1}^{i} f_j x_j^*$ for higher fares for $i = 1, \cdots, m$ if $\beta > 1$.*

Proof. By applying the Envelope Theorem to LPAUX for $\bar{G}^j_{d(j,\beta)}$, there is

$$\frac{\partial \bar{G}^j_{d(j,\beta)}}{\partial \beta} = R^*_j, \tag{27}$$

where $R^*_j = R^*(D^j_{d(j,\beta)})$.

For $\beta \leq 1$, it suffices to show $\partial b_i/\partial \beta \geq 0$ for $i = 1, \cdots, m$. From the definition of g^+_j, it is clear that for $i < \bar{u}$,

$$\frac{\partial b_i}{\partial \beta} = \sum_{j=1}^{i} \frac{\partial g^+_j}{\partial \beta} = \sum_{j=1}^{i} \frac{R^*_j - R^*_{j+1}}{f_j}. \tag{28}$$

Clearly $\partial b_i/\partial \beta \geq 0$ if $\{R^*_j : j = 1, \cdots, i+1\}$ is a decreasing sequence, i.e., $R^*_j \geq R^*_{j+1}$ for $j = 1, \cdots, i$, which is indeed the case when $\beta \leq 1$.

For $\beta > 1$, it suffices to show $\partial \sum_{j=1}^{i} f_j x_j/\partial \beta \geq 0$ for $j = 1, \cdots, m$. From the definition of g^+_j, it is clear that for $i < \bar{u}$,

$$\frac{\partial b_i}{\partial \beta} = \sum_{j=1}^{i} \frac{\partial f_j g^+_j}{\partial \beta} = \sum_{j=1}^{i} R^*_j - R^*_{j+1} = R^*_1 - R^*_{i+1} \geq 0. \tag{29}$$

□

Clearly, form (1) implies and thus is stronger than form (2). When $\beta > 1$ form (1) can be violated in this counter-example with three fares: $L = (0, 0, 0), U = (5, 5, 5), f = (100, 49, 24), n = 10$ and $\beta = 1.2$.

Theorem 3 gives a clear direction of policy adjustment with β, which may work with "displacement rate", a very important KPI according to Vinod (2021a). It is the ratio of spill (passengers rejected by the policy) to unconstrained demand by booking class. If higher displacement rates are observed for higher fare classes, then the policy is obviously not protecting enough seats for the higher fare requests and losing potential revenues, which serves as a signal to take corrective action by increasing β.

Theorem 4 (Optimality of SNBL Policies)**.** *For both the* continuous and discrete *m-fare capacity control problems with demand bounds, if the adversary can observe the booking decisions and manipulate future booking requests, then no algorithm can provide a regret guarantee strictly less than an optimal SNBL policy.*

Proof. Consider an arbitrary online algorithm, which does not have to be a SNBL policy. Let z^* be the regret guarantee by an optimal SNBL policy found by LPARM (with integrity constraints in case of discrete problems). The adversary can start with $j = m$, and go through the following steps:

1. Send in U_j requests, observe x_j^+, the number of requests accepted, and calculate $z_j^+ = \beta R^*(D^j_{d(j,\beta)}) - (R_j^+(D^j_{d(j,\beta)}) + \sum_{i=j}^m f_i x_i^+)$.
2. If $z_j^+ \geq z^*$, then send in the rest of $D^j_{d(j,\beta)}$ and STOP.
3. Let $j = j - 1$, and GOTO step 1 if $j > 0$.

This procedure must end with some $j > 0$ and $z_j^+ \geq z^*$, where z_j^+ is the best regret with all future requests in $D^j_{d(j,\beta)}$ accepted. If it stops with $j = 0$, a feasible solution $\bar{x} = x^+, \bar{z} = \max_{j=1}^m z_j^+$ is obtained with $\bar{z} < z^*$, which contradicts z^* being the optimal objective of LPARM. □

This shows that SNBL policies are surprisingly powerful when it comes to regret guarantees. Note that all online algorithms are considered, static or dynamic, deterministic or randomized. Deterministic algorithms always yield the same output for the same input, whereas "randomized" algorithms make some choices based on the draw of a random number. Note that for discrete problems with an adversary agnostic to the randomized booking decisions, randomization can help achieve the regret guarantee given Theorem (2) after relaxing the integrity constraint. The scheme in Lan et al. (2008) can be readily employed to randomize among a set of discrete SNBL policies and achieve on average a better guarantee.

It is possible to implement dynamic policies that recalibrate the booking limits in the booking process to achieve better guarantees, while static policies maintain constant booking limits throughout the booking horizon. Dynamic policies can be helpful in practice when the scenario is not worst-case, otherwise they can not improve on the regret guarantee. A dynamic policy can be derived as a series of "static" policies at each point of recalibration, in the same way as shown in Lan et al. (2008).

5. Simulation Study

The major interest is to find out if the control parameter β can be used to adapt to different environments, and if less conservative policies can be obtained by adjusting β. Three environments are set up with the demand in all fare classes simultaneously weak, medium, or strong, by having different levels of mean demand, while keeping the same variances to minimize their influences. Each environment is driven by a beta distribution: for a fare class $i = 1, \cdots, m$ a random sample $V_i \in [0, 1]$ is independently drawn from it, with the number of class i requests given by $L_i + (U_i - L_i)V_i$. Three distributions are chosen: Beta$(23/16, 69/16)$ for weak case, Beta$(4, 4)$ for medium case, and Beta$(69/16, 23/16)$ for strong case, with the same standard variation of $1/6$ and $0.25, 0.5, 0.75$ as their respective mean. Each environment generates 10000 demand scenarios, which are fed to the LPARM policies whose β values are $\beta_i = i/30, i = 1, \cdots, 90$ to estimate average revenues and the 99% confidence intervals. These numerical results faciliate finding a best β with maximal average revenue, the policy of which may be less conservative and comparable in

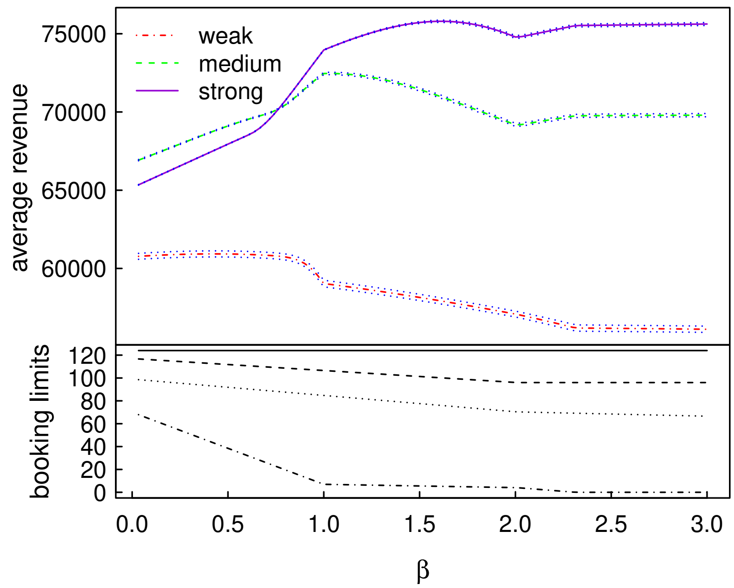

Figure 1: Average revenues with 99% confidence intervals in blue dotted lines closely wrapping the revenue curves, and the booking limits of LPARM policies for $\beta \in (0, 3]$.

average revenues to the popular EMSR policies of EMSRa (Belobaba 1989) and EMSRb (Belobaba 1992).

The example is adapted from Talluri et al. 2004, §2.2.3.4, with $f = (1050, 567, 527, 350), n = 124$. The demand is independent across fare classes with mean $\mu = (17.3, 45.1, 73.6, 19.8)$, and the arrivals are LBH. The bounds are given by $L = 0.4\mu, U = 1.6\mu$, assuming that the demand for all fare classes has the same coefficient of variation of 0.2, and the bounds are three standard deviations from the mean.

The upper plot in Figure 1 shows that as demand gets stronger, the best β^* (in the second column of Table 1) to maximize the average revenue indeed gets bigger, while the lower plot shows the booking limits get more aggressive as β increases, protecting more seats for higher fares. Table 1 also provides a comparison of revenues, which shows that with the ideal β^* value, a less conservative policy can be obtained to

Demand	Best β^*	Average revenue $\pm$ Standard deviation			
		$\beta = \beta^*$	$\beta = 1$	EMSRa	EMSRb
Weak	0.433	60930±75	59035±82	60932±75	60922±75
	gap %	0.00%	3.11%	0.00%	0.02%
Medium	1.033	72463±29	72459±28	73392±36	73337±38
	gap %	1.27%	1.27%	0.00%	0.08%
Strong	1.600	75797±29	73964±9	77336±28	77159±34
	gap %	1.99%	4.36%	0.00%	0.23%

Table 1: The best β^*, average revenues for LPARM at $\beta = \beta^*$ and $\beta = 1$ (for absolute regret), EMSRa, EMSRb, and the percentage gap from the highest revenue of the four entries.

achieve up to 3% improvement in average revenues over the absolute regret ($\beta = 1$) policy, and a less than 2% revenue gap from the EMSR policies. These revenue differences are quite reliable as common random numbers are used and the standard deviations are around 0.1% of the revenues.

6. Conclusion

This paper develops an adaptive robust method for the classical single-leg capacity control problem with the practically desirable ability to fine-tune the level of conservatism and respond to changing environments. A new technique of joint reduction based on the concept of joint dominance is employed to reduce the scenarios to very few extreme ones, which may depend on the ARM parameter β. A linear program LPARM is formulated based on these extreme scenarios, and the optimal solution is obtained in closed-form. The SNBL policy recommended by LPARM is proved optimal among all online policies. The closed-form optimal policy becomes more aggressive as the β parameter increases, giving a clear direction of policy change by adjusting β. The simulation study shows that as the environment has a more optimistic outlook on demand, the β parameter should increase to take the opportunity. It also finds that less conservative policies with more revenues can be found when β

is chosen properly. The average revenues of such policies can be very close to those of the EMSR policies.

From a research perspective, the adaptive robust optimization with competitive analysis of online algorithms is very promising in RM. Some directions for further reasearch can be seen. As this paper only considers capacity control, it is natural to extend it to consider joint decisions with overbooking. The interface between an adaptive robust model and the KPIs monitored in real-time is another interesting direction. How to develop a robust adaptive method for the network RM problem remains a challenging future research topic.

References

Albers, S., 2003. Online algorithms: a survey. Mathematical Programming 97, 3–26.

Ball, M.O., Queyranne, M., 2009. Toward robust revenue management: Competitive analysis of online booking. Operations Research 57, 950–963.

Barz, C., 2007. Risk-averse capacity control in revenue management. volume 597. Springer Science & Business Media.

Belobaba, P.P., 1989. Or practice—application of a probabilistic decision model to airline seat inventory control. Operations research 37, 183–197.

Belobaba, P.P., 1992. Optimal vs. heuristic methods for nested seat allocation, in: Presentation at ORSA/TIMS Joint National Meeting.

Birbil, Ş.İ., Frenk, J., Gromicho, J.A., Zhang, S., 2009. The role of robust optimization in single-leg airline revenue management. Management Science 55, 148–163.

Feng, Y., Xiao, B., 1999. Maximizing revenues of perishable assets with a risk factor. Operations Research 47, 337–341.

Gönsch, J., 2017. A survey on risk-averse and robust revenue management. European Journal of Operational Research 263, 337–348.

Kouvelis, P., Yu, G., 2013. Robust discrete optimization and its applications. volume 14. Springer Science & Business Media.

Lan, Y., 2021. Adjustable regret for continuous control of conservatism and competitive ratio analysis. URL: https://arxiv.org/abs/2105.05536, doi:10.48550/ARXIV.2105.05536.

Lan, Y., Ball, M.O., Karaesmen, I.Z., 2011. Regret in overbooking and fare-class allocation for single leg. Manufacturing & Service Operations Management 13, 194–208.

Lan, Y., Gao, H., Ball, M.O., Karaesmen, I., 2008. Revenue management with limited demand information. Management Science 54, 1594–1609.

Lancaster, J., 2003. The financial risk of airline revenue management. Journal of Revenue and Pricing Management 2, 158–165.

Ma, W., Simchi-Levi, D., Teo, C.P., 2021. On policies for single-leg revenue management with limited demand information. Operations Research 69, 207–226.

Palmatier, G.E., Crum, C., 2003. Enterrise sales and operations planning. Ross Poublishing Inc. and APICS .

Perakis, G., Roels, G., 2008. Regret in the newsvendor model with partial information. Operations research 56, 188–203.

Perakis, G., Roels, G., 2010. Robust controls for network revenue management. Manufacturing & Service Operations Management 12, 56–76.

Rusmevichientong, P., Topaloglu, H., 2012. Robust assortment optimization in revenue management under the multinomial logit choice model. Operations research 60, 865–882.

Sierag, D., van der Mei, R., 2016. Single-leg choice-based revenue management: a robust optimisation approach. Journal of Revenue and Pricing Management 15, 454–467.

Talluri, K.T., Van Ryzin, G., Van Ryzin, G., 2004. The theory and practice of revenue management. volume 1. Springer.

Vinod, B., 2021a. Airline revenue planning and the covid-19 pandemic. Journal of Tourism Futures .

Vinod, B., 2021b. An approach to adaptive robust revenue management with continuous demand management in a covid-19 era. Journal of Revenue and Pricing management 20, 10–14.