Robust Controls for Network Revenue Management

Georgia Perakis • Guillaume Roels

Sloan School of Management, MIT, E53-359, Cambridge, MA 02139, USA
Anderson School of Management, UCLA, B511, Los Angeles, CA 90095, USA.
georgiap@mit.edu • groels@anderson.ucla.edu

Revenue management models traditionally assume that future demand is unknown, but can be represented by a stochastic process or a probability distribution. Demand is however often difficult to characterize, especially in new or nonstationary markets.

In this paper, we develop robust formulations for the capacity allocation problem in revenue management, using the maximin and the minimax regret criteria, under general polyhedral uncertainty sets. Our approach encompasses the following open-loop controls: partitioned booking limits, nested fare classes by origin-destination pairs, Displacement-Adjusted Virtual Nesting, and fixed bid prices. We also characterize the optimal booking policy under interval uncertainty; while partitioned booking limits are optimal under the maximin criterion, some nesting is desirable under the minimax regret criterion.

Our numerical analysis reveals that robust controls can outperform the classical heuristics for network revenue management, while achieving the best performance in the worst case. Our models are scalable to solve practical problems, because they combine efficient solution methods (small mixed-integer and linear optimization problems) with very modest data requirements.

1. Introduction

The field of revenue management (RM) originated in the airline industry as a way to efficiently allocate fixed capacity to different classes of customers. Since then, its scope has expanded, combining capacity rationing with pricing tactics, and the concept has been applied to a variety of industries, such as hotels, rental cars, and media. See Talluri and van Ryzin (2004b) for an overview of the field.

The decision to accept or reject an incoming customer is often made without knowing future demand. Traditional RM models assume that future demand is unknown but can be characterized by either a stochastic process representing the customer arrival process (the *dynamic* models) or a probability distribution representing the aggregate number of customers (the *static* models).

Accurate forecasting is key to effective RM. The best forecasts are typically obtained by gathering demand information from different sources (e.g., historical sales data, recent bookings, competitive environment), interpreting it carefully (e.g., sales data are only censored demand data), and combining alternative forecasting methods, such as time series, regression models, and subjective opinion (Boyd and Bilegan 2003). In general, quantitative forecast methods are favored in stable business environments, where large amounts of historical data are available and can be used to calibrate econometrics models. On the other hand, simple forecasting methods are preferable when the business environment is new, nonstationary, or subject to random shocks (e.g., pandemic crises), and when there is a large amount of demand quantities to be estimated (about 2 million every day for a medium-size airline, see Talluri and van Ryzin 2004b). Given the reliance of RM models on quantitative demand information, one may wonder if RM is really effective without relevant historical data.

In this paper, we investigate the problem of allocating fixed network capacity to different classes of customers without historical data. Instead of assuming that demand follows a probability distribution, as is traditionally done, we only assume that it lies in a *polyhedral uncertainty set*, giving enough flexibility to model information about the range, moments, shape, correlation of the demand, data censorship, as well as subjective opinions. This representation of uncertainty captures the stochastic nature of the problem, but remains simple to estimate.

We consider the *maximin* and the *minimax regret* criteria for decision-making under uncertainty. The maximin criterion guarantees a minimum level of profit, and is more appropriate for risk-averse decision-makers. In contrast, the minimax regret criterion minimizes the opportunity cost from not knowing the demand distributions and gives rise to less conservative recommendations.

The booking policy considered in this paper generalizes the following controls, which are frequently used in practice: *nested booking limits*, *partitioned booking limits*, *Displacement-Adjusted Revenue Virtual Nesting*, and *fixed bid prices*. We develop simple formulations to compute the worst-case performance (minimum revenue or maximum regret) of any policy based on these controls. We also characterize the structure of the most robust policy under interval uncertainty: While partitioned booking limits are optimal under the maximin criterion, some nesting is desirable under the minimax regret criterion. Our numerical study suggests that the proposed robust policies outperform the traditional open-loop controls, especially in the presence of correlation or data censorship. Our approach is scalable to solve large network RM problems as it combines efficient solution procedures with very modest data requirements.

Literature Review. The single-resource capacity control problem was introduced by Littlewood (1972) with two classes of customers arriving sequentially, and was subsequently extended to multiple classes of customers (see Talluri and van Ryzin 2004b for a review). With sequential arrivals, the optimal control can be achieved with nested protection levels, nested booking limits, or bid-price tables (Talluri and van Ryzin 2004b). Several heuristics, such as EMSR-a (Belobaba 1987) and EMSR-b (Belobaba 1992), also perform well in practice.

The network RM problem is significantly more complex and little is known about the optimal policy. Consequently, the network RM problem is often solved heuristically, either by approximating the revenue-to-go function in the dynamic program (Bertsimas and Popescu 2003, Adelman 2006), or by restricting the set of feasible policies. Commonly used controls are partitioned booking limits, virtual nesting controls, bid prices (see Williamson 1988, 1992 for a numerical comparison among these controls), and nesting fare classes by itinerary (Curry 1990 and Chi 1995). Partitioned booking limits and bid prices are usually obtained with mathematical programming formulations, such as the Deterministic Linear Program (DLP) (Dror et al. 1988), the Randomized Linear Program (RLP) (Talluri and van Ryzin 1999), and the Probabilistic Nonlinear Program (PNLP) (Wollmer 1986). Recent research however shows that improved bid prices can be obtained with continuous-time optimal control formulations (Akan and Ata 2006) and approximate dynamic programming formulations (Adelman 2006, Topaloglu 2006, Talluri 2007). The controls can then be fine-tuned with a stochastic gradient algorithm based on demand samples (Bertsimas and de Boer 2005 and van Ryzin and Vulcano 2005).

Traditional models of demand in RM have been subject to criticism, because they do not capture sell-ups, buy-downs, demand correlation, group arrivals, and nonsequential order of arrivals among others. To overcome these limitations, the following approaches have been proposed: consumer-choice behavior, data-driven optimization, and robust optimization. Consumer-choice models aim at better understanding the behavior of customers, by disaggregating demand at the customer level; See Talluri and van Ryzin (2004a) and Zhang and Cooper (2005) for recent developments in RM. In these models, the source of uncertainty is no longer the number of requests for fares, but the customers' probabilities of purchase when a limited set of products is offered. In contrast, data-driven optimization and robust optimization are distribution-free techniques. While data-driven optimization use historical data to learn and update the decisions, robust optimization only requires limited or no information about demand. Hence, data-driven optimization is more effective in stable business environments with large amount of available data, while robust optimization is more suitable in nonstationary or new business environments where expert judgment is more critical.

Data-driven optimization for RM was pioneered by van Ryzin and McGill (2000), who developed an adaptive algorithm to determine booking limits using sales data. More recently, Bertsimas and de Boer (2005) and van Ryzin and Vulcano (2005) proposed stochastic gradient methods for improving booking policies in network RM, using demand samples. In stochastic inventory management, nonparametric methods have recently received a lot of attention, after the work by Godfrey and Powell (2001), Levi et al. (2005), and Huh and Rusmevichientong (2006). In dynamic pricing, Eren and Maglaras (2006) used historical data to estimate the entropy-maximizing demand distribution while Rusmevichientong et al. (2006) and Besbes and Zeevi (2006) analyzed multiproduct pricing problems for a single and multiple resources, respectively.

In contrast, robust optimization models do not require historical data. Robust optimization has recently received a lot of attention since Ben-Tal and Nemirovski (1999) and Bertsimas and Sim (2004) among others developed a methodology to make robust but not too conservative decisions. Robust approaches have been widely used in inventory control with the maximin criterion (e.g., Scarf 1958, Gallego and Moon 1993, Gallego et al. 2001, Bertsimas and Thiele 2006, Ben-Tal et al. 2005) and the minimax regret criterion (Yue et al. 2006, Perakis and Roels 2006). In RM, Birbil et al. (2006), Ball and Queyranne (2006), and Lan et al. (2006) analyzed robust nested booking limits on a single leg. Birbil et al. (2006) developed efficient algorithms to compute the maximin booking limits, partitioned or nested, under ellipsoidal uncertainty. Ball and Queyranne (2006) studied nested booking limits in a single problem using the competitive ratio, with no information about the demand, and Lan et al. (2006) generalized their results to interval uncertainty. While they also cover the minimax regret, our study is however more general as it also addresses network problems, general polyhedral uncertainty sets, and general open-loop booking limit controls.

Outline. The paper is organized as follows. Section 2 reviews the classical network RM problem, and §3 introduces the decision-theoretic framework. In §4, we propose a mixed-integer formulation for the minimax regret network RM problem, characterize its complexity and the structure of its optimal solution under interval uncertainty, and then develop simpler formulations and approximations when fare classes are either nested by origin-destination or partitioned. In §5, we propose a mixed-integer formulation for the maximin problem, and develop a simpler linear formulation when classes are partitioned. Numerical examples in §6 illustrate the performance of the proposed policies, in comparison to existing heuristics. Finally, §7 provides concluding remarks.

Notations. We begin by introducing some notational conventions. Vector (resp. matrices) are denoted in small (resp. capital) bold letters. For a vector $\mathbf{x}$, x_j denotes its jth component; similarly, for a matrix $\mathbf{A}$, $\mathbf{A}_j$ represents the jth column and $\mathbf{a}_i$ the ith row. All vectors are column vectors, and $\mathbf{x}'$ is the vector transpose. The function $\min\{\mathbf{x}, \mathbf{y}\}$ takes the componentwise minimum of vectors $\mathbf{x}$ and $\mathbf{y}$, and the function $\mathbf{x}^+$ takes the componentwise maximum of $\mathbf{x}$ and $\mathbf{0}$. Let $\mathbf{1}$ be a vector of ones. Finally, we often use the terminology of the airline industry (e.g., seats), but our analysis can be applied to any network RM problem.

2. Problem Statement

We first introduce the dynamic network RM problem following the presentation by Cooper (2002). Consider a network with K resources (e.g., flights, night stays) and N products, differentiated by origin-destinations and fare classes (ODF). Customers arrive according to a certain (continuous or discrete) stochastic process over a finite time interval; let d_j be the random total demand for product j. There are $\mathbf{c}$ units available of resources. Each product j has a unit revenue r_j and consumes $\mathbf{A}_j$ units of resources.

We seek a policy π that maximizes the expected revenues $\mathbf{r}'E[\mathbf{n}^\pi]$, where $\mathbf{n}^\pi$ is the vector of total number of accepted requests when policy π is in use. The policy needs to satisfy (almost surely, denoted by a.s.) the capacity constraints, i.e., $\mathbf{A}\mathbf{n}^\pi \leq \mathbf{c}$. The accepted requests are nonnegative and cannot exceed the total demand, i.e., $\mathbf{0} \leq \mathbf{n}^\pi \leq \mathbf{d}$. In addition, the policy is required to be non-anticipating. That is, the acceptance/rejection decision at each time t should be based only on the information acquired up to time t. Let Π be the set of non-anticipating policies. The problem can then be formulated as follows:

$$\begin{aligned} \sup_{\pi \in \Pi} \quad & \mathbf{r}'E[\mathbf{n}^\pi], \\ s.t. \quad & \mathbf{A}\mathbf{n}^\pi \leq \mathbf{c} \qquad (a.s.), \end{aligned} \tag{1}$$

$$\mathbf{0} \le \mathbf{n}^{\pi} \le \mathbf{d} \quad (a.s.).$$

When the arrival process is discrete, Problem (1) can be formulated as a dynamic program: At each (discrete) time t, given the level of available capacity, one needs to decide whether to accept or reject the requests arriving in period t, in order to maximize the total expected revenues until the end of the time horizon. The dynamic program formulation highlights the structure of the optimal policy: a request for product j at time t is accepted if and only if its fare r_j exceeds the opportunity cost from consuming $\mathbf{A}_j$ units of capacity, given the level of available capacity at time t (e.g., see Talluri and van Ryzin 2004b).

However, the dynamic program is rarely solved to optimality in practice due to its large size. Its complexity can be reduced either by approximating the revenue-to-go function (e.g., see Bertsimas and Popescu 2003 and Adelman 2006), or by considering a subset of feasible policies, such as booking limits and bid prices.

Booking Limits. Booking limits set a maximum on the number of requests that can be accepted for a set of products. In a single-leg problem, products are naturally ordered by fare and it is optimal to define booking limits over a *nest* of products. The booking limit on nest j imposes a maximum on the total number of requests that can be accepted for products with a fare lower than or equal to r_j.

In a network environment, however, there is no natural ordering of products, and it is not clear how to choose the nests. One alternative is to define booking limits for every product. The total capacity is accordingly *partitioned* into N buckets. With partitioned booking limits, Problem (1) simplifies to a PNLP, see Talluri and van Ryzin (2004b). When the random demands are replaced by their mean, the problem of finding partitioned booking limits reduces to the following DLP:

$$\begin{aligned} \max_{\mathbf{y}} \quad & \mathbf{r}'\mathbf{y}, \\ s.t. \quad & \mathbf{A}\mathbf{y} \le \mathbf{c}, \\ & \mathbf{y} \le \boldsymbol{\mu}, \qquad (DLP) \\ & \mathbf{y} \ge \mathbf{0}. \end{aligned}$$

Another approach is to decompose the network problem into K single-leg problems, one for each resource, and to define different nests on each resource. This methodology, called the *Displacement-Adjusted Virtual Nesting* (DAVN), performs extremely well in practice (Smith et al. 1992). The approach relies on modifying the fares to account for network effects; specifically, the fare of product j on leg k is calculated as $r_j - \sum_{l \neq k : a_{lj}=1} p_l$, where $\mathbf{p}$ are the shadow prices associated with the capacity constraints in the DLP. For every leg, products are ordered in decreasing order of adjusted fare, and then grouped into a certain number of buckets. Nested booking limits for these buckets are then computed, on every leg, using either EMSR or DLP (see Bertsimas and de Boer 2005 for details). An incoming request is accepted if there is sufficient capacity and if, on every leg, the booking limit of the associated bucket has not been reached.

Bid-Price Controls. Also widely used are bid-price controls. In every period and for every vector of resources, a price p_k is associated with each resource k. A request for product j is accepted at time t if and only if the collected fare r_j exceeds the implicit cost of consuming resources, $\mathbf{A}_j'\mathbf{p}$. Bid price controls are in general not optimal (Talluri and van Ryzin 1998), yet they are widely used in practice. Usually, the shadow prices of the DLP (possibly randomized, see Talluri and van Ryzin 1999) are chosen as bid prices, but more recent methods use approximate dynamic programming (Adelman 2006 and Topaloglu 2006).

3. Decision-Theoretic Framework

In this paper, we assume only partial information about demand. Specifically, we assume that the aggregate demand $\mathbf{d}$ belongs to a polyhedral uncertainty set $\mathcal{P}$ and make no assumption about the arrival sequence. Let $\mathcal{D}$ be the set of multivariate stochastic processes, such that $\mathbf{d} \in \mathcal{P}$, and let $\mathcal{F}$ be the feasible decision set, assumed to be compact. We denote by $R(\mathbf{y}, \mathbf{D})$ the revenue associated with a decision $\mathbf{y} \in \mathcal{F}$ when the demand process $\mathbf{D}$ is realized.

Decision-Making Criteria. Because the expected utility maximization criterion has no meaning in a distribution-free environment, different decision criteria need to be considered. In this paper, we use the maximin and the minimax regret criteria, and refer to Ball and Queyranne (2004) and Lan et al. (2006) for an analysis of the competitive ratio. The minimax regret criterion is less conservative than the maximin criterion because it benchmarks the decision made under uncertainty $\mathbf{y}$ against what would have been optimal to do in more informed circumstances. In fact, the maximin criterion is better suited for risk-averse decision-makers, as it guarantees a minimum level of revenue.

- The **maximin** criterion selects the decision that maximizes the worst-case revenue, where the worst case is taken over all demand processes under consideration, that is,

$$\varphi^* = \max_{\mathbf{y}\in\mathcal{F}} \min_{\mathbf{D}\in\mathcal{D}} R(\mathbf{y}, \mathbf{D}). \tag{2}$$

- The **minimax regret** criterion selects the decision that minimizes the maximum regret, where the maximum is taken over all demand processes from $\mathcal{D}$, that is,

$$\rho^* = \min_{\mathbf{y}\in\mathcal{F}} \rho(\mathbf{y}), \tag{3}$$

  where the regret $\rho(\mathbf{y})$ is defined as the maximum additional revenue that could have been obtained with full information about the demand process, i.e.,

$$\rho(\mathbf{y}) = \max_{\mathbf{D}\in\mathcal{D}} \left\{ \max_{\mathbf{z}\in\mathcal{F}} R(\mathbf{z}, \mathbf{D}) - R(\mathbf{y}, \mathbf{D}) \right\}. \tag{4}$$

  The competitive ratio is similar to the minimax regret, but measures the regret in relative, and not absolute, terms. It is in fact more risk averse than the minimax regret because maximizing the competitive ratio with a linear utility function is equivalent to minimizing the maximum regret with a logarithmic utility function.

Uncertainty Set. We now discuss the choice of the uncertainty set $\mathcal{P}$. If the range of the demand is known to be equal to $[l_j, u_j]$, for every product j, then $\mathcal{P} = \{\mathbf{d} : \mathbf{l} \leq \mathbf{d} \leq \mathbf{u}\}$. In particular, $\mathcal{P}$ represents the set of all demand realizations.

When demand is characterized differently than by its range, however, one may want to consider a smaller set than the set of all possible demand realizations. For instance, when only the mean of all demands is known, the set of all demand realizations is the nonnegative orthant, which is arguably conservative. Consistently with the recent developments in robust optimization (see Ben-Tal and Nemirovski 1999 and Bertsimas and Sim 2004), we denote by $\mathcal{P}(\eta)$ the polyhedral uncertainty set, such that $\Pr[\mathbf{d} \in \mathcal{P}(\eta)] \geq \eta$. Accordingly, the maximin revenue φ^* and the minimax regret ρ^*, defined over the all multivariate stochastic processes with $\mathbf{d} \in \mathcal{P}(\eta)$, are guaranteed only with probability η. In particular, φ^* can be viewed as the distribution-free analog of the Value-at-Risk criterion.

When the polyhedron is only defined by intervals, the probability of the demand vector satisfying the constraints can be computed explicitly, e.g., using Markov's, Chebyshev's, or Gauss' inequalities (Popescu 2005). For general uncertainty sets, the probabilistic guarantee can be found by solving a moment bound problem (see Bertsimas and Popescu 2005 and Popescu 2005). In effect, moment constraints can specify (or bound) the mean, variance, and correlation among the different demands, as well as the probability that demand exceeds a certain threshold (as when only censored demand information is available), and shape constraints can specify the symmetry or the mode of the distribution. If $\mathcal{P}$ is represented as the intersection of a polynomial (in the problem data) number of hyperplanes, the moment bound problem can be solved in polynomial time (in the problem data), see Bertsimas and Popescu (2005).

In general, building the uncertainty set $[l_j, u_j]$ around the median is more robust than around its mean. Indeed, Perakis and Roels (2006) showed that, in the context of the newsvendor model, knowing the median is in general more informative than knowing the mean. Because of the strong connection between RM and the newsvendor model, we expect the same information levels to hold here. Incidentally, the median is also easier to estimate than the mean, for it is less affected by censored demand data.

Open Loop Booking Limit Controls. In this paper, we investigate the robustness of booking limit controls and assume standard nesting (as opposed to theft nesting, see Bertsimas and de Boer 2005). For every set of products $S \in \mathcal{S}$, we define a booking limit y_S. Let x_j be the realized sales of product j. Then, the booking limit control ensures that $\sum_{j \in S} x_j \leq y_S$, for every set $S \in \mathcal{S}$. We further assume that the controls are open loop, i.e., they are not state-dependent.

These controls generalize the partitioned booking limits (take $\mathcal{S} = \cup_{j=1}^{N}\{j\}$), the nested booking limits on a single leg (take $\mathcal{S} = \cup_{j=1}^{N}\{j, j+1, ..., N\}$), the DAVN booking limits (take $\mathcal{S} = \cup_{k=1}^{K} \cup_{j=1}^{N^{\max}} \{B_j^k, B_{j+1}^k, ..., B_{N^{\max}}^k\}$ where B_j^k is the set of classes in bucket j on leg k), and the fixed bid prices (take $\mathcal{S} = \cup_{j=1}^{N}\{j\}$, and set $y_j = u_j$ if $r_j \geq \mathbf{p}'\mathbf{A}_j$, and $y_j = 0$ otherwise). They cannot, however, substitute for bid price tables (in which a vector of bid prices is specified for each level of capacity) or theft nesting, because these controls are state-dependent.

4. Minimax Regret

In this section, we first formulate and characterize the general minimax regret problem. We then analyze in more details a specific booking limit policy, which nests fare classes by origins and destinations. We show that, for this particular policy, the minimax regret problem (3) can be formulated as a a linear optimization problem (LP) with an exponential number of constraints, which leads to a sequence of lower bound approximations. We also provide an explicit solution for the minimax regret nested booking limits on a single leg and propose an upper bound approximation on the minimax regret for partitioned booking limits.

4.1 Problem Formulation

We first formulate the inner problem in (3) as a mixed-integer optimization problem (MIP), for given open-loop booking limits $\mathbf{y}$, and characterize its complexity. Let $\mathbf{x}$ be the realized sales, in the worst case, under policy $\mathbf{y}$. The policy $\mathbf{y}$ is benchmarked against the *perfect hindsight* policy $\mathbf{z}$, determined *after* observing the demand process. Clearly, the realized sales under policy $\mathbf{z}$ are exactly equal to $\mathbf{z}$. The maximum regret $\rho(\mathbf{y})$ measures the maximum difference in revenues between the perfect hindsight policy $\mathbf{z}$ and policy $\mathbf{y}$, i.e., $\max\{\mathbf{r}'\mathbf{z}-\mathbf{r}'\mathbf{x}\}$, where $\mathbf{x}$ are the realized sales under the booking policy $\mathbf{y}$.

We now formulate the constraints that $\mathbf{z}$ and $\mathbf{x}$ need to satisfy. Let $\mathbf{d}$ be the demand vector. From the preceding discussion, we assume that $\mathbf{d} \in \mathcal{P}(\eta)$, where $\mathcal{P}(\eta)$ is a polyhedron. By definition of our booking limit policy, $\sum_{j \in S} x_j \leq y_S$, for all $S \in \mathcal{S}$. One cannot sell more than the demand; therefore, $\mathbf{z} \leq \mathbf{d}$ and $\mathbf{x} \leq \mathbf{d}$. Moreover, one cannot sell more than the available capacity; therefore $\mathbf{Az} \leq \mathbf{c}$ and $\mathbf{Ax} \leq \mathbf{c}$. Sales are also nonnegative, i.e., $\mathbf{z} \geq \mathbf{0}$ and $\mathbf{x} \geq \mathbf{0}$.

Finally, we model the dynamic dimension of the system. Because arrivals are sequential, if the realized sales x_i are less than the optimal sales z_i for some product i, while the booking limits are not constraining, i.e. $\sum_{j \in S} x_j < y_S$ for all $S \in \mathcal{S}$ such that $i \in S$, then one should have run out of capacity. In particular, one should have accepted requests for other products, which would have been rejected under the optimal policy $\mathbf{z}$, and these requests have depleted the resources to the point that all requests for i cannot be met. That is, the maximum regret needs to be optimized over all possible sequences of arrivals. Although there is virtually an infinite number of sequences, we show next how to formulate this problem with at most $|\mathcal{S}| + N + K$ binary variables.

Let α_j, $j = 1, ..., N$, be a binary variable, equal to 1 if $x_j = d_j$, equal to zero otherwise. Because $\mathbf{x} \leq \mathbf{d}$, we can formulate this condition as follows: $\mathbf{d} \leq \mathbf{x} + M(\mathbf{1} - \boldsymbol{\alpha})$, where M is a large number.

Let β_S, $S = 1, ..., |\mathcal{S}|$, be a binary variable, equal to 1 if the booking limit on products in S has been reached, i.e., $\sum_{j \in S} x_j = y_S$, and equal to zero otherwise. Because $\sum_{j \in S} x_j \leq y_S$, for all $S \in \mathcal{S}$, we can formulate this condition as follows: $\sum_{j \in S} x_j \geq \beta_S y_S$ for all $S \in \mathcal{S}$.

Finally, let γ_k, $k = 1, ..., K$, be a binary variable, equal to 1 if the k-th capacity constraint is binding with the realized sales $\mathbf{x}$, i.e., $\mathbf{a}_k'\mathbf{x} = c_k$, equal to zero otherwise. Because $\mathbf{Ax} \leq \mathbf{c}$ and $\mathbf{Ax} \geq \mathbf{0}$, we can formulate this condition as follows: $\mathbf{a}_k'\mathbf{x} \geq c_k \gamma_k$ for all $k = 1, ..., K$.

If the realized sales are less than the optimal sales, i.e., $x_j < z_j$, or alternatively, if $x_j < d_j$ (because $z_j \leq d_j$), then either one of the booking limits has been reached, i.e., $\sum_{S: j \in S} \beta_S \geq 1$, or one of the resources has been depleted, i.e., $\mathbf{A}_j' \gamma \geq 1$. on the other hand, if no booking limit has been reached, then either the sales equal the demand, i.e., $\alpha_j = 1$, or one of the resources has been depleted. Finally, if no capacity constraint is binding, then either the sales equal the demand, or one of the booking limits have been reached. Combining these statements observations to the following constraint:

$$\mathbf{A}_j' \boldsymbol{\gamma} + \alpha_j + \sum_{S: j \in S} \beta_S \geq 1, \quad j = 1, ..., N.$$

We are now ready to present the MIP formulation of the maximum regret:

$$\begin{aligned}
\rho(\mathbf{y}) = \max_{\mathbf{z}, \mathbf{x}, \mathbf{d}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}} \quad & \mathbf{r}'\mathbf{z} - \mathbf{r}'\mathbf{x}, & \\
s.t. \quad & \mathbf{Az} \leq \mathbf{c}, & \\
& \mathbf{0} \leq \mathbf{z} \leq \mathbf{d}, & \\
& \mathbf{Ax} \leq \mathbf{c}, & \\
& \textstyle\sum_{j \in S} x_j \leq y_S, & S \in \mathcal{S}, \\
& \mathbf{0} \leq \mathbf{x} \leq \mathbf{d}, & \\
& \mathbf{d} \leq \mathbf{x} + M(\mathbf{1} - \boldsymbol{\alpha}), & \\
& \textstyle\sum_{j \in S} x_j \geq \beta_S y_S, & S \in \mathcal{S}, \\
& \mathbf{a}_k'\mathbf{x} \geq c_k \gamma_k, & k = 1, ..., K, \\
& \mathbf{A}_j' \boldsymbol{\gamma} + \alpha_j + \textstyle\sum_{S: j \in S} \beta_S \geq 1, & j = 1, ..., N, \\
& \mathbf{d} \in \mathcal{P}(\eta), & \\
& \boldsymbol{\alpha} \in \{0,1\}^N, \ \boldsymbol{\beta} \in \{0,1\}^{|\mathcal{S}|}, \ \boldsymbol{\gamma} \in \{0,1\}^K. &
\end{aligned} \tag{5}$$

The maximum regret (5) allows direct comparison of different policies, including DAVN, partitioned booking limits, and bid prices, without having to estimate the demand distribution and arrival processes, at a moderate computational and forecasting cost.

The next lemma presents an alternate formulation of the maximum regret problem.

Lemma 1. *Problem (5) is equivalent to the following bilevel linear optimization problem (BLP):*

$$\begin{aligned}
\rho(\mathbf{y}) = \max_{\mathbf{z},\mathbf{d},\mathbf{x}} \quad & \mathbf{r}'\mathbf{z} - \mathbf{r}'\mathbf{x}, \\
s.t. \quad & \mathbf{A}\mathbf{z} \leq \mathbf{c}, \\
& \mathbf{0} \leq \mathbf{z} \leq \mathbf{d}, \\
& \mathbf{d} \in \mathcal{P}(\eta), \\
& \mathbf{x} \in \arg\max_{\mathbf{x}} \mathbf{1}'\mathbf{x}, \\
s.t. \quad & \mathbf{A}\mathbf{x} \leq \mathbf{c}, \\
& \textstyle\sum_{j\in S} x_j \leq y_S, \qquad S \in \mathcal{S}, \\
& \mathbf{0} \leq \mathbf{x} \leq \mathbf{d}.
\end{aligned}$$

Proof. Let $(\boldsymbol{\gamma}, \boldsymbol{\beta}, \boldsymbol{\alpha})$ be boolean variables associated with the constraints of the lower level problem, equal to 1 if the respective constraint is tight, and equal to 0 otherwise. By Theorem 4.1 in Hansen et al. (1992), every optimal solution of the lower level problem satisfies $\mathbf{A}_j'\boldsymbol{\gamma} + \alpha_j + \sum_{S:j\in S} \beta_S \geq 1,\ j = 1, ..., N$. Conversely, consider an optimal solution to (5), and suppose that $\mathbf{x}$ is not a maximum flow; then some x_j can be increased, a contradiction because at least one of the constraints involving x_j is tight. □

Incidentally, the lower level problem is a maximum flow problem and is therefore always feasible, for any upper level solution $(\mathbf{z}, \mathbf{d})$. Consequently, there exists an optimal solution to the maximum regret problem that is an extreme point of the following polyhedron (see Hansen et al. 1992):

$$\left\{ (\mathbf{z}, \mathbf{d}, \mathbf{x}) : \mathbf{A}\mathbf{z} \leq \mathbf{c}, \mathbf{0} \leq \mathbf{z} \leq \mathbf{d}, \mathbf{d} \in \mathcal{P}(\eta), \mathbf{A}\mathbf{x} \leq \mathbf{c}, \sum_{S:j\in S} x_j \leq y_S\ \forall S \in \mathcal{S}, \mathbf{0} \leq \mathbf{x} \leq \mathbf{d} \right\}.$$

In particular, the worst-case demand vector $\mathbf{d}$ is an extreme point of the polyhedron $\mathcal{P}(\eta)$. For instance, under interval uncertainty, the worst-case demand is either at its lower bound or at its upper bound.

The next proposition characterizes the complexity of Problem (4): Not only is the inner problem NP-hard, but the outer problem is also non convex. As a result, it will be critical to introduce some simplifications, or develop good approximations to make the method practical.

Proposition 1.

(a) Evaluating $\rho(\mathbf{y})$ is strongly NP-hard.

(b) $\rho(\mathbf{y})$ may not be quasiconvex.

Proof. (a) From Lemma 1, the evaluation of $\rho(\mathbf{y})$ involves the solution of a BLP, which has been shown to be strongly NP-hard (Hansen et al. 1992).

(b) Consider the following three-class single-leg problem with partitioned booking limits (i.e., $\mathcal{S} = \{\{1\}, \{2\}, \{3\}\}$), with $\mathbf{r} = [3, 2, 1]$, $c = 6$, and $\mathcal{P} = \{\mathbf{d} | d_1 = 6, d_2 = 5, d_3 = 2\}$, that is, demand is deterministic. With $\mathbf{y}^1 = [6, 5, 1]$, $\rho(\mathbf{y}^1) = (6\times3) - (5\times2+1\times1) = 7$ and with $\mathbf{y}^2 = [6, 3, 3]$, $\rho(\mathbf{y}^2) = (6 \times 3) - (1 \times 3 + 3 \times 2 + 2 \times 1) = 7$. In contrast, with $\mathbf{y} = (\mathbf{y}^1 + \mathbf{y}^2)/2 = [6, 4, 2]$, $\rho(\mathbf{y}) = (6 \times 3) - (4 \times 2 + 2 \times 1) = 8 > \max\{\rho(\mathbf{y}^1), \rho(\mathbf{y}^2)\}$. Therefore, the function is not quasiconvex.

□

Structure of the Optimal Booking Policy. We now proceed to characterizing the structure of the optimal booking limit policy under interval uncertainty. Proposition 2 considers substitute products, i.e., products sharing common resources, while Proposition 3 considers complementary products.

Proposition 2. *If $l_j \leq d_j \leq u_j$, with $l_j < u_j$, $\forall j$, a nested booking policy on classes i and j, i.e., $x_j \leq y_j$ and $x_i + x_j \leq y_i + y_j$, is optimal only if for all $\mathcal{K} \subseteq \{1, ..., N\}$ such that $\mathbf{A}_i = \mathbf{A}_j + \sum_{k\in\mathcal{K}} \mathbf{A}_k$, $r_i \geq r_j + \sum_{k\in\mathcal{K}} r_k$. Otherwise, if $\min\{\mathbf{A}_i, \mathbf{A}_j\} \neq \mathbf{0}$, a partitioned booking limit policy on classes i and j, i.e., $x_i \leq y_i$ and $x_j \leq y_j$, is optimal.*

Proof. We first show that the nesting strategy is optimal if the condition holds. Consider sending a unit of flow through a subnetwork consisting of three nodes, with $\mathbf{A}_i = [1, 1]'$, $\mathbf{A}_j = [1, 0]'$ and $\mathbf{A}_k = [0, 1]'$, where product k is obtained from aggregating all products in $\mathcal{K}$. Accordingly, $r_i \geq r_j + r_k$. We study the performance of a partitioned booking limit policy on the first resource, that is, $x_j \leq y_j$ and $x_i \leq y_i$, with no restriction on product k. Because there is only one unit of capacity available, $x_i + x_j \leq 1$. Under the partitioned booking limit strategy, the maximum regret is attained at one of the following:

$$\begin{aligned}\rho_1 &= r_i(1-y_i),\\ \rho_2 &= (r_j+r_k)(1-y_j),\\ \rho_3 &= r_i - r_j y_j - r_i \min\{1-y_j, y_i\},\end{aligned}$$

where the first regret is incurred when there is only demand for i, the second when there is only demand for j and k, and the third when there is demand for products i and j, with that for i arriving before that for j. Clearly, it is optimal to set $y_i + y_j \geq 1$, that is, the sum of booking limits is larger than the available capacity, and $\rho_3 = (r_i - r_j)y_j$. The first regret is minimized when $y_i = 1$ and the second and third when they are set equal, i.e., when $y_j = (r_j + r_k)/(r_i + r_k) < 1$. In other words, the optimal booking policy must be nested. Clearly, for more general networks, i.e., nonunit capacity and constraining demand upper bounds, the *structure* of the optimal policy remain the same, while the actual *values* of the optimal booking limits may be different.

We now show that partitioned booking limits are optimal whenever the condition does not hold, that is, when either there exists some $\mathcal{K}$ such that $\mathbf{A}_i = \mathbf{A}_j + \sum_{k\in\mathcal{K}} \mathbf{A}_k$ and $r_i \leq r_j + \sum_{k\in\mathcal{K}} r_k$, or more generally, when $\mathbf{A}_i \not\subseteq \mathbf{A}_j$. To see this, consider sending a unit of flow through a four-node subnetwork, with $\mathbf{A}_i = [0,1,1]'$, $\mathbf{A}_j = [1,1,0]'$, $\mathbf{A}_k = [0,0,1]'$, and $\mathbf{A}_l = [1,0,0]'$. Without loss of generality, let the revenues satisfy $r_j + r_k > r_i + r_l > r_j$, where r_l may or may not be equal to zero. When $r_l = 0$, we can aggregate products i and l to get $\mathbf{A}_i = \mathbf{A}_j + \mathbf{A}_k$ and $r_i < r_j + r_k$; when $r_l > 0$, then $\mathbf{A}_i \not\subseteq \mathbf{A}_j$ and $\mathbf{A}_j \not\subseteq \mathbf{A}_i$ (therefore, if $r_i + r_l > r_j + r_k$, swapping indices would leave the structure of the network unchanged). We study the performance of a partitioned booking limit policy on the second resource, that is, $x_j \leq y_j$ and $x_i \leq y_i$, without restrictions on products k and l. Because there is only one unit of capacity available, $x_i + x_j \leq 1$. Under the partitioned booking limit strategy, the maximum regret is attained at one of the following:

$$\begin{aligned}\rho_1 &= (r_i + r_l)(1-y_i),\\ \rho_2 &= (r_j+r_k)(1-y_j),\\ \rho_3 &= (r_i + r_l) - r_j y_j - (r_i + r_l) \min\{1-y_j, y_i\},\\ \rho_4 &= (r_j + r_k) - r_i y_i - (r_j + r_k) \min\{1-y_i, y_j\},\end{aligned}$$

where the first regret is incurred when there is only demand for i and l, the second when there is only demand for j and k, the third when there is demand for products i, j, and l, with that for i arriving before that of j, and the last when there is demand for products i, j, and k, with that for j arriving before that of i. Clearly, it is optimal to set $y_i + y_j \geq 1$, that is, the sum of booking limits is larger than the available capacity. The first and fourth regrets are functions of y_i only, respectively decreasing and increasing with y_i; accordingly, the maximum regret is minimized when both functions are set equal, that is, when $y_i = (r_i + r_l)/(r_j + r_k + r_l) < 1$. Similarly, the second and third regrets are functions of y_j only, respectively decreasing and increasing with y_j; the maximum regret is therefore minimized when they are set equal, that is, when $y_j = (r_j + r_k)/(r_i + r_k + r_l) < 1$. Therefore, an optimal booking policy must be partitioned. Clearly, the structure of the optimal booking policy remains unchanged in more general networks. □

Therefore, nesting booking limits may be optimal, but it needs to be used parsimoniously. In particular, nesting is optimal among all products with the same OD pair but different fares, proving the robustness of the policy proposed by Curry (1990) and Chi (1995) and corroborating the result by Lan et al. (2006) for single-leg networks. Accordingly, the next section will be devoted to approximating the minimax booking limits under this policy.

However, too much nesting may hurt more than help. In fact, nesting classes gives flexibility not only to the decision-maker, but also to the "clairvoyant adversary" who maximizes the regret by choosing the worst demand scenario and sequence of arrivals. If the condition of Proposition 2 is not met, the adversary will always first release a large demand for low-revenue products, in order to consume capacity, before releasing a large demand for high-revenue products, which would then have to be rejected by lack of available capacity. In this case, it is optimal to restrict the flexibility of both the decision-maker and the adversary.

The next proposition develops a necessary optimality condition on the booking limits of complementary products. To the best of our knowledge, this is the first structural result exploiting complementarity in a network.

Proposition 3. *Suppose that $\mathcal{K}$ is a set of complementary products, that is $\min_{i\in\mathcal{K}}\{\mathbf{A}_i\} = \mathbf{0}$. Moreover, suppose that there exists a product k such that $\sum_{i\in\mathcal{K}} \mathbf{A}_i = \mathbf{A}_k$ with $r_k > r_i$, $\forall i \in \mathcal{K}$, and $\sum_{i\in\mathcal{K}} r_i > r_k$. Then, if $l_j \leq d_j \leq u_j$, $\forall j$, it is optimal to have $x_i \leq y_i + x_j$ for every $i, j \in \mathcal{K}$.*

Proof. Without loss of generality, consider sending a unit of flow through the three-node subnetwork, with $\mathbf{A}_i = [1,0]'$, $\mathbf{A}_j = [0,1]'$, and $\mathbf{A}_k = [1,1]'$. Because of the existence of product k, it is optimal to impose a booking limit on each product, by Proposition 2, that is, $x_i \le y_i$ and $x_j \le y_j$, where $y_i < 1$ and $y_j < 1$. When there is demand for only products i and j, the regret equals

$$\rho = (r_i + r_j) - r_i y_i - r_j y_j.$$

Suppose that the booking limits are modified as follows: $x_i \le y_i + x_j$ and $x_j \le y_j + x_i$. Then the regret under this demand scenario is lowered to:

$$\rho = \max\left\{(r_i + r_j) - r_i y_i - r_j \min\{1, y_j + y_i\}, (r_i + r_j) - r_j y_j - r_i \min\{1, y_j + y_i\}\right\},$$

while the regrets under the other demand scenarios (e.g., with positive demand for product k), remain unchanged. Hence, the modified booking limits, exploiting the complementarity between products, are optimal. □

Although the condition under which Proposition 3 holds may seem restrictive, it applies in many situations. In hotel RM for instance, customers can book per night or per stay, and discounts may be offered for longer stays.

As a consequence of Proposition 2, neither partitioned booking limits nor bid prices are in general optimal for a network problem. In fact, partitioned booking limits are always dominated by booking limits nested by OD pairs, by Proposition 2, and bid prices are always dominated by partitioned booking limits, since any bid-price policy $\mathbf{p}$ can be expressed as a partitioned booking limit policy, by setting $y_j = u_j$ whenever $r_j > \mathbf{p}'\mathbf{A}_j$ and $y_j = 0$ otherwise, for all $j = 1, ..., N$. [1]

4.2 Nesting Fare Classes by Origin-Destination Pairs

In this section, we characterize the minimax regret policy when booking limits are set only on products with the same OD pair. Formally, $\{i, j\} \subseteq S \in \mathcal{S}$ only if $\mathbf{A}_i = \mathbf{A}_j$. In

[1] However, if the partitioned booking limits are required to satisfy the capacity constraints, i.e., $\mathbf{Ay} \le \mathbf{c}$, neither policy dominates. To see this, consider the following three-class single-leg problem, with a unit capacity, $\mathbf{l} = \mathbf{0}$ and $\mathbf{u} = \mathbf{1}$. When $\mathbf{r} = [3, 2, 1]'$, the optimal bid-price policy, setting $p = 3/2$, leads to a regret of 1 while the optimal partitioned booking limit policy, setting $y_1 = 3/5$, $y_2 = 2/5$ and $y_3 = 0$, yields a regret of 6/5. On the other hand, when $\mathbf{r} = [3, 2, 3/2]'$, the optimal bid-price policy, setting $p = 3/2$, leads to a regret of 3/2 while the optimal partitioned booking limit policy, setting $y_1 = 5/9$, $y_2 = 1/3$ and $y_3 = 1/9$, yields a regret of 4/3.

addition, we assume that the booking limits on the same products are nested, i.e., if $i \in S, S' \in \mathcal{S}$, then either $S' \subset S$ or $S \subset S'$. This booking limit policy is similar to the policies proposed by Curry (1990) and Chi (1995). In fact, it generalizes the classical nesting policy $S = \cup_{j=1}^{N}\{j, j+1, ..., N\}$ in a single-leg problem, as well as partitioned booking limit policy $S = \cup_{j=1}^{N}\{j\}$ in a general network problem.

Moreover, we assume that the booking limits satisfy the capacity constraints, i.e., $\mathbf{Ay} \leq \mathbf{c}$. While this simplification may be suboptimal in general (e.g., consider a single leg two-class problem with partitioned booking limits: if $r_1 = r_2$, it is optimal to set $y_1 = y_2 = c$), it can be made without loss of optimality for a the single-leg problem with nested booking limits. Moreover, most existing booking limit controls (such as those derived from the DLP) satisfy this assumption. Under this assumption, and under interval uncertainty, i.e., $\mathcal{P} = \{\mathbf{d} : \mathbf{l} \leq \mathbf{d} \leq \mathbf{u}\}$, we demonstrate that Problem (3) simplifies to an LP with an exponential number of constraints.

In the sequel, we focus on maximal subsets from $\mathcal{S}$, i.e., sets $S \in \mathcal{S}$ such that there exists no other $S' \in \mathcal{S}$ with $S \subset S'$. In particular, let $\mathcal{S}^{\max}$ be the set of maximal subsets. Without loss of generality, we assume that products using the same resources are ordered by fare, and that maximal subsets comprise adjacent indices only, that is, if $\{i, j\} \subseteq S \in \mathcal{S}^{\max}$, with $i < j$, then $\mathbf{A}_i = \mathbf{A}_j$, $r_i < r_j$, and $\{i, i+1, ..., j-1, j\} \subseteq S$.

Lemma 2. *Suppose that $\mathcal{P} = \{\mathbf{d} : \mathbf{l} \leq \mathbf{d} \leq \mathbf{u}\}$. Then, the maximum regret on set $S \in \mathcal{S}^{\max}$ from following policy $\mathbf{y}$ instead of policy $\mathbf{z}$ equals*

$$\rho_S(\mathbf{y}, \mathbf{z}) = \sum_{j \leq \tau : j \in S, z_j < l_j} r_j(z_j - \min\{l_j, y_j\}) + \sum_{j > \tau : j \in S} r_j(z_j - y_j), \tag{6}$$

where $\tau = \min\{j < t : j \in S, y_j < l_j\}$ and $t = \min\{j \in S : z_j > y_j\}$. If $z_j \leq y_j$ for all $j \in S$, we set $t = N + 1$. If $t = \min\{j \in S\}$, or if $y_j \geq l_j$ for all $j < t$, $j \in S$, set $\tau = t - 1$.

Proof. For all $j \geq t$, $j \in S$, the worst-case demand is equal to $\max\{z_j, y_j\}$. To see this, suppose that $d_j < z_j$ for some j with $z_j \geq y_j$; the regret can be increased by $r_j(z_j - d_j)$ by increasing d_j, a contradiction. Suppose now that $d_j < y_j$ for some j with $y_j > z_j$. In particular, let $j^* = \min\{j \geq t, j \in S : d_j < y_j, y_j > z_j\}$ and let $k^* = \max\{j < j^*, j \in S : z_j > y_j\}$; such a k^* exists by definition of t. In this case, $(y_{j^*} - d_{j^*})$ units of capacity are available to satisfy the demand of product k^*; increasing the demand d_{j^*} by $\Delta = \min\{y_{j^*} - d_{j^*}, z_{k^*} - y_{k^*}\}$ increases the regret by $(r_{k^*} - r_{j^*})\Delta > 0$, a contradiction. As a result, for all $j \geq t$, $j \in S$, the regret equals $r_j(z_j - y_j)$.

For all $j \in S$, $\tau < j < t$, the worst-case demand equals $\max\{y_j, l_j\}$ To see this, suppose that $d_j < y_j$ for some j with $y_j > l_j$. In particular, let $j^* = \min\{j > \tau, j \in S : d_j < y_j, y_j > l_j\}$ and let $k^* = \max\{j < j^*, j \in S : l_j > y_j\}$; such a k^* exists by definition of τ. In this case, $(y_{j^*} - d_{j^*})$ units of capacity are available to satisfy the demand of product k^*; increasing the demand d_{j^*} by $\Delta = \min\{y_{j^*} - d_{j^*}, l_{k^*} - y_{k^*}\}$ increases the regret by $(r_{k^*} - r_{j^*})\Delta > 0$, a contradiction. As a result, for all $\tau < j < t$, $j \in S$, the regret equals $r_j(z_j - y_j)$.

Finally, for all $j \leq \tau$, the worst-case demand is equal to $\max\{z_j, l_j\}$. Because $y_j \geq z_j$, the regret is therefore equal to zero when $z_j \geq l_j$, and to $r_j(z_j - \min\{y_j, l_j\})$ otherwise. □

Hence, in the worst case, the demand pattern is such that capacity allocated to a particular set of products S is never used for other products. That is, if $S \subset S'$ and there exists some product $j \in S' \setminus S$ for which $z_j \geq y_j$, then the booking limit on S is always reached in the worst case.

The next proposition demonstrates that, when the demand uncertainty set is defined by bound constraints, and when the booking policy nests only products with the same OD pair, the minimax regret problem can be formulated as an LP.

Proposition 4. *When fare classes are nested by OD pair, when* $\mathcal{P} = \{\mathbf{d} : \mathbf{l} \leq \mathbf{d} \leq \mathbf{u}\}$*, and when* $\mathbf{Ay} \leq \mathbf{c}$*, the minimax regret problem (3) can be formulated as the following LP:*

$$\begin{aligned}
\min_{\rho,\mathbf{y},\mathbf{w}} \quad & \rho, \\
s.t. \quad & \rho \geq R(\{t_S\}_{S\in\mathcal{S}^{\max}}) - \textstyle\sum_{j=1}^{N} r_j w_j(\{t_S\}_{S\in\mathcal{S}^{\max}}), && \forall \{t_S\}_{S\in\mathcal{S}^{\max}} \\
& w_j(\{t_S\}_{S\in\mathcal{S}^{\max}}) \leq y_j, && \forall j, \forall \{t_S\}_{S\in\mathcal{S}^{\max}} \\
& w_j(\{t_S\}_{S\in\mathcal{S}^{\max}}) \leq l_j, && \forall j < t_S : \{j, t_S\} \subseteq S \in \mathcal{S}^{\max}, \forall \{t_S\}_{S\in\mathcal{S}^{\max}} \\
& \mathbf{Ay} \leq \mathbf{c}, \\
& \mathbf{0} \leq \mathbf{y} \leq \mathbf{u},
\end{aligned} \tag{7}$$

where $\{t_S\}_{S\in\mathcal{S}^{\max}}$ *is a set of products, with only one product per set* $S \in \mathcal{S}^{\max}$*, such that* $t_S \in S \cup \{N+1\}$ *for every* $S \in \mathcal{S}^{\max}$ *and where*

$$\begin{aligned}
R(\{t_S\}_{S\in\mathcal{S}^{\max}}) = \max_{\mathbf{z}} \quad & \mathbf{r}'\mathbf{z}, \\
s.t. \quad & \mathbf{Az} \leq \mathbf{c}, \\
& \mathbf{0} \leq \mathbf{z} \leq \mathbf{u}, \\
& z_j \leq l_j \qquad \forall j < t_S : \{j, t_S\} \subseteq S, S \in \mathcal{S}^{\max}.
\end{aligned} \tag{8}$$

Proof. The maximum regret equals

$$\rho(\mathbf{y}) = \max_{\mathbf{z}} \sum_{S\in\mathcal{S}^{\max}} \rho_S(\mathbf{y},\mathbf{z}),$$
$$\mathbf{A}\mathbf{z} \le \mathbf{c},$$
$$\mathbf{0} \le \mathbf{z} \le \mathbf{u},$$

where, by Lemma 2, the maximum regret for each $S \in \mathcal{S}^{\max}$ equals

$$\rho_S(\mathbf{y},\mathbf{z}) = \max_{t\in S\cup\{N+1\}} \left\{ \sum_{j<t:j\in S, z_j<l_j} r_j(z_j - \min\{l_j, y_j\}) + \sum_{j\ge t:j\in S} r_j(z_j - y_j) \right\}.$$

Alternatively, one can invert the order of maximization, by first maximizing with respect to the indices t, for each set $S \in \mathcal{S}^{\max}$, denoted by $\{t_S\}_{S\in\mathcal{S}^{\max}}$, and then maximizing with respect to the policy $\mathbf{z}$, that is

$$\rho(\mathbf{y}) = \max_{\{t_S\}_{S\in\mathcal{S}^{\max}}} \left\{ R(\{t_S\}_{S\in\mathcal{S}^{\max}}) - \sum_{S\in\mathcal{S}^{\max}} \sum_{j<t_S:j\in S} r_j \min\{y_j, l_j\} - \sum_{S\in\mathcal{S}^{\max}} \sum_{j\ge t_S:j\in S} r_j y_j \right\},$$

where $R(\{t_S\}_{S\in\mathcal{S}^{\max}})$ is defined by (8).

The minimax regret problem can then be formulated as

$$\min_{\mathbf{y}} \max_{\{t_S\}_{S\in\mathcal{S}^{\max}}} \left\{ R(\{t_S\}_{S\in\mathcal{S}^{\max}}) - \sum_{S\in\mathcal{S}^{\max}} \sum_{j<t_S:j\in S} r_j \min\{y_j, l_j\} - \sum_{S\in\mathcal{S}^{\max}} \sum_{j\ge t_S:j\in S} r_j y_j \right\},$$
$$s.t. \quad \mathbf{A}\mathbf{y} \le \mathbf{c},$$
$$\mathbf{0} \le \mathbf{y} \le \mathbf{u},$$

and can therefore be transformed into an LP. □

When demand is deterministic, $l_j = u_j = \mu_j$ for all $j = 1, ..., N$, the regret ρ equals a constant, equal to the optimal solution of the DLP, from which $\sum_{j=1}^{N} r_j y_j$ is subtracted. Equivalently, the problem consists in maximizing $\mathbf{r}'\mathbf{y}$, subject to the capacity and the upper bound constraints. That is, Problem (7) reduces to the DLP.

Problem (7) has an exponential number of constraints, corresponding to the number of possible choices for $\{t_S\}_{S\in\mathcal{S}^{\max}}$. Specifically, because only one index per set S is selected, and that there is a choice among $|S|+1$ possible indices (i.e., all products from S as well as $\{N+1\}$), the optimization problem has $(|S_1|+1)\times(|S_2|+1)\times ... \times(|S_{|\mathcal{S}^{\max}|}|+1)$, constraints. In particular, when classes are partitioned, there are 2^N constraints.

The exponential number of constraints is not surprising in light of Proposition 1 (a). Nevertheless, the explicit formulation of the minimax regret problem as an LP gives rise to a sequence of lower bounds on the minimax regret, by considering only a subset of the constraints in (7). In addition, there is a case where the size of the optimization problem remains tractable, which we examine next.

4.2.1 Single-Leg Problem with Nested Booking Limits

When there is only one leg, and that all classes are nested together, i.e., $S = \cup_{j=1}^{N}\{j, j+1, ..., N\}$, there are at most $N+1$ constraints in (7). Moreover, Subproblem (8) has an explicit solution.

Corollary 1. *When $\mathcal{P} = \{\mathbf{d} : \mathbf{l} \leq \mathbf{d} \leq \mathbf{u}\}$, the minimax regret problem (3) on a single leg can be formulated as the following LP:*

$$\begin{array}{lll} \min & \rho & \\ s.t. & \rho \geq \sum_{j=1}^{t-1} r_j w_j + g_t - \sum_{j=t}^{N} r_j y_j & t = 1, ..., N', \\ & \sum_{j=1}^{N} y_j \leq c, & \\ & w_j \geq l_j - y_j & j = 1, ..., N, \\ & y_j \leq u_j & j = 1, ..., N \\ & w_j \geq 0, y_j \geq 0 & j = 1, ..., N. \end{array} \tag{9}$$

where N' is the smallest integer $t \leq N$ such that $c \leq \sum_{j=1}^{t} l_j$, or equal to N otherwise and where

$$g_t = \sum_{j=t}^{N} r_j \Big(\min\{u_j, c - \sum_{i=1}^{t-1} l_i - \sum_{i=t}^{j-1} u_i\} \Big)^+.$$

Proof. When $t_S = t$, the optimal solution to (8) is defined recursively as follows:

$$z_j = \min\{c - \sum_{i<j} z_i, l_j\} \quad j < t,$$
$$z_j = \min\{c - \sum_{i<j} z_i, u_j\} \quad j \geq t.$$

Therefore, when $t \leq N'$, $R(t) = \sum_{j<t} r_j l_j + g_t$. Substituting the optimal value of $R(t)$ into (7) completes the proof. $\square$

In fact, Problem (9) has an explicit solution, as shown in the next proposition. Lan et al. (2006) independently proved the same result, using a competitive analysis argument.

Proposition 5. *When* $\mathcal{P} = \{\mathbf{d} : \mathbf{l} \le \mathbf{d} \le \mathbf{u}\}$*, the minimax regret nested booking limits, for the single-leg RM problem are equal to the following:*

$$\begin{array}{l} y_j = \left(\min\{c - \sum_{i<j} y_i, \frac{1}{r_j}(g_j - g_{j+1})\}\right)^+, \quad \text{if } j < N, \\ y_N = c - \sum_{i<N} y_i \end{array}$$

Proof. Let $\mathbf{y}^*$ be an optimal solution of (9), and let $t = \arg\max\{j : y_j > 0\}$. Accordingly, $w_j = 0$ for all $j < t$ (because $y_j \ge l_j$). With $\mathbf{y}^*$, the $t-1$ constraints (in which the terms y_j, $j < t$, only appear) are tight, together with the capacity constraint. Solving this system of t equations with t unknowns, one obtains that $y_j^* = (g_j - g_{j+1})/r_j$ and $y_t^* = c - \sum_{j<t} y_j^*$. □

Full Uncertainty. In the case of complete uncertainty, i.e., $l_j = 0$ and $u_j \ge c$, the minimax regret booking limits simplify to

$$\begin{array}{l} y_j = \min\{c(1 - \frac{r_{j+1}}{r_j}), (c - \sum_{i<j} y_i)^+\} \quad \text{if } j < N, \\ y_N = c - \sum_{i<N} y_i. \end{array}$$

Therefore, if the total capacity is ample and if the spread between fares is small, every class but the last one is allocated some fraction $(1 - r_{j+1}/r_j)$ of the total capacity. The proportionality of the capacity allocation to the ratio of fares is similar to Littlewood's formula (Littlewood 1972) or EMSR-a (Belobaba 1987).

4.2.2 Minimax Randomized Regret for Partitioned Booking Limits

As mentioned earlier, Problem (7) has an exponential number of constraints, equal to 2^N, when classes are partitioned. In this section, we propose an approximation procedure, different from considering a subset of constraints in (7), by randomizing the perfect hindsight solution $\mathbf{z}$. Specifically, instead of formulating the problem of choosing the booking limits $\mathbf{z}$ as a MIP, we assume that the booking limits $\mathbf{z}$ can be randomized and the capacity constraint must hold only in expectation, that is $\mathbf{A}E[\mathbf{z}] \le \mathbf{c}$. Under this assumption, the minimax regret network RM problem (3) is relaxed to

$$\min_{\mathbf{y}:\mathbf{A}\mathbf{y}\le\mathbf{c},\mathbf{y}\ge\mathbf{0}} \left\{ \max_{\mathbf{Z}\in\mathcal{Z}} \max_{\mathbf{D}\in\mathcal{D}} E_{\mathbf{Z}}[R(\mathbf{Z},\mathbf{D}) - R(\mathbf{y},\mathbf{D})] \right\}, \tag{10}$$

where $\mathcal{Z}$ is the set of nonnegative multivariate distributions such that $\mathbf{A}E[\mathbf{z}] \le \mathbf{c}$. We call (10) the *minimax randomized regret.* Because $\mathbf{Az} \le \mathbf{c}$ (a.s.) $\Rightarrow \mathbf{A}E[\mathbf{z}] \le \mathbf{c}$, the minimax randomized regret is an upper bound on the minimax regret (3). The next proposition shows that the minimax randomized regret problem can be efficiently solved.

Proposition 6. *The minimax randomized regret network RM problem (10) with partitioned booking limits can be formulated as the following LP:*

$$\begin{aligned}
\min_{\mathbf{p},\mathbf{q},\mathbf{y}} \quad & \mathbf{p}'\mathbf{c} + \mathbf{q}'\mathbf{1} \\
s.t. \quad & \mathbf{A}\mathbf{y} \leq \mathbf{c}, \\
& \mathbf{p}'\mathbf{A}\mathbf{U} + \mathbf{q}' \geq \mathbf{r}'(\mathbf{U} - \mathbf{Y}), \\
& \mathbf{p}'\mathbf{A}\mathbf{L} + \mathbf{q}' \geq \mathbf{0}, \\
& \mathbf{p}'\mathbf{A}\mathbf{L} + \mathbf{q}' \geq \mathbf{r}'(\mathbf{L} - \mathbf{Y}), \\
& \mathbf{q}' \geq -\mathbf{r}'\mathbf{L}, \\
& \mathbf{q}' \geq -\mathbf{r}'\mathbf{Y}, \\
& \mathbf{y} \leq \mathbf{u}, \\
& \mathbf{p}, \mathbf{y} \geq \mathbf{0}.
\end{aligned} \tag{11}$$

where $\mathbf{U}, \mathbf{L}, \mathbf{Y}$ *are diagonal matrices in which the diagonal elements correspond to* $\mathbf{u}, \mathbf{l}$*, and* $\mathbf{y}$ *respectively.*

Proof. Let δ_j^0, δ_j^l, and δ_j^u the probabilities that the optimal booking limit z_j equals 0, l_j, and u_j respectively. Since any feasible value for z_j can be expressed as a convex combination of these three points, $\delta_j^0 + \delta_j^l + \delta_j^u = 1$, and these probabilities are between 0 and 1. With partitioned booking limits, the maximum regret associated with product j as defined in Lemma 2 simplifies to

$$\rho(z_j, y_j) = \begin{cases} r_j(z_j - y_j) & \text{if } y_j < l_j, \\ r_j\left\{z_j - l_j, (z_j - y_j)^+\right\} & \text{if } y_j \geq l_j, \end{cases}$$

which is piecewise increasing. The maximum randomized regret therefore corresponds to the concave envelope of this function, which is also piecewise increasing with at most three breakpoints at zero, l_j, and u_j. (The regret when $z_j = y_j$ can always be replicated, or dominated, by randomizing z_j.) In particular, the regret equals $r_j(u_j - y_j)$ when $z_j = u_j$, $r_j \max\{0, l_j - y_j\}$ when $z_j = l_j$, and $-r_j \min\{y_j, l_j\}$ when $z_j = 0$. Therefore, the maximum randomized regret problem can be formulated as follows:

$$\max \quad \mathbf{r}'(\mathbf{U} - \mathbf{Y})\boldsymbol{\delta}^u + \mathbf{r}' \max\{\mathbf{0}, \mathbf{L} - \mathbf{Y}\}\boldsymbol{\delta}^l - \mathbf{r}' \min\{\mathbf{L}, \mathbf{Y}\}\boldsymbol{\delta}^0,$$

$$
\begin{aligned}
s.t. \qquad & \boldsymbol{\delta}^u + \boldsymbol{\delta}^l + \boldsymbol{\delta}^0 = \mathbf{1}, \\
& \mathbf{A}(\mathbf{U}\boldsymbol{\delta}^u + \mathbf{L}\boldsymbol{\delta}^l) \leq \mathbf{c}, \\
& \boldsymbol{\delta}^u, \boldsymbol{\delta}^l, \boldsymbol{\delta}^0 \geq \mathbf{0}.
\end{aligned}
$$

Let $\mathbf{p}$ and $\mathbf{q}$ be the dual variables respectively associated with the capacity constraints and the probability normalization constraints. By strong duality, the above problem is equivalent to its dual, which is a minimization problem. Plugging this inner problem into the general minimax regret problem completes the proof. □

Despite the regret randomization, Problem (11) simplifies to the DLP when demand is deterministic, as Problem (7).

Problem (11) has $(K + 2N)$ variables and $(K + 6N)$ constraints, which is a considerable improvement from (7). In comparison, the DLP has N variables and $K + N$ constraints. The larger size of the problem is the price to pay to capture demand stochasticity.

Bid Prices Based on Robust Booking Limits. The variables $\mathbf{p}$ in (11) are the dual variables associated with the constraint $\mathbf{A}E[\mathbf{z}] \leq \mathbf{c}$. In fact, the optimal value of p_k measures the additional revenue that could be obtained if, in addition to knowing the demand distributions, the malevolent adversary were also given an additional unit of capacity c_k. Therefore, the optimal value of $\mathbf{p}$ can be used as a proxy for the marginal value of capacity. Although they are obtained from an LP, the variables $\mathbf{p}$ capture the stochastic nature of the demand, in contrast to the dual values of DLP.

5. Maximin Revenue

In this section, we first formulate the maximin problem (2) and then characterize the structure of an optimal policy.

5.1 Problem Formulation

We first formulate the inner minimization problem in (2) as a MIP, for given open-loop booking limits $\mathbf{y}$, and characterize its complexity. Using the same notations as in §4.1, the minimum revenue problem can be formulated as follows:

$$
\varphi(\mathbf{y}) = \min_{\mathbf{x},\mathbf{d},\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma}} \qquad \mathbf{r}'\mathbf{x},
$$

$$
\begin{aligned}
s.t. \quad & \mathbf{Ax} \le \mathbf{c}, & \\
& \textstyle\sum_{j \in S} x_j \le y_S, & S \in \mathcal{S}, \\
& \mathbf{0} \le \mathbf{x} \le \mathbf{d}, & \\
& \mathbf{d} \le \mathbf{x} + M(\mathbf{1} - \boldsymbol{\alpha}), & \\
& \textstyle\sum_{j \in S} x_j \ge \beta_S y_S, & S \in \mathcal{S}, \\
& \mathbf{a}'_k \mathbf{x} \ge c_k \gamma_k, & k = 1, ..., K, \\
& \mathbf{A}'_j \boldsymbol{\gamma} + \alpha_j + \textstyle\sum_{S: j \in S} \beta_S \ge 1, & j = 1, ..., N, \\
& \mathbf{d} \in \mathcal{P}(\eta), & \\
& \boldsymbol{\alpha} \in \{0,1\}^N, \ \boldsymbol{\beta} \in \{0,1\}^{|\mathcal{S}|}, \ \boldsymbol{\gamma} \in \{0,1\}^K. &
\end{aligned} \tag{12}
$$

The next proposition characterizes the complexity of the maximin revenue problem. The proof is similar to that of Proposition 1 and is therefore omitted.

Proposition 7.

(a) Evaluating $\varphi(\mathbf{y})$ *is NP-hard.*

(b) $\varphi(\mathbf{y})$ *may not be quasiconcave.*

5.2 Optimal Solution Characterization

We now proceed to characterizing the structure of an optimal solution to (12).

Interval Uncertainty. Interestingly, even with interval uncertainty, the revenue is not necessarily minimized when $d_j = l_j$ for all j. [2] Nevertheless, when booking limits are chosen optimally, the worst-case demand is always equal to its lower bound. It is indeed optimal to set the partitioned booking limits $\mathbf{y}$ so as to maximize $\mathbf{r}'\mathbf{y}$ subject to $\mathbf{Ay} \le \mathbf{c}$ and $\mathbf{0} \le \mathbf{y} \le \mathbf{l}$ (i.e., a modified DLP). Increasing y_j above l_j would never pay off as the adversary can always choose $d_j = l_j$; on the other hand, the unused capacity $y_j - l_j$ could be used to serve the (deterministic) demand for another product, potentially leading to an increase in revenue. Moreover, because only the deterministic component of demand will materialize, nesting fare classes is never be profitable.

[2]For instance, consider a two-class problem, with $l_1 = 1$, $l_2 = 0$, and $c = 1$. If $y_1 = 0$ and $y_2 = 1$, the worst-case demand scenario is $d_1 = 1$, $d_2 = 1$, leading to a profit of r_2. In contrast, with $\mathbf{d} = \mathbf{l}$, the profit would have been equal to $r_1 > r_2$.

As a result, the maximin solution is pessimistic, as it anticipates a demand no larger than its lower bound. To reduce the level of conservatism of the maximin approach, one can restrict the total amount of variability to a "budget of uncertainty." Indeed, it is unlikely that all demands are significantly different from their nominal value. Let μ_j be the nominal value of demand d_j, and let σ_j be the maximum deviation of d_j from μ_j. That is, $l_j = \mu_j - \sigma_j$ and $u_j = \mu_j + \sigma_j$. Bertsimas and Sim (2004) suggested to limit the sum of relative deviations from the nominal values to some budget Γ, i.e., $\sum_{j=1}^{N} |(d_j - \mu_j)/\sigma_j| \leq \Gamma$. The parameter Γ measures the amount of uncertainty captured by the model. If $\Gamma = 0$, there is no uncertainty and the problem is deterministic (in which case all demands d_j are fixed to their nominal values μ_j). At the other extreme, if $\Gamma = N$, there is complete uncertainty on demand, and the worst case occurs when $d_j = \mu_j - \sigma_j$, $\forall j$. Intermediate values of Γ specify a moderate level of conservativeness, while accounting for some uncertainty. For simplicity, we assume that Γ is an integer.

Interval Uncertainty with a Budget Constraint. We now show that the maximin problem with partitioned booking limits and under a budget of uncertainty can be formulated as an LP with a polynomial size (in the problem data). Similarly to §4.2, we assume that the booking limits satisfy the capacity constraints, i.e., $\mathbf{Ay} \leq \mathbf{c}$. While this may not be optimal in general (e.g., consider the single-leg two-class problem with unit capacity: under the demand constraint $d_1 + d_2 \geq 1$, it is optimal to set $y_1 = 1 = y_2$), this assumption is commonly made with existing methods. It is also optimal when either $\Gamma = 0$ or $\Gamma = N$ because then, the booking limits also play the role of protection levels (as in DLP).

Because revenue is minimized only with negative deviations from the nominal values when $\mathbf{Ay} \leq \mathbf{c}$ (so that capacity reserved for high-revenue products cannot be consumed by cheap products), it is never optimal to set a booking limit above the nominal demand value, and the maximin problem with uncertainty budget can be simplified to:

$$\max_{\mathbf{y}:\mathbf{Ay}\leq\mathbf{c},\mathbf{0}\leq\mathbf{y}\leq\boldsymbol{\mu}} \; \min_{\boldsymbol{\delta}:\mathbf{0}\leq\boldsymbol{\delta}\leq\mathbf{1},\mathbf{1}'\boldsymbol{\delta}\leq\Gamma} \sum_{j=1}^{N} r_j \min\{\mu_j - \sigma_j\delta_j, y_j\}. \tag{13}$$

Proposition 8. *If* $\mathbf{Ay} \leq \mathbf{c}$*, the maximin revenue problem with partitioned booking limits, with a budget of uncertainty* Γ *on the sum of relative deviations from the nominal values, (13), can be formulated as the following LP:*

$$\max \qquad \alpha\Gamma + \boldsymbol{\beta}'\mathbf{1} + \mathbf{r}'\mathbf{y}$$

$$
\begin{aligned}
s.t. \quad & \mathbf{A}\mathbf{y} \leq \mathbf{c}, \\
& \mathbf{y} \leq \boldsymbol{\mu}, & (14) \\
& r_j y_j + \alpha + \beta_j \leq r_j(\mu_j - \sigma_j), \quad \forall j, \\
& \alpha \leq 0, \boldsymbol{\beta} \leq 0, \mathbf{y} \geq \mathbf{0}. & (15)
\end{aligned}
$$

Proof. The minimization problem in (13) consists in minimizing a concave function (minimum of linear functions) over a polyhedron. Therefore, an optimal solution is an extreme point of the polyhedron. Since Γ is integer, every extreme point of the polyhedron $\{\boldsymbol{\delta} : \mathbf{0} \leq \boldsymbol{\delta} \leq \mathbf{1}, \mathbf{1}'\boldsymbol{\delta} \leq \Gamma\}$ different from zero has Γ coordinates equal to 1 and $N - \Gamma$ coordinates equal to zero. Because δ_j can only take two values in an optimal solution, the inner problem is equivalent to:

$$
\begin{aligned}
\min \quad & \mathbf{r}'\mathbf{y} + \textstyle\sum_{j=1}^{N} r_j \min\{0, \mu_j - \sigma_j - y_j\}\delta_j, \\
s.t. \quad & \mathbf{1}'\boldsymbol{\delta} \leq \Gamma, \\
& \mathbf{0} \leq \boldsymbol{\delta} \leq \mathbf{1}.
\end{aligned}
$$

Let α and $\boldsymbol{\beta}$ be the dual variables associated with the budget constraint and the upper bound constraints respectively. By strong duality, the above problem is equivalent to its dual, which is a maximization problem. Plugging this inner maximization problem into the outer maximin problem proves the result. □

There is unfortunately no equivalent linear formulation for nested booking limits, as all $\binom{N}{\Gamma}$ demand scenarios must be enumerated to find the minimum revenue. Indeed, the decision-maker and the malevolent adversary face different types of capacity constraints (in product units for the decision-maker and in half-intervals σ_i for the adversary), which misaligns their optimal strategies, giving rise to multiple local optima.

6. Numerical Examples

6.1 Single-Leg Example

We first consider the single-leg four-class example in pages 48-50 in Talluri and van Ryzin (2004b). Each class j is characterized by a fare r_j, a mean demand μ_j, and a standard deviation σ_j (see Table 1).

Table 1: Problem data and nested booking limits for the single-leg example of Talluri and van Ryzin (2004b)

		Demand Statistics		Nested Booking Limits			
j	r_j	μ_j	σ_j	Minimax Regret	Maximin	EMSR-a	EMSR-b
1	1,050	17.3	5.8	119	119	119	119
2	567	45.1	15.0	102	108	102	102
3	534	39.5	13.2	69	77	80	68
4	520	34.0	11.3	37	51	63	35
Minimum Revenue				55,043	55,043	55,043	55,043
Maximum Regret				4,695	7,838	5,039	5,356

Uncertainty Set. We assume that only the range of the demand is known, taken as $[\mu_j - \sigma_j, \mu_j + \sigma_j]$. Assuming a normal demand distribution, this interval covers 68% of the possible demand realizations, for each fare class. Figure 1 demonstrates that the performance of the minimax regret and the maximin nested booking limits (the details of the simulation are explained below) are not too sensitive (i.e., within a few percents) to the amount of uncertainty captured by the model. Notice the concave shape of the functions, characteristic of a trade-off between flexibility and conservativeness. Observe also that the minimax regret is less sensitive to the interval size than the maximin, demonstrating its better ability to deal with uncertainty.

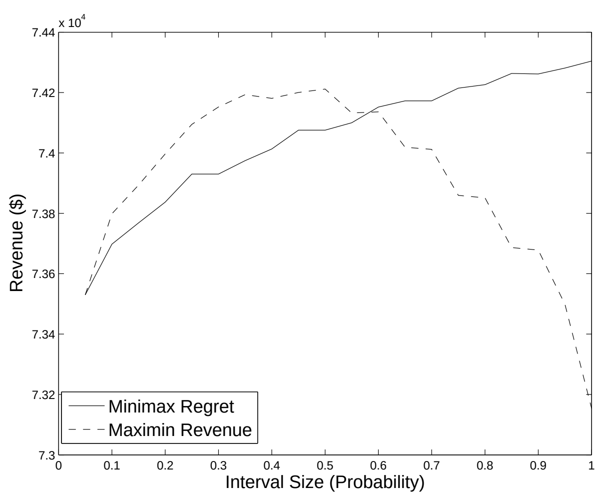

Figure 1: Simulated expected revenue for different interval sizes, using the minimax regret and maximin nested booking limits in the single-leg example by Talluri and van Ryzin (2004b) when $c = 130$.

Policy Comparison. Table 1 displays the nested protection levels obtained with the minimax regret, the maximin, EMSR-a (Belobaba 987), and EMSR-b (Belobaba 1992) when there are 119 seats in the aircraft. Normal demand distributions were assumed to compute the EMSR booking limits. The maximin criterion is pessimistic and protects little capacity for high-fare classes. Note that nesting classes is not necessary from a maximin revenue standpoint (see §5.2). The minimax regret is in contrast less conservative and leads to booking limits comparable to the EMSR heuristics. The two last rows of Table 1 compares the minimum revenue and the maximum regret guarantees with the different policies, by solving (5) and (12) respectively. From the minimum revenue objective, all policies are alike. In fact, the performance guarantee is extremely low, especially when compared with the optimal value of the DLP, equal to 73,723 (which is an upper bound on the expected revenue). The ranges of the maximum regrets, on the other hand, are extremely narrow, indicating that all policies are expected to perform well in expectation.

Performance Simulation. To measure the robustness of our policies, we generate 1,000 demand scenarios, and simulate the airline booking process under the proposed policies, starting 150 days before departure. Similarly to de Boer et al. (2002) and Bertsimas and de Boer (2005), we model the arrival process as a non-homogeneous Poisson process. The arrival intensity for ODF j in period t satisfies the following relationship:

$$\lambda_j(t) = B_j(t)G_j,$$

where $B_j(t)$ follows a standardized beta distribution, and G_j follows a gamma distribution. The choice of the beta distribution for modeling the arrival pattern is motivated by its flexible shape (e.g., increasing, decreasing, unimodal, bimodal) while the factor G_j creates some correlation between the demands for product j across all booking periods. Under this demand model, the total number of booking requests follows a negative binomial distribution.

In our simulations, the shape parameters of the beta distribution were such that high-fare demand almost certainly arrives after low-fare demand (specifically, the shape parameters were set equal to $\mathbf{a} = [10, 5, 2, 2, 1]$ and $\mathbf{b} = [1, 2, 2, 5, 10]$). On the other hand, the parameters of the gamma distribution were derived from the mean and standard deviation presented in Table 1 (specifically, the shape parameter was set to $\mu_j^2/(\sigma_j^2 - \mu_j)$ and the scale parameter was set to $\sigma_j^2/\mu_j - 1$).

Table 2 displays the 95% confidence intervals for the mean revenues generated with the maximin, the minimax regret, EMSR-a and EMSR-b, when capacity is varied from 80 to 150, creating demand factors (DF) from 1.7 to 0.9, under standard nesting.

Table 2: 95% confidence intervals for the mean revenues for the single-leg example of Talluri and van Ryzin (2004) with $\mathbf{r} = [1050, 567, 534, 520]$

c	DF	Regret	Maximin	EMSR-a	EMSR-b
80	1.70	49,999±230	49,098±109	49,806±189	49,890±247
90	1.51	55,193±230	54,308±111	55,063±191	55,011±261
100	1.36	60,382±232	59,538±122	60,326±199	60,209±207
110	1.24	65,453±249	64,688±164	65,458±227	65,329±269
120	1.13	70,164±310	69,581±249	70,230±294	70,129±313
130	1.05	74,214±409	73,869±362	74,350±393	74,296±398
140	0.97	77,292±522	77,286±485	77,607±506	77,584±507
150	0.91	79,351±624	79,699±601	79,867±614	79,867±614

In general, the robust approaches perform almost equally well than the EMSR heuristics. Their good performance is remarkable despite the fact that they focus on the worst cases. The maximin criterion nevertheless tends to underperform the other approaches, but the optimality gap is surprisingly small given its the level of conservatism. One should nevertheless point out that the performance of the maximin booking limits depends critically on how they are used: While the maximin revenue is not affected by whether the booking limits are nested or partitioned, the simulated revenue would have been significantly lower with partitioned booking limits.

Observe also that the minimax regret approach tends to underperform the EMSR heuristics when the demand factor decreases. Intuitively, the worst-case demand scenarios foreseen by the minimax regret (see §4.2.1) are characterized by a high demand load factor (in particular, there is one worst-case demand scenario for which all demands are equal to their upper bound), which explains its superior performance under large demand factors.

Evenly-Spaced Fares. Similar observations hold when the fares are more evenly spaced, as shown in Table 3. The dominance of the EMSR heuristics over the robust approaches, is however stronger in this case, as well as the dominance of the minimax regret over the maximin criterion. In fact, with the original fares, all policies concord on protecting a certain number of seats for class 1, because the opportunity cost of missing a sale is high, while there is more confusion when fares are more spread out.

Table 3: 95% confidence intervals for the mean revenues for the single-leg example of Talluri and van Ryzin (2004) with $\mathbf{r} = [1050, 950, 699, 520]$

c	DF	Regret	Maximin	EMSR-a	EMSR-b
80	1.70	89,405±541	88,389±382	89,695±469	89,859±543
90	1.51	73,650±355	70,245±146	74,592±383	74,588±411
100	1.36	79,263±392	76,074±181	80,013±412	80,298±506
110	1.24	84,696±436	82,517±261	85,227±415	85,500±506
120	1.13	90,269±531	89,011±376	90,429±448	90,649±519
130	1.05	94,970±678	94,836±524	95,426±550	95,578±584
140	0.97	98,544±804	99,469±686	99,722±695	99,772±708
150	0.91	101,580±902	102,720±841	102,810±844	102,860±848

Demand Correlation. We now investigate the impact of correlation among demands. In particular, we assume that the distributions of the demand rates G_j are perfectly correlated gamma distributions. As shown in Table 4, the minimax regret approach tends to outperform the EMSR heuristics in this case. In effect, the worst-case demand scenarios, against which the booking limits are protected, typically assume a large degree of (positive or negative) correlation among demands (see §4.2.1).

Table 4: 95% confidence intervals for the mean revenues for the single-leg example of Talluri and van Ryzin (2004) with $\mathbf{r} = [1050, 950, 699, 520]$ with perfectly correlated demand rates

c	DF	Regret	Maximin	EMSR-a	EMSR-b
80	1.70	66,975±425	64,406±221	66,725±394	66,965±430
90	1.51	72,472±485	69,519±251	72,172±543	72,532±559
100	1.36	77,812±534	74,622±334	77,294±576	77,259±674
110	1.24	82,996±594	79,900±458	82,420±587	82,448±676
120	1.13	87,712±725	84,958±606	86,804±665	87,301±713
130	1.05	91,559±897	89,491±769	90,379±795	91,008±823
140	0.97	94,817±1,050	93,436±941	93,807±952	94,130±966
150	0.91	97,865±1,193	96,683±1,083	96,854±1,114	97,002±1,121

Censored Sales Data. The robust approaches have clearly the largest appeal when only limited information about demand is available. So far, we have assumed complete information about the demand distributions; however, in practice, only sales data is measured. We now investigate the performance of the robust approaches with demand censoring.

We assume that the medians $\hat{\boldsymbol{\mu}}$ and the standard deviations $\hat{\boldsymbol{\sigma}}$ of the demand are estimated from past sales data. The estimates are naive, in the sense that they are not adjusted for the missing demand data, yet the median estimates tend to be less sensitive to outlier data than alternative estimation methods. We simulate the booking process with EMSR-b booking limits until the sales data converge (within 5%, with the Euclidean norm) to the estimates that are used to compute the booking limits. Once the estimates are consistent with the sales data, we measure the performance of the minimax regret and the maximin with the estimated range $[(\hat{\boldsymbol{\mu}}-\hat{\boldsymbol{\sigma}})^+, (\hat{\boldsymbol{\mu}}+\hat{\boldsymbol{\sigma}})]$, as well as the EMSR-a and EMSR-b heuristics, assuming normal demand with mean $\hat{\boldsymbol{\mu}}$ and standard deviation $\hat{\boldsymbol{\sigma}}$.

As shown in Table 5, the minimax regret approach tends to outperform the EMSR heuristics when demand data is censored and the demand factor is high. The performance of the maximin booking limits is on the other hand strongly dependent on accurate range estimation, as they are based on one worst-case scenario (i.e., all demands equal to their lower bounds), in contrast to the minimax regret, which balances several worst cases.

Table 5: 95% confidence intervals for the mean revenues for the single-leg example of Talluri and van Ryzin (2004) with $\mathbf{r} = [1050, 950, 699, 520]$ with censored sales observations

c	DF	Regret	Maximin	EMSR-a	EMSR-b
80	1.70	65,927±207	58,724±81	62,220±124	63,039±131
90	1.51	74,557±381	69,893±131	73,557±299	74,002±381
100	1.36	80,338±513	75,151±165	78,577±286	79,204±333
110	1.24	85,551±532	81,740±255	83,919±306	84,457±344
120	1.13	90,619±589	88,699±374	89,784±397	90,158±424
130	1.05	95,273±676	94,792±523	95,210±534	95,382±549
140	0.97	98,997±794	99,465±686	99,644±690	99,713±696
150	0.91	101,580±902	102,720±841	102,810±844	102,860±849

6.2 Small Network Example

We next consider the network example of de Boer et al. (2002), which consists of four cities arranged in series. Flights are assumed to go only in one direction; accordingly, there are six possible OD pairs. Each OD pair has three fare classes, giving rise to 18 different products. There are 150 booking periods, and demand is assumed to follow a non-homogeneous Poisson process. The fares, means, standard deviations, and shape parameters of the booking arrival process of all products appear in Tables 8 and 9 of de Boer et al. (2002). In addition to the base case, we consider a situation with larger variances (Table 10 in de Boer et al.) and smaller fare spread (Table 11 in de Boer et al.). Each aircraft has 200 seats.

As before, we assume that the only information available about the demand distributions is their ranges, defined as $[(\boldsymbol{\mu} - \boldsymbol{\sigma})^+, \boldsymbol{\mu} + \boldsymbol{\sigma}]$ where $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ are the median and the standard deviation of the demand distributions. We benchmark the performance of the robust approaches against the the following methods, which are commonly used in practice:

- DAVN EMSR/DLP: DAVN booking limits, were classes are clustered in at most 10 buckets, according to the algorithm described in appendix of Bertsimas and de Boer (2005), and where the booking limits on each resource are respectively computed with EMSR-b (Belobaba 1992) and the DLP.

- Nested BL DLP: booking limits where fare classes are nested by OD pair (Curry 1990), and the booking limits are obtained from the DLP.

- Bid DLP/RLP: bid prices, respectively set equal to the shadow prices of the DLP, and the average shadow prices of 100 randomized DLPs, where the demand samples are independently generated from a negative binomial distribution.

The robust controls (DAVN-ρ/φ, Nested BL-ρ/φ,Bid-ρ/φ) are defined similarly, but are computed by minimizing the maximum regret (5) and maximizing the minimum revenue (12).

The MIPs were solved with the branch-and-bound algorithm of CPLEX 10, and the outer optimization problem was solved with the Sequential Quadratic Programming (SQP) method developed in Matlab 2007. Note that, because of lack of convexity (Propositions 1 and 7), the SQP method is only guaranteed to converge to a local optimum.

Table 6 presents the maximum regret $\rho(\mathbf{y})$ and the minimum revenue $\varphi(\mathbf{y})$ for the different methods, in the base case investigated by de Boer et al. (2002). We also report the 95% confidence intervals for the mean revenues over 1,000 simulations of the booking process.

Performance of the Minimax Regret. In general, the maximum regret is lower when fare classes are nested by OD pair, consistently with Proposition 2. In contrast, the mean simulated revenue is typically larger with DAVN booking limits than when fare classes are nested by OD pair. Therefore, the most robust policy has not necessarily the same *structure* as the best policy on average. Within a class of policies, however, the minimax regret performs extremely well, and is sometimes superior to the traditional approaches.

Table 6: Maximum regret, minimum revenue and 95% confidence intervals for the mean simulated revenues, for the network example of de Boer et al. (2002).

Control	Base Case			Lower Fare Spread	Inflated Variance
	$\rho(\mathbf{y})$	$\varphi(\mathbf{y})$	$E[R(\mathbf{y})]$	$E[R(\mathbf{y})]$	$E[R(\mathbf{y})]$
DAVN-ρ	19,995	36,205	76,429±469	63,232±293	74,142±467
DAVN-φ	52,650	42,150	60,098±174	57,200±90	60,747±225
DAVN EMSR	23,875	36,130	77,1297±468	63,802±300	76,231±491
DAVN DLP	22,740	37,450	76,423±431	63,245±318	75,580±460
Nested BL-ρ	16,140	43,350	73,647±402	61,522±205	72,668±413
Nested BL-φ	32,270	44,870	74,042±349	57,660±89	62,444±237
Nested BL DLP	20,945	41,350	73,115±398	61,363±283	72,218±433
Bid-ρ/φ	39,080	44,870	73,545±297	64,072±254	72,913±314
Bid DLP	39,080	44,870	73,545±297	63,334±192	72,545±333
Bid RLP	32,695	41,190	76,513±381	64,498±250	75,854±402

From these observations, we conclude that RM models are more sensitive to the sequence of arrivals (which significantly affects the structure of the minimax regret policy) than to the demand variability (which affects the values of the booking limits, for a given policy). That is, RM is more sensitive to the *combinatorial* nature of the sequence of arrivals than the *stochastic* nature of demand.

Only single-leg problems reduce to stochastic optimization problems, because their combinatorial structure is completely absorbed through the use of nested booking limit controls. (There however remains the unsolved question on how to use these controls, using standard or theft nesting, see Bertsimas and de Boer 2005.) In contrast, network RM is mostly a combinatorial problem, as demonstrate the pathological sequences of arrivals characterizing the maximum regret, as well as the superiority of the DAVN approaches even when the stochastic nature of demand is ignored in the computation of the booking limits.

The example proposed by de Boer et al. (2002) not only assumes a strict low-before-high fare arrivals (which is defendable given the advance-purchase restrictions) but also impose the *same arrival pattern for all products*, i.e., that the *mix of itineraries remains constant over time*. This latter assumption, which may or may not be true in practice, strikingly contrasts with the worst-case sequences of arrivals considered by the minimax regret. Clearly, under more extreme sequence of arrivals, the performance of the minimax regret (and in particular, of the Nested BL-ρ) would dramatically improve over more traditional approaches.

Performance of the Maximin. The maximin criterion often performs poorly. One shall however mention that the maximin controls are generally not uniquely defined, and that not all perform as poorly. For instance, in the single-leg case, partitioned and nested booking limits lead to the same maximin revenue but have very different average-case performances. For general network problems, any control that covers the deterministic portion of demand **l** is guaranteed the maximin revenue. One would then need to define a secondary criterion for selecting the best policy, from an average-case perspective, among those that achieve the maximin criterion.

Correlation. We now investigate the impact of correlation by assuming perfect correlation among the demands rates G_j. As shown in Table 7, the performance gap between the DAVN and the nested booking limit policies disappears when demands are correlated. In fact, the situation with perfectly correlated demand rates is closer to the worst-case scenarios considered by the minimax regret approach, highlighting the superiority of nested booking limits (Proposition 1). As in the case with independent demands, the minimax regret performs very well, within a class of policies while the maximin controls tend to be conservative (but can be optimized with a secondary objective).

Table 7: 95% confidence intervals for the mean simulated revenues, for the network example of de Boer et al. (2002) with perfectly correlated demand rates.

Control	Base Case	Lower Fare Spread	Inflated Variance
DAVN-ρ	71,723±1,084	60,973±646	70,857±1,003
DAVN-φ	58,855±196	56,518±164	58,493±340
DAVN EMSR	71,839±1,023	60,894±687	70,751±1,078
DAVN DLP	71,796±952	60,878±703	70,573±1,030
Nested BL-ρ	71,793±927	60,813±475	71,002±938
Nested BL-φ	69,739±713	56,888±181	59,701±352
Nested BL DLP	71,342±947	60,459±685	70,047±1,1017
Bid-ρ/φ	69,817±595	60,979±931	66,459±741
Bid DLP	68,633±616	60,424±472	66,223±708
Bid RLP	71,678±817	60,979±634	66,459±741

Censored Data. We now evaluate the performance of the robust approaches in the presence of censored sales observations. With censored observations, the demand parameters are estimated from the sales data directly. To obtain the demand estimates, we simulate the booking process, using the DAVN DLP booking limits, and estimate the mean, median, and standard deviation from the sales data, without adjustment for the missign data; we then iterate this procedure until the demand estimates used for deriving the DAVN DLP booking limits are within 5% (in Euclidean norm) of the demand estimates obtained from the sales data. After convergence, we simulate the booking process using those demand estimates, to compare the robust approaches with the more traditional controls. Table 8 shows that, similarly to the single-leg example, the minimax regret tends to dominate the DLP-based controls in the presence of censored observations. Interestingly, the DAVN EMSR booking limits are performing extremely well, despite the fact that both the mean and the standard deviation of the demand are ill estimated. In comparison, the bid prices obtained from the RLP have the poorest performance.

Table 8: 95% confidence intervals for the mean simulated revenues, for the network example of de Boer et al. (2002) with censored sales observations.

Control	Base Case	Lower Fare Spread	Inflated Variance
DAVN-ρ	69,205±278	56,864±147	69,702±299
DAVN-φ	66,576±145	56,634±99	68,234±191
DAVN EMSR	70,007±252	58,289±191	69,810±277
DAVN DLP	66,729±194	55,883±139	66,884±209
Nested BL-ρ	68,649±242	57,596±123	68,670±250
Nested BL-φ	67,313±134	58,357±68	69,467±179
Nested BL DLP	64,616±210	55,186±136	64,491±214
Bid-ρ/φ	60,019±180	55,553±108	60,132±238
Bid DLP	59,910±187	55,688±113	60,429±245
Bid RLP	21,156±479	14,302±310	21,158±448

Minimax Randomized Regret. Computing the minimax regret booking limits becomes quickly intractable as the network size grows, since evaluating the maximum regret involves solving a MIP (5). It is therefore important to develop good approximation schemes to compute the minimax regret booking limits efficiently. In §4.2.2, we showed that, with partitioned booking limits, evaluating the maximum regret reduces to solving a simple LP (6) when the regret is randomized. Because of the randomization, the optimal value of (6) is an upper bound on the maximum regret obtained with the same partitioned booking limits. The upper bound is in fact very tight: in the three scenarios considered, we never observed a gap between the upper bound and the actual maximum regret larger than 0.03%. However, because partitioned booking limits are in general not optimal, the maximum regret can be lowered by considering different booking limit policies; for instance, in the base case, nested booking limits lead to a maximum regret of 16,140 (see Table 6) while partitioned booking limits give rise to a maximum regret of 16,562.

Problem (6) gives rise to booking limits, which can be subsequently used in a nested fashion, either using the DAVN method or by nesting fare classes with the same OD pair, as well as bid prices. Table 9 displays the simulated mean revenues obtained with these controls (where the tilde refers to the randomization of the regret). Comparing these revenues with those obtained with the minimax (deterministic) regret (see Table 6) reveals that nearly the same level of performance can be attained with the randomized regret, at a much lower computational cost.

Table 9: 95% confidence intervals for the mean simulated revenues with the booking limits obtained from the minimax randomized regret problem (6), for the network example of de Boer et al. (2002).

Control	Base Case	Inflated Variance	Lower Fare Spread
DAVN-$\tilde{\rho}$	77,065±384	62,152±160	76,306±374
Nested BL-$\tilde{\rho}$	72,989±398	61,308±215	72,496±411
Bid-$\tilde{\rho}$	73,497±293	64,723±264	72,818±324

6.3 Large Network Examples

We now investigate the viability of the robust approach for solving large-scale network RM problems. Because of the large computational cost involved with the MIP formulation of the maximum regret (5), we only use the minimax randomized regret (6). The minimax randomized regret provides an upper bound on the maximum regret for partitioned booking limits, and can be used heuristically to choose nested booking limits and bid prices.

The first problem is the hub-and-spoke problem reported in Table 5.3. in Williamson (1988). Four cities are connected with a hub. Considering all possible origin-destination pairs, there is a total of 20 itineraries on 8 legs. In addition, there are four classes per itinerary. Each class on a given itinerary is associated with a mean demand, a standard deviation, and a fare. We use the same nonhomogenous Poisson model as before, assuming the same beta-distributed arrival pattern for all itineraries with the same fare class and selecting the shape and scale parameters of the gamma distribution of the demand rate to match the given means and standard deviations. We consider five variants of this problem, by taking the aircraft capacity equal to 100, 125, 150, 175, and 200 seats, corresponding to demand factors varying from 125% to 55.4%.

Table 10 reports the mean simulated revenues with the proposed policies. (For the sake of clarity, we omit the confidence intervals here, but the range is about 600.) As before, the minimax regret approach is at least comparable to more traditional approaches (except the bid price controls which behave unevenly). Interestingly, the DAVN EMSR booking limits perform poorly in this example, in comparison to the DAVN DLP and minimax regret booking limits.

Table 10: Mean simulated revenues, for the network example of Williamson (1988).

Control	Capacity				
	100	125	150	175	200
DAVN-$\tilde{\rho}$	131,440	146,580	156,570	165,460	176,360
DAVN-φ	104,010	112,980	113,160	113,350	114,210
DAVN EMSR	133,940	147,830	155,900	160,070	160,240
DAVN DLP	132,750	147,470	156,570	167,510	177,280
Nested BL-$\tilde{\rho}$	127,900	141,690	152,860	162,890	171,710
Nested BL-φ	101,700	111,520	112,640	112,630	112,640
Nested BL DLP	127,610	141,110	152,470	162,540	171,640
Bid-$\tilde{\rho}$	129,380	141,780	127,690	149,730	172,620
Bid DLP	127,300	111,200	138,470	147,420	172,620
Bid RLP	127,310	144,090	152,560	143,980	172,620

The second problem is a real airline network problem, with 517 different itineraries, each with 11 fare classes, and 67 segments. Arrivals follow a non-homogenous Poisson process, with 16 changes of rates. The aggregate DF is 64.5%. Table 11 reports the mean revenues under 100 simulation runs.

Summary. We now provide a brief summary of our main observations:

- The structure of the most robust policies may not be the same as the structure of the policies that perform the best, on average. However, within a certain booking policy, the robust approaches perform as well as, and sometimes even better than, the more traditional approaches, especially in the presence of correlation or censored data. The sensitivity of the minimax regret to pathological sequences of arrivals highlights the criticality of addressing the combinatorial nature of network RM models (i.e., all

Table 11: 95% confidence intervals for the mean simulated revenues, for the real-world hub-and-spoke network problem.

Control	Revenue
DAVN-$\tilde{\rho}$	3,549,400±9,623
DAVN-φ	2,997,900±5,323
DAVN EMSR	3,612,000±9,369
DAVN DLP	3,612,000±9,043
Nested BL-$\tilde{\rho}$	3,413,100±8,893
Nested BL-φ	2,932,200±4,378
Nested BL DLP	3,392,100±7,472
Bid-$\tilde{\rho}$	3,554,700±9,622
Bid DLP	3,614,700±9,117
Bid RLP	3,601,500±9,336

possible sequences of arrivals), as well as its stochastic nature (i.e., all possible demand realizations).

- The performance of the robust controls is rather insensitive to the size of the uncertainty set, especially for the minimax regret which equally weighs the extremes.
- Maximin booking limits may ne be uniquely defined. Given the large spread in performance among the maximin solutions, there is a need for another decision criterion, to be optimized over all maximin solutions.
- Computing the minimax regret booking limits is in general not efficient, as it involves the solution of a MIP. However, randomizing the minimax regret gives rise to efficient heuristics to obtain robust controls.

7. Conclusions

In this paper, we develop robust formulations for the capacity allocation problem in RM, using the maximin and the minimax regret criteria, under polyhedral uncertainty. We consider generic booking limit controls, including, among others, partitioned booking limits, nested booking limits, DAVN, and fixed bid prices. Our models allow for a quick performance comparison of different policies, because they only involve the solution of LPs or small MIPs. In addition, when demand uncertainty is represented by intervals, we provide closed-form solutions for the minimax regret nested booking limits in a single leg, propose a sequence of upper bounds when fare classes are nested by OD pairs, develop an efficient heuristic to compute partitioned booking limits, and show that the maximin booking limits can be obtained by solving an LP. Our numerical study reveals that the robust policies generally perform as well as, and sometimes even better than, most traditional approaches.

Robust approaches are pragmatic. They are scalable to solve practical network RM problems, because they combine efficient solution methods (closed-form solutions, LPs, or small MIPs) with modest data requirements. They are also flexible, because they do not require anything about the uncertainty set, as long as it is polyhedral. Accordingly, various levels of information about demand can be incorporated into the uncertainty set representation, (such as range, moments, shape, correlation, demand censorship) at no (modeling or computational) cost, in contrast to traditional probabilistic models. Moreover, a probabilistic guarantee can be derived for any type of uncertainty set and any level of information, by solving a moment bound problem (Bertsimas and Popescu 2005 and Popescu 2005). Because of its unique combination of flexibility and scalability, the robust optimization paradigm has a tremendous potential for improving operations and increasing revenues, especially in fast-changing business environments.

References

Adelman, D. (2006). Dynamic bid-prices in revenue management. *Oper. Res.* To Appear.

Akan, M. and Ata, B. (2006). On bid price controls for network revenue management. Working Paper, Northwestern University.

Ball, M. and Queyranne, M. (2006). Toward robust revenue management: Competitive analysis of online booking. Robert H. Smith School Research Paper No. RHS 06-021.

Belobaba, P. P. (1987). *Air Travel Demand and Airline Seat Inventory Management.* PhD thesis, Flight Transportation Laboratory, MIT, Cambridge, MA.

Belobaba, P. P. (1992). Optimal vs. heuristic methods for nested seat allocation. In *Proc. AGIFORS Reservations and Yield Management Study Group*, Brussels, Belgium.

Ben-Tal, A., Golany, B., Nemirovski, A., and Vial, J.-P. (2005). Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing & Service Oper. Management*, 7(3):248–271.

Ben-Tal, A. and Nemirovski, A. (1999). Robust solutions of uncertain linear programs. *Oper. Res. Lett.*, 25(1):1–13.

Bertsimas, D. and de Boer, S. (2005). Simulation-based booking limits for airline revenue management. *Oper. Res.*, 53(1):90–106.

Bertsimas, D. and Popescu, I. (2003). Revenue management in a dynamic network environment. *Transp. Sci.*, 37(3):257–277.

Bertsimas, D. and Popescu, I. (2005). Optimal inequalities in probability theory: A convex optimization approach. *SIAM J. of Optimization*, 15(3):780–804.

Bertsimas, D. and Sim, M. (2004). The price of robustness. *Oper. Res.*, 52(1):35–53.

Bertsimas, D. and Thiele, A. (2006). A robust optimization approach to inventory theory. *Oper. Res.*, 54(1):150–168.

Besbes, O. and Zeevi, A. (2006). Blind nonparametric revenue management: Asymptotic optimality of a joint learning and pricing method. Columbia Working Paper.

Birbil, S. I., Frenk, J. B. G., Gromicho, J. A. S., and Zhang, S. (2006). An integrated approach to single-leg airline revenue management: The role of robust optimization. Working Paper, Erasmus University Rotterdam.

Boyd, E. A. and Bilegan, I. C. (2003). Revenue management and e-commerce. *Management Sci.*, 49(10):1363–1386.

Chi, Z. (1995). *Airline Yield Management in a Dynamic Network Environment*. PhD thesis, MIT, Cambridge, MA.

Cooper, W. L. (2002). Asymptotic behavior of an allocation policy for revenue management. *Oper. Res.*, 50(4):720–727.

Curry, R. E. (1990). Optimal airline seat allcoation with fare classes nested by origins and destinations. *Transp. Sci.*, 24(3):193–204.

de Boer, S. V., Freling, R., and Piersma, N. (2002). Mathematical programming for network revenue management revisited. *Eur. J. Oper. Res.*, 137(1):72–92.

Dror, M., Trudeau, P., and Ladany, S. P. (1988). Network models for seat allocation on flights. *Transp. Res. B*, 22(4):293–250.

Erken, S. and Maglaras, C. (2006). Revenue management heuristics under limited market information: A maximum entropy approach. Presentation at the 6th Annual INFORMS Revenue Management and Pricing Section Conference, Columbia University.

Gallego, G. and Moon, I. (1993). The distribution free newsboy problem: Review and extensions. *J. Oper. Res. Soc.*, 44(8):825–834.

Gallego, G., Ryan, J. K., and Simchi-Levi, D. (2001). Minimax analysis for finite-horizon inventory models. *IIE Transactions*, 33:861–874.

Godfrey, G. and Powell, W. B. (2001). An adaptive, distribution-free algorithm for the newsvendor problem with censored demands, with application to inventory and distribution problems. *Management Sci.*, 47:1101–1112.

Hansen, P., Jaumard, B., and Savard, G. (1992). New branch-and-bound rules fro linear bilevel programming. *SIAM J. Sci. Stat. Comput.*, 13(5):1194–1217.

Huh, W. T. and Rusmevichientong, P. (2006). An asymptotic analysis of inventory planning with censored demand. Columbia Working Paper.

Lan, Y., Gao, H., Ball, M., and Karaesmen, I. (2006). Revenue management with limited demand information. University of Maryland, College Park.

Levi, R., Roundy, R., and Shmoys, D. B. (2005). Provably near-optimal sample-based policies for stochastic inventory control models. Cornell Working Paper.

Littlewood, K. (1972). Forecasting and control of passenger bookings. In *Proceedings of the Twelfth Annual AGIFORS Symposium*, Nathanya, Israel.

Perakis, G. and Roels, G. (2006). Regret in the newsvendor problem with partial information. *Oper. Res.* To Appear.

Popescu, I. (2005). A semidefinite approach to optimal-moment bounds for convex class of distributions. *Math. Oper. Res.*, 30(3):632–657.

Rusmevichientong, P., Van Roy, P., and Glynn, P. W. (2006). A non-parametric approach to multi-product pricing. *Oper. Res.*, 54(1):89–98.

Scarf, H. E. (1958). A min-max solution to an inventory problem. In Arrow, K. J., Karlin, S., and Scarf, H. E., editors, *Studies in Mathematical Theory of Inventory and Production*, pages 201–209. Stanford University Press, Stanford, CA.

Smith, B. C., Leimkuhler, J. F., and Darrow, R. M. (1992). Yield management at American Airlines. *Interfaces*, 22(1):8–31.

Talluri, K. (2007). RLP revisited. Presentation at the 7th Annual INFORMS Revenue Management and Pricing Section Conference, University of Pompeu-Fabra.

Talluri, K. I. and van Ryzin, G. J. (1998). An analysis of bid-price controls for network revenue management. *Management Sci.*, 44(11):1577–1593.

Talluri, K. I. and van Ryzin, G. J. (1999). A randomized linear programming formulation method for computing network bid prices. *Transp. Sci.*, 33(2):207–216.

Talluri, K. I. and van Ryzin, G. J. (2004a). Revenue management under a general discrete choice model of consumer behavior. *Management Sci.*, 50(1):15–33.

Talluri, K. T. and van Ryzin, G. J. (2004b). *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers, Boston, MA.

Topaloglu, H. (2006). Using Lagrangian relaxation to compute capacity-dependent bid-prices in network revenue management. Working Paper, Cornell University.

van Ryzin, G. J. and McGill, J. I. (2000). Revenue management without forecasting and optimization: An adaptive algorithm for determining seat protection levels. *Management Sci.*, 46(6):760–775.

van Ryzin, G. J. and Vulcano, G. (2005). Simulation-based optimization of virtual nesting controls for network revenue management. NYU Working Paper.

Williamson, E. L. (1988). Comparison of optimization techniques for origin-destination seat inventory comtrol. Master's thesis, MIT, Cambridge, MA.

Williamson, E. L. (1992). *Airline Network Seat Control: Methodologies and Revenue Impacts.* PhD thesis, MIT, Cambridge, MA.

Wollmer, R. D. (1986). A hub-and-spoke seat management model. Technical report, Douglas Aircraft Company, McDonnell Douglas Corporation.

Yue, J., Chen, B., and Wang, M.-C. (2006). Expected value of distribution information for the newsvendor problem. *Oper. Res.*, 6(54):1128–1136.

Zhang, D. and Cooper, W. L. (2005). Revenue management for parallel flights with customer choice behavior. *Oper. Res.*, 53(3):415–431.