Tractable Open Loop Policies for Joint Overbooking and Capacity Control over a Single Flight Leg with Multiple Fare Classes

Huseyin Topaloglu
School of Operations Research and Information Engineering,
Cornell University, Ithaca, New York 14853, USA
topaloglu@orie.cornell.edu

S. Ilker Birbil, J. B. G. Frenk, Nilay Noyan
Faculty of Engineering and Natural Sciences,
Sabanci University, Orhanli, Tuzla 34956, Istanbul, Turkey
{sibirbil, frenk, nnoyan}@sabanciuniv.edu

December 19, 2011

Abstract

In this paper, we consider the joint overbooking and capacity control problem over a single flight leg with multiple fare classes. The objective is to maximize the net expected revenue, which is given by the difference between the expected revenue from the accepted requests and the expected penalty cost from the denied reservations. We study a class of open loop policies that accept the requests for each fare class with a fixed acceptance probability. In this case, the challenge becomes that of finding a set of acceptance probabilities that maximize the net expected revenue. We derive a simple expression that can be used to compute the optimal acceptance probabilities, despite the fact that the problem of finding the optimal acceptance probabilities is a high dimensional optimization problem. We show that the optimal acceptance probabilities randomize the acceptance decisions for at most one fare class, indicating that the randomized nature of our open loop policies is not a huge practical concern. We bound the performance loss of our open loop policies when compared with the optimal policy. Computational experiments demonstrate that open loop policies perform remarkably well, providing net expected revenues within two percent of the optimal on average.

Revenue management operations of an airline deal with the problem of allocating the limited seats available on the flight legs among different fare classes. Perhaps, the two most important ingredients of these operations are overbooking and capacity control. Overbooking decides how much the physically available seats on the flight legs should be oversold given that not all reservations show up at the departure time, either due to prior cancellations or last minute no shows. On the other hand, capacity control decides which fare classes should be left open for purchase by the passengers, trying to keep a balance between not selling too many seats to low fare class passengers and making sure that the flight legs do not depart with too many empty seats. Naturally, the overbooking and capacity control ingredients have effects on each other. Overbooking determines the working capacities available on the flight legs, which influence the lowest fare class that should be left open for purchase, whereas capacity control determines the fare class mix of the passengers, which influences the penalty cost of denying boarding to a reservation and the likelihood that a reservation shows up at the departure time.

In this paper, we study a joint overbooking and capacity control model over a single flight leg. In our problem setting, the requests for different fare classes arrive according to a Poisson process. We need to decide whether to accept or reject each arriving request. An accepted request provides a revenue and becomes a reservation, whereas a rejected request leaves the system. At the departure time of the flight leg, a portion of the reservations show up. If we cannot accommodate all of these reservations on the flight leg, then we incur a penalty cost for the denied reservations. The objective is to maximize the net expected revenue, which is the difference between the expected revenue from the accepted requests and the expected penalty cost from the denied reservations. A dynamic programming model for this problem can be computationally difficult since it requires keeping track of the number of accepted requests for each fare class, resulting in a high dimensional state variable. Instead, we work with a class of open loop policies that accept a request for each fare class with a fixed acceptance probability. In this case, the goal is to find a set of acceptance probabilities that maximize the net expected revenue.

We make contributions along the dimensions of computational tractability of our model, structure of the optimal acceptance probabilities and performance of open loop policies. First, we note that there is one acceptance probability for each fare class and the problem of finding a set of acceptance probabilities that maximize the net expected revenue is a high dimensional optimization problem. We show that it is possible to come up with an essentially closed form expression for the optimal acceptance probabilities. This expression requires only a few lookups to the cumulative distribution function of the Poisson distribution. Besides computing the optimal acceptance probabilities in a tractable fashion, this closed form expression suggests that fare to show up probability ratio should be used as a metric to prioritize the fare classes when deciding which requests to accept. Second, the class of policies that we use are randomized policies as they employ probabilistic decision rules. Therefore, a possible cause for concern may be that such randomized policies, while acceptable in theory, may not be ideal for practical implementation. We alleviate this concern by showing that the optimal acceptance probabilities randomize the acceptance decisions for at most one fare class. Third, we develop a performance bound for our class of open loop policies. This performance bound allows us to identify the crucial problem parameters that affect the performance of open loop policies. Fourth, we point out several extensions of our model that can be important from a practical perspective. In particular, our initial development assumes that we incur the same penalty cost for each reservation that we deny boarding, but we make extensions to general convex penalty cost functions that allow a progressively higher penalty cost for each denied reservation. Furthermore, our model assumes that the requests for different fare classes arrive according to a Poisson processes. This arrival process implies that the coefficient of variation of the number of requests is equal to the reciprocal of the square root of the mean number of requests. For certain practical settings, such a coefficient of variation may be considered too small, especially when the mean number of requests is large. We show how to incorporate more variable demand processes into our model by using arrival rates that are themselves random variables. This approach essentially corresponds to using mixture distributions for the request arrivals. Ultimately, our model allows us to address the uncertainty in the joint overbooking and capacity control problem without sacrificing computational tractability and with reasonable data requirements.

Besides providing an overbooking and capacity control policy over a single flight leg, we feel that our model fills an important gap in the revenue management literature. Under the assumption that overbooking is not possible, there is a natural progression that spans the two fare class capacity control model of Littlewood (1972), its extension to multiple fare classes in the expected marginal seat revenue heuristic of Belobaba (1987) and the dynamic programming based capacity control model of Brumelle and McGill (1993). Littlewood (1972) studies a two fare class model where the demand for the low fare class arrives before the demand for the high fare class and derives an expression for the optimal number of seats to protect for the high fare class. Later on, Belobaba (1987) builds on this expression in a multiple fare class setting. He develops the well known expected marginal seat revenue heuristic, giving a simple expression for the amount of capacity to protect for all of the later arriving fare classes when making the capacity control decisions here and now for a particular fare class. Due to its ease of implementation, minimal data requirements and intuitive appeal, the expected marginal seat revenue heuristic has seen considerable attention from both practitioners and academics. Finally, Brumelle and McGill (1993) formulate the capacity control problem with multiple fare classes as a dynamic program and derive numerous structural properties of the optimal policy.

In contrast, one crucial piece of a similar progression is missing when we consider the case where overbooking is possible. On the one hand, there exist overbooking models with a single fare class. On the other hand, there exist overbooking models with multiple fare classes that are based on dynamic programming formulations. However, there does not appear to exist any multiple fare class overbooking models that admit a simple solution. A useful byproduct of our work is to fill this gap by building the missing bridge between single fare class models and dynamic programming models. In this respect, our model plays a role similar to the one played by the expected marginal seat revenue heuristic of Belobaba (1987) in the case where overbooking is not possible. Section 4.2.1.2 in Talluri and van Ryzin (2005) describes a single fare class overbooking model and derives an expression for the optimal number of requests to accept. Our model can be visualized as an extension of this model to multiple fare classes. After deriving the optimality conditions for our model, we carefully compare it with the one in Section 4.2.1.2 in Talluri and van Ryzin (2005). Also, our overbooking approach is different from those that compute an overbooking pad to obtain a desired overbooking risk, since our approach is driven by the penalty cost for the denied reservations. Belobaba (2006) gives a comparison of risk and cost driven approaches. Subramanian et al. (1999) study dynamic programming formulations of overbooking problems with multiple fare classes. Such dynamic programming formulations incorporate early cancellations as well as last minute no shows, though early cancellations are likely to be less crucial in practice as one can address them by reoptimization during the booking process. Dynamic programming formulations are typically not computationally tractable as they require high dimensional state variables. In addition, we need to know the dynamics of the arrivals of the requests over time to be able to solve dynamic programming formulations and it is usually difficult to pinpoint how the optimal policy depends on the problem parameters. Our model requires only forecasting the total numbers of requests for different fare classes and its optimality conditions clearly demonstrate how the optimal acceptance probabilities depend on the problem parameters. Thus, there are practical reasons that may make our model preferable to dynamic programming formulations.

To sum up, we make the following research contributions in this paper. 1) We develop a joint overbooking and capacity control model with multiple fare classes. Our model assumes that the requests for different fare classes arrive according to a Poisson process. We focus on a class of open loop policies that accept a request for each fare class with a fixed probability. 2) We derive an essentially closed form expression for the optimal acceptance probabilities. 3) We show that the optimal acceptance probabilities randomize the acceptance decisions for at most one fare class. 4) We develop a performance bound for our class of open loop policies. 5) We extend our model to incorporate general convex penalty cost functions and demand processes that display more variability than the Poisson process. 6) Our computational experiments indicate that open loop policies perform quite well. Furthermore, a heuristic modification of our open loop policies does not randomize the acceptance decisions at all and this modification suffers from minimal performance loss.

The rest of the paper is organized as follows. In Section 1, we review the relevant literature. In Section 2, we formulate the joint overbooking and capacity control problem over a single flight leg as a dynamic program. This formulation gives a precise description of the problem that we want to solve and clarifies the difficulties associated with dynamic programming formulations. In Section 3, we introduce our open loop policies and show that the optimal acceptance probabilities randomize the acceptance decisions for at most one fare class. In Section 4, we derive a simple expression that can be used to compute the optimal acceptance probabilities. In Section 5, we give a performance bound for our open loop policies. In Section 6, we extend our model by incorporating general convex penalty cost functions and by showing how to capture arrival processes with more variability than a Poisson process. In Section 7, we provide computational experiments. In Section 8, we conclude.

1 Literature Review

We begin by reviewing overbooking models that take place over a single flight leg. Beckmann (1958) studies a model with a single fare class, where the goal is to decide how many requests to accept given that not all reservations show up at the departure time and we incur a penalty cost whenever we deny boarding to a reservation. Section 4.2.1.2 in Talluri and van Ryzin (2005) analyzes a similar single fare class model and the optimality condition of their model is similar to ours. We compare the two optimality conditions later in the paper, but it is worthwhile to note that our model can be visualized as an extension of the one in Section 4.2.1.2 in Talluri and van Ryzin (2005) to deal with multiple fare classes. Thompson (1961) and Coughlan (1999) develop overbooking models with multiple fare classes, where they treat the requests for different fare classes as static random variables, ignoring the temporal dynamics of the arrivals of the requests. These models do not admit closed form solutions. There exist overbooking models over a single flight leg that use dynamic programming formulations to capture the temporal dynamics more accurately. Rothstein (1971) gives a single fare class dynamic programming formulation to carefully model the cancellations. Rothstein (1974) points out extensions to multiple fare classes. Chatwin (1998) characterizes the structure of the optimal policy in an overbooking problem with a single fare class. Chatwin (1999) studies a model where the fare and the refund are prefixed functions of the time left in the selling horizon. Subramanian et al. (1999) illustrate the computational difficulties when one uses dynamic programming formulations under fare class specific cancellation and no show probabilities. Lan et al. (2011) and Ball and Queyranne (2009) approach overbooking problems by using online algorithms with regret criterion. Aydin et al. (2010) study both static and dynamic models, provide upper and lower bounds on the optimal net expected revenue in the static case and characterize the structure of the optimal policy in the dynamic case.

As for overbooking models that take place over a network of flight legs, Gallego and van Ryzin (1997) consider the case where the primary control is the price charged for the different itineraries and give a deterministic approximation to the problem. Similarly, Bertsimas and Popescu (2003) propose a deterministic linear programming formulation, using which they derive an overbooking and capacity control policy. Karaesmen and van Ryzin (2004*a*) study network overbooking models, where they estimate the expected revenue from the accepted requests by using a deterministic approximation. Karaesmen and van Ryzin (2004*b*) focus on multiple flight legs that can serve as substitutes of each other, which is the case for parallel flight legs that are operated between the same origin destination pair on a single day. Kunnumkal and Topaloglu (2008) and Erdelyi and Topaloglu (2009) show that it is possible to approximate the dynamic programming formulation of the network overbooking problem, as long as one can develop an approximation to the penalty cost that is separable by the itineraries. Erdelyi and Topaloglu (2010) give a strategy to decompose the network overbooking problem by the flight legs, so that they can obtain accurate approximations by solving a sequence of overbooking problems, each taking place over a single flight leg. Therefore, the single leg problem that we study in this paper appears as a subproblem in their decomposition strategy.

There are papers that give performance bounds for revenue management policies by relating these policies to deterministic approximations. The performance bound that we give for our open loop policies follows this line of research. For a single product pricing problem, Gallego and van Ryzin (1994) bound the performance of the prices that are obtained from a deterministic approximation. Gallego and van Ryzin (1997) extend this work to multiple products. For network revenue management problems without overbooking, Talluri and van Ryzin (1998), Cooper (2002) and Topaloglu (2009) give performance bounds for policies that are based on deterministic linear programs. Levi and Radovanovic (2010) bound the performance of policies in a revenue management setting with reusable resources and these policies are also derived from deterministic linear programs.

2 Problem Formulation

We have a single flight leg that can be used to serve the requests for different fare classes arriving randomly over time. We need to decide whether to accept or reject each arriving request as the requests arrive. An accepted request generates a revenue that reflects its fare class and becomes a reservation, whereas a rejected request simply leaves the system. At the departure time of the flight leg, a certain portion of the reservations show up. If all of the reservations that show up cannot be accommodated on the flight leg, then we deny boarding to some of them by incurring a penalty cost. The goal is to maximize the net expected revenue, which is the difference between the expected revenue from the accepted requests and the expected penalty cost from the denied reservations.

The problem takes place over the finite selling horizon $[0, \tau]$. Time 0 corresponds to the beginning of the selling horizon and time τ corresponds to the departure time of the flight leg. There are n fare classes indexed by $1, \ldots, n$. The requests for fare class j arrive according to the nonstationary Poisson process with the intensity function $\{\Lambda_j(t) : t \in [0, \tau]\}$. We assume that the arrival processes for the requests for different fare classes are independent. If we accept a request for fare class j, then we generate a revenue of f_j. A reservation for fare class j shows up at the departure time with probability q_j. We assume that the show up decisions of different reservations are independent of each other. The available capacity on the flight leg is C. If a reservation shows up at the departure time and we cannot accommodate it on the flight leg, then we incur a penalty cost of θ for each reservation that we deny boarding. Throughout the paper, we assume that we do not give refunds to the reservations that do not show up at the departure time and the reservations are not canceled during the selling horizon, but these assumptions are only for notational brevity. It is straightforward to incorporate refunds that are specific to each fare class and we can generalize our approach to the case where a reservation stays in the system for an exponentially distributed amount of time before it is canceled, unless the flight leg departs first. Using an exponential distribution for the time until cancellation does not provide the most general cancellation model and may not be appropriate in some practical settings. However, this way of modeling the time until cancellation closely follows the standard discrete time models that assume that a reservation is canceled with a fixed probability at each time period, in which case, the time until cancellation ends up having a geometric distribution.

We use x_{jt} to denote the total number of reservations for fare class j at time t so that the vector $x_t = (x_{1t}, \ldots, x_{nt})$ captures the state of the reservations. Given that the number of reservations for fare class j at time τ is $x_{j\tau}$, we use $S_j(x_{j\tau})$ to denote the number of reservations for fare class j that show up at the departure time of the flight leg. Noting the assumption that the show up decisions of different reservations are independent, $S_j(x_{j\tau})$ has a binomial distribution with parameters $(x_{j\tau}, q_j)$. The vector $S(x_\tau) = (S_1(x_{1\tau}), \ldots, S_n(x_{n\tau}))$ captures the state of the reservations at the departure time of the flight leg and the assumption that the show up decisions of different reservations are independent also implies that the elements of the vector $S(x_\tau)$ are independent of each other.

To write the Hamilton Jacobi Bellman equation for the problem, we divide the selling horizon into time intervals of length Δ. Over an interval of length Δ around time t, we observe a request for fare class j with probability $\Delta \, \Lambda_j(t) + o(\Delta)$, where $o(\cdot)$ satisfies $\lim_{\Delta \to 0} o(\Delta)/\Delta = 0$. With the remaining probability $1 - \sum_{j=1}^n \Delta\,\Lambda_j(t) + o(\Delta)$, we do not observe a request for any fare classes. In this case, we can compute the optimal policy in discrete time by solving the optimality equation

$$V_t(x_t) = \sum_{j=1}^n \Delta\,\Lambda_j(t)\ \max\Big\{f_j + V_{t+\Delta}(x_t + e^j), V_{t+\Delta}(x_t)\Big\} + \Big[1 - \sum_{j=1}^n \Delta\,\Lambda_j(t)\Big] V_{t+\Delta}(x_t) + o(\Delta), \quad (1)$$

where $e^j \in \Re^n$ is the unit vector with a one in the element corresponding to j. The state variable in the optimality equation keeps track of the numbers of reservations for different fare classes. The two terms in the curly brackets correspond to accepting and rejecting a request for fare class j. Letting $[\cdot]^+ = \max\{\cdot, 0\}$, the boundary condition of the optimality equation is $V_\tau(x_\tau) = -\theta\,\mathbb{E}\big\{[\sum_{j=1}^n S_j(x_{j\tau}) - C]^+\big\}$, which computes the expected penalty cost of denying boarding to the reservations. Subtracting $V_{t+\Delta}(x_t)$ from both sides of the optimality equation, dividing it by Δ and taking the limit as Δ approaches to zero, we obtain the Hamilton Jacobi Bellman equation

$$\frac{\partial V_t(x_t)}{\partial t} = -\sum_{j=1}^n \Lambda_j(t)\ \max\Big\{f_j + V_t(x_t + e^j), V_t(x_t)\Big\} + \sum_{j=1}^n \Lambda_j(t)\, V_t(x_t) \quad (2)$$

with the boundary condition $V_\tau(x_\tau) = -\theta\,\mathbb{E}\big\{[\sum_{j=1}^n S_j(x_{j\tau}) - C]^+\big\}$. Our derivation of the Hamilton Jacobi Bellman equation is somewhat heuristic since we do not verify the differentiability of $V_t(x_t)$ with respect to t or justify the interchange of the order of the limit and the maximization operator, but these somewhat heuristic arguments can be made precise by appealing to the general theory in Bremaud (1980). We prefer to omit the details since we never directly work with the Hamilton Jacobi Bellman equation throughout the paper. The main goal of this equation is to prevent any ambiguities regarding the statement of the problem that we want to solve.

The Hamilton Jacobi Bellman equation characterizes the optimal policy, but solving this equation requires dealing with a high dimensional function in continuous time and knowing the exact dynamics of the requests. Furthermore, this equation does not yield clear structural results or give a feel for how the optimal policy depends on the problem parameters. In the next section, we focus on policies that are easily computable and have intuitive relationships with the problem parameters.

3 Open Loop Policy

We consider a policy that accepts a request for a particular fare class with a fixed probability. In particular, we let p_j be the probability with which we accept a request for fare class j. Since the total number of requests for fare class j is a Poisson random variable with mean $\lambda_j = \int_0^\tau \Lambda_j(t)\,dt$ and we accept each request for fare class j with probability p_j, the number of accepted requests for fare class j is a Poisson random variable with mean $\lambda_j\,p_j$. Furthermore, since each reservation for fare class j shows up with probability q_j, the number of reservations for fare class j that show up at the departure time is a Poisson random variable with mean $q_j\,\lambda_j\,p_j$. Using $\mathsf{Pois}(\alpha)$ to denote a Poisson random variable with mean α, if we use the policy characterized by the acceptance probabilities $p = (p_1, \ldots, p_n)$, then the net expected revenue that we obtain can be written as

$$\Pi(p) = \sum_{j=1}^n f_j\,\lambda_j\,p_j - \theta\,\mathbb{E}\big\{[\mathsf{Pois}(\textstyle\sum_{j=1}^n q_j\,\lambda_j\,p_j) - C]^+\big\}, \quad (3)$$

where the first term computes the expected revenue from the accepted requests and the second term computes the expected penalty cost of denying boarding to the reservations. In this case, we can compute the best set of acceptance probabilities by solving the problem

$$\max_{p \in [0,1]^n} \Pi(p). \tag{4}$$

Although the decision variable in problem (4) is a high dimensional vector, we shortly show that solving this problem turns out to be quite tractable. Furthermore, we may be concerned about the fact that the optimal solution to problem (4) potentially includes many fractional decision variables, indicating that the policy characterized by the best acceptance probabilities is highly randomized. Such randomized policies may not be desirable in practice. We alleviate this concern by showing that there exists an optimal solution to problem (4) with at most one fractional decision variable.

The net expected revenue function in (3) involves expectations of the form $\mathbb{E}\{[\mathsf{Pois}(\alpha) - C]^+\}$ for a Poisson random variable $\mathsf{Pois}(\alpha)$. In the next lemma, we show that $\mathbb{E}\{[\mathsf{Pois}(\alpha) - C]^+\}$ is a differentiable and convex function of α and give a closed form expression for the derivative of $\mathbb{E}\{[\mathsf{Pois}(\alpha) - C]^+\}$ with respect to α in terms of the cumulative distribution function of the Poisson distribution. The proof of this lemma follows from a more general result due to Aydin et al. (2010). We defer both this general result and the proof of the lemma to Appendix A.

Lemma 1 *The expectation* $\mathbb{E}\{[\mathsf{Pois}(\alpha) - C]^+\}$ *is a differentiable and convex function of* α *and the derivative of* $\mathbb{E}\{[\mathsf{Pois}(\alpha) - C]^+\}$ *with respect to* α *is given by* $\mathbb{P}\{\mathsf{Pois}(\alpha) \geq C\}$.

We can also show the same differentiability and convexity result by using general theorems on stochastic convexity from Appendix B in Talluri and van Ryzin (2005), but the expression for the derivative of $\mathbb{E}\{[\mathsf{Pois}(\alpha) - C]^+\}$ with respect to α given in Lemma 1 may not hold for other stochastically convex random variables. An implication of Lemma 1 is that the net expected revenue is a concave function of the acceptance probabilities. To see this, we note that Section 3.2.2 in Boyd and Vandenberghe (2005) shows that if $g(\cdot) : \Re \to \Re$ is convex and $h(\cdot) : \Re^n \to \Re$ is linear, then $g(h(\cdot)) : \Re^n \to \Re$ is convex. Therefore, letting $g(\alpha) = \mathbb{E}\{[\mathsf{Pois}(\alpha) - C]^+\}$ and $h(p) = \sum_{j=1}^n q_j \lambda_j p_j$, $g(\alpha)$ is a convex function of α by Lemma 1 so that $g(h(p)) = \mathbb{E}\{[\mathsf{Pois}(\sum_{j=1}^n q_j \lambda_j p_j) - C]^+\}$ is a convex function of p. Since the first term in $\Pi(p)$ is linear in p, $\Pi(p)$ is concave in p.

We can use Lemma 1 to obtain gradients of the net expected revenue function. Differentiating $\Pi(p)$ with respect to p_j and using Lemma 1 and the chain rule, we have

$$\frac{\partial \Pi(p)}{\partial p_j} = f_j \lambda_j - \theta q_j \lambda_j \mathbb{P}\{\mathsf{Pois}(\textstyle\sum_{i=1}^n q_i \lambda_i p_i) \geq C\} = q_j \lambda_j \left\{ \frac{f_j}{q_j} - \theta \mathbb{P}\{\mathsf{Pois}(\textstyle\sum_{i=1}^n q_i \lambda_i p_i) \geq C\} \right\}. \tag{5}$$

Considering the expression in the curly brackets above, the magnitude of f_j/q_j is crucial in determining the sign of the derivative $\partial \Pi(p)/\partial p_j$. Throughout the paper, we assume that the fare classes are indexed such that we have $f_1/q_1 > f_2/q_2 > \ldots > f_n/q_n$. This assumption comes with slight loss of generality as it does not allow multiple fare classes with the same fare to show up probability ratio, but this loss of generality is not a huge practical concern since the estimates of the show up probabilities are subject to noise and one should be comfortable with perturbing these estimates to obtain a strict ordering. Furthermore, noting that the net expected revenue function is continuous in the show up probabilities, perturbing the show up probabilities by small amounts changes the optimal objective value of problem (4) also by a small amount. In the next proposition, we show that there exists an optimal solution to problem (4) with at most one fractional decision variable.

Proposition 2 *There exists an optimal solution p^* to problem* (4) *such that we have $p_1^* = 1, \ldots, p_{j-1}^* = 1, p_j^* \in (0,1], p_{j+1}^* = 0, \ldots, p_n^* = 0$ for some $j = 1, \ldots, n$.*

Proof An alternative statement of the proposition is that there exists an optimal solution p^* to problem (4) such that if $p_j^* \in [0,1)$ for some $j = 1, \ldots, n$, then we have $p_k^* = 0$ for all $k > j$. To obtain a contradiction, we assume that there exists $j = 1, \ldots, n$ and $k > j$ such that $p_j^* \in [0,1)$ and $p_k^* \in (0,1]$. Since p^* is an optimal solution to problem (4), $p_j^* \in [0,1)$ must maximize the net expected revenue when we fix all of the acceptance probabilities at p^* except the one for fare class j and optimize only over the acceptance probability for fare class j. That is, $p_j^* \in [0,1)$ must be an optimal solution to the problem $\max_{p_j \in [0,1]} \Pi(e^j (p_j - p_j^*) + p^*)$. Since the optimal solution to the last optimization problem is either the left point of the interval $[0,1]$ or in the interior, the derivative of the objective function at the optimal solution must be less than or equal to zero. In other words, we have $d\Pi(e^j (p_j - p_j^*) + p^*)/dp_j \,\big|_{p_j = p_j^*} \leq 0$, which implies that

$$\frac{d\Pi(e^j (p_j - p_j^*) + p^*)}{dp_j}\bigg|_{p_j = p_j^*} = \frac{\partial \Pi(p)}{\partial p_j}\bigg|_{p = p^*} = q_j\, \lambda_j \left\{ \frac{f_j}{q_j} - \theta\, \mathbb{P}\big\{ \mathsf{Pois}(\textstyle\sum_{i=1}^n q_i\, \lambda_i\, p_i^*) \geq C \big\} \right\} \leq 0,$$

where the second equality follows by (5). Therefore, we have $f_j/q_j \leq \theta\, \mathbb{P}\big\{\mathsf{Pois}(\sum_{i=1}^n q_i\, \lambda_i\, p_i^*) \geq C\big\}$. By following a similar argument, we observe that $p_k^* \in (0,1]$ must be an optimal solution to the problem $\max_{p_j \in [0,1]} \Pi(e^k (p_k - p_k^*) + p^*)$. Since the optimal solution to the last optimization problem is either the right point of the interval $[0,1]$ or in the interior, the derivative of the objective function at the optimal solution must be greater than or equal to zero, which implies that $f_k/q_k \geq \theta\, \mathbb{P}\big\{\mathsf{Pois}(\sum_{i=1}^n q_i\, \lambda_i\, p_i^*) \geq C\big\}$ and the last two inequalities yield $f_k/q_k \geq f_j/q_j$. Thus, we obtain a contradiction to the assumption that $f_1/q_1 > f_2/q_2 > \ldots > f_n/q_n$ and $k > j$. □

Proposition 2 is reassuring since it shows that if we choose the acceptance probabilities by solving problem (4), then the resulting policy randomizes the acceptance decisions for at most one fare class. We either accept or reject all of the requests for other fare classes. Furthermore, this proposition indicates that if the fare to show up probability ratio is larger for a certain fare class, then we are more likely to accept the requests for this fare class. This trend is sensible since if the fare associated with a fare class is larger and the show up probability is smaller, then we generate a larger revenue by accepting the requests for this fare class and the reservations for this fare class are less likely to incur a penalty cost by being denied boarding. Therefore, Proposition 2 suggests f_j/q_j for $j = 1, \ldots, n$ as a metric according to which we should prioritize the fare classes. A similar metric for prioritizing the fare classes appears in Talluri and van Ryzin (2004), where the authors study a revenue management problem under customer choice behavior. In their setting, the customers are offered a set of fare classes, among which they make a choice according to a general choice model. Talluri and van Ryzin (2004) order the fare classes such that $f_1 > f_2 > \ldots > f_n$ and show that if it is optimal to offer fare class j, then it is also optimal to offer fare classes $1, \ldots, j-1$ under certain conditions. Their work assumes that overbooking is not allowed and it is not clear that f_j / q_j for $j = 1, \ldots, n$ is the correct metric for prioritizing the fare classes when overbooking is allowed and the customers choose among the offered fare classes.

In the next section, we build on Proposition 2 to develop an efficient algorithm for solving problem (4) despite the fact that the decision variable in this problem is a high dimensional vector.

4 Solution Algorithm

Proposition 2 allows us to focus on solutions for problem (4) where there is a particular fare class such that the acceptance probabilities are all equal to one before the fare class and the acceptance probabilities are all equal to zero after the fare class. To facilitate the discussion, we define the vector $\delta^j = (\delta^j_1, \ldots, \delta^j_n)$ such that $\delta^j_1 = 1, \ldots, \delta^j_{j-1} = 1, \delta^j_j = 0, \ldots, \delta^j_n = 0$. In this case, for $p_j \in [0,1]$, the solution $e^j p_j + \delta^j$ has an acceptance probability of one for fare classes $1, \ldots, j-1$ and an acceptance probability of zero for fare classes $j+1, \ldots, n$. To choose the best value for p_j, we can solve the problem $\max_{p_j \in [0,1]} \Pi(e^j p_j + \delta^j)$. This optimization problem has a scalar objective function and since $\Pi(p)$ is a concave function of p, we can solve it by using bisection search. Proposition 2 implies that if we check the optimal objective value of the problem $\max_{p_j \in [0,1]} \Pi(e^j p_j + \delta^j)$ for all $j = 1, \ldots, n$ and choose the largest optimal objective value, then we obtain the optimal objective value of problem (4) and the corresponding optimal solution is an optimal solution to problem (4).

In the next proposition, we show that we do not need to solve the problem $\max_{p_j \in [0,1]} \Pi(e^j p_j + \delta^j)$ for all $j = 1, \ldots, n$ to be able to find the one with the largest optimal objective value. We can actually find the problem with the largest optimal objective value by using a closed form expression.

Proposition 3 *Letting ζ^j be the optimal objective value of the problem $\max_{p_j \in [0,1]} \Pi(e^j p_j + \delta^j)$ and j^* be the largest $j = 1, \ldots, n$ that satisfies*

$$\frac{f_j}{q_j} - \theta \, \mathbb{P}\big\{\mathsf{Pois}(\textstyle\sum_{i=1}^{j-1} q_i \, \lambda_i) \geq C\big\} > 0, \tag{6}$$

we have $\zeta^{j^*} \geq \zeta^j$ for all $j = 1, \ldots, n$.

Proof We only show that $\zeta^{j^*} \geq \zeta^j$ for all $j = 1, \ldots, j^* - 1$. One can use a similar argument to show that $\zeta^{j^*} \geq \zeta^j$ for all $j = j^* + 1, \ldots, n$. For any $j < j^*$, differentiating $\Pi(e^j p_j + \delta^j)$ with respect to p_j by using (5), noting that $f_j / q_j > f_{j^*} / q_{j^*}$ and using the fact that the complementary cumulative distribution function of the Poisson distribution is increasing in its mean, we obtain

$$\begin{aligned} \frac{d\Pi(e^j p_j + \delta^j)}{dp_j} &= q_j \, \lambda_j \left\{ \frac{f_j}{q_j} - \theta \, \mathbb{P}\big\{\mathsf{Pois}(\textstyle\sum_{i=1}^{j-1} q_i \, \lambda_i + q_j \, \lambda_j \, p_j) \geq C\big\} \right\} \\ &> q_j \, \lambda_j \left\{ \frac{f_{j^*}}{q_{j^*}} - \theta \, \mathbb{P}\big\{\mathsf{Pois}(\textstyle\sum_{i=1}^{j^*-1} q_i \, \lambda_i) \geq C\big\} \right\} > 0, \end{aligned} \tag{7}$$

where the last inequality follows from the definition of j^*. Therefore, the derivative of the objective function of the problem $\max_{p_j \in [0,1]} \Pi(e^j p_j + \delta^j)$ is always positive for any $j < j^*$ so that one must be an optimal solution to this problem. In other words, $\zeta^j = \Pi(e^j + \delta^j) = \Pi(\delta^{j+1})$ for any $j < j^*$. On the other hand, zero is a feasible but not necessarily an optimal solution to the problem $\max_{p_j \in [0,1]} \Pi(e^j p_j + \delta^j)$ so that $\zeta^j \geq \Pi(\delta^j)$ for any $j < j^*$. Therefore, we obtain $\Pi(\delta^1) \leq \zeta^1 = \Pi(\delta^2) \leq \zeta^2 = \Pi(\delta^3) \leq \ldots = \Pi(\delta^{j^*-1}) \leq \zeta^{j^*-1} = \Pi(\delta^{j^*}) \leq \max_{p_{j^*} \in [0,1]} \Pi(e^{j^*} p_{j^*} + \delta^{j^*}) = \zeta^{j^*}$. □

Proposition 3 implies that we can use the following algorithm to solve problem (4). In the first step, we compute j^* as the largest $j = 1, \ldots, n$ that satisfies (6). In the second step, we use bisection search to solve the problem $\max_{p_{j^*} \in [0,1]} \Pi(e^{j^*} p_{j^*} + \delta^{j^*})$. Letting $p^*_{j^*}$ be an optimal solution to this problem, $e^{j^*} p^*_{j^*} + \delta^{j^*}$ is an optimal solution to problem (4). Computing j^* requires at most n lookups to the cumulative distribution function of the Poisson distribution. Noting that the expression in (7) gives the derivative of the objective function of the problem $\max_{p_{j^*} \in [0,1]} \Pi(e^{j^*} p_{j^*} + \delta^{j^*})$, each iteration of bisection search to solve the problem $\max_{p_{j^*} \in [0,1]} \Pi(e^{j^*} p_{j^*} + \delta^{j^*})$ requires one lookup to the cumulative distribution function of the Poisson distribution.

The expression in (6) has interesting connections with the earlier literature. Section 4.2.1.2 in Talluri and van Ryzin (2005) describes an overbooking model with a single fare class. In this model, an accepted request generates a revenue of f and a reservation shows up at the departure time with probability q. The capacity on the flight leg is C. If a reservation that shows up at the departure time cannot be accommodated on the flight leg, then we incur a penalty cost of θ for each reservation that we deny boarding. Using $\mathsf{Binom}(n, q)$ to denote a Binomial random variable with parameters (n, q), Section 4.2.1.2 in Talluri and van Ryzin (2005) shows that it is optimal to accept at most n^* reservations, where n^* is the largest integer that satisfies $f/q - \theta\, \mathbb{P}\big\{\mathsf{Binom}(n^*, q) \geq C\big\} > 0$. The form of this optimality condition closely resembles the expression in (6). The random variable $\mathsf{Binom}(n^*, q)$ captures the number of reservations that show up given that we accept n^* requests and we compare the complementary cumulative distribution of this random variable with the fare to show up probability ratio. Similarly, the random variable $\mathsf{Pois}(\sum_{i=1}^{j-1} q_i\, \lambda_i)$ in (6) captures the number of reservations that show up given that we accept the requests for only fare classes $1, \ldots, j-1$ and we compare the complementary cumulative distribution of this random variable with the fare to show up probability ratio for fare class j. Therefore, our model can be interpreted as an extension of this earlier work to multiple fare classes.

The expression in (6) is also related to two fare class capacity control models. Littlewood (1972) considers a model with no overbooking, where C units of capacity on the flight leg needs to be allocated between the requests from the discount and full fare classes, which respectively provide revenues of p_d and p_f for each accepted request. The demand from the discount fare class arrives first, followed by the demand from the full fare class. Using the random variable D_f to denote the demand from the full fare class, Littlewood (1972) shows that it is optimal to accept at most b^* reservations from the discount fare class, where b^* is the smallest integer that satisfies $p_d - p_f\, \mathbb{P}\big\{D_f \geq C - b^*\big\} < 0$. Therefore, quantile solutions similar to those in (6) appear in the earlier literature.

In the next section, we use the properties of the Poisson distribution to give a performance bound for the policies that use fixed acceptance probabilities.

5 Performance Bound

The policy that we describe in Section 3 is an open loop policy as it accepts the requests for different fare classes with fixed probabilities and its decisions do not depend on how many reservations have already been accepted. In this section, we develop a performance bound for such open loop policies. In particular, we study an open loop policy obtained from a somewhat crude deterministic approximation that is formulated under the assumption that all of the random variables take on their expected values. Using the decision variables $x = (x_1, \ldots, x_n)$ with the interpretation that x_j is the probability with which we accept the requests for fare class j, this deterministic approximation is given by

$$\max_{x \in [0,1]^n} \sum_{j=1}^{n} f_j \, \lambda_j \, x_j - \theta \, [\textstyle\sum_{j=1}^n q_j \, \lambda_j \, x_j - C]^+. \tag{8}$$

We note that the objective function of problem (8) is the deterministic analogue of the net expected revenue function in (3). Therefore, problem (8) can be interpreted as a deterministic approximation to the original overbooking problem that is formulated under the assumption that the arrivals of the requests and the show up decisions of the reservations take on their expected values. In this section, we show that it is possible to come up with a reasonable performance bound for the policies obtained by problem (8) despite the fact that this problem disregards the uncertainty. In the revenue management literature, it is customary to formulate deterministic approximations to estimate the optimal net expected revenue and there are numerous papers that develop performance bounds for these approximations. While our analysis in this section parallels the ones in the earlier literature, the performance bound that we derive allows us to make several useful observations about the problem parameters that may affect the performance of open loop policies.

In the next lemma, we show that the optimal objective value of problem (8) is an upper bound on the optimal net expected revenue. In other words, letting $\bar{0} = (0, \ldots, 0) \in \Re^n$ and noting the Hamilton Jacobi Bellman equation in (2), the optimal net expected revenue is $V_0(\bar{0})$ and the optimal objective value of problem (8) turns out to be an upper bound on $V_0(\bar{0})$. The proof of this lemma follows from a standard argument that uses Jensen's inequality and we defer it to Appendix A.

Lemma 4 *Letting Z^* be the optimal objective value of problem (8), we have $Z^* \geq V_0(\bar{0})$.*

In the next proposition, we give a performance bound for the acceptance probabilities that we obtain by solving problem (8). Letting x^* be an optimal solution to problem (8), the net expected revenue that we obtain by using the acceptance probabilities x^* is $\Pi(x^*)$. Since $V_0(\bar{0})$ is the optimal net expected revenue, we have $\Pi(x^*)/V_0(\bar{0}) \leq 1$. In the next proposition, we give a lower bound on the ratio $\Pi(x^*)/V_0(\bar{0})$ in terms of the problem parameters.

Proposition 5 *Letting x^* be an optimal solution to problem (8), we have*

$$\frac{\Pi(x^*)}{V_0(\bar{0})} \geq 1 - \frac{\theta \, \max\{1, \sum_{j=1}^n q_j \, \lambda_j / C\}}{\sqrt{2\,\pi} \, \frac{f_n}{q_n} \, \sqrt{C}}. \tag{9}$$

Proof Since we have $\Pi(x^*)/V_0(\bar{0}) \geq \Pi(x^*)/Z^*$ by Lemma 4, it is enough to show that the expression on the right side of (9) provides a lower bound on $\Pi(x^*)/Z^*$. We derive an upper bound on the expectation $\mathbb{E}\big\{[\mathsf{Pois}(\alpha) - C]^+\big\}$ by noting that

$$\mathbb{E}\big\{[\mathsf{Pois}(\alpha) - C]^+\big\} - [\alpha - C]^+ = \sum_{k=0}^{\infty}[[k-C]^+ - [\alpha - C]^+]\,\frac{e^{-\alpha}\,\alpha^k}{k!}$$
$$\leq \sum_{k=C+1}^{\infty}[[k-C]^+ - [\alpha - C]^+]\,\frac{e^{-\alpha}\,\alpha^k}{k!} \leq \sum_{k=C+1}^{\infty}[k-\alpha]\,\frac{e^{-\alpha}\,\alpha^k}{k!}$$
$$= \alpha\,\frac{e^{-\alpha}\,\alpha^C}{C!} \leq \alpha\,\frac{e^{-C}\,C^C}{C!} \leq \alpha\,\frac{e^{-C}\,C^C}{\sqrt{2\,\pi\,C}(C/e)^C} = \frac{\alpha}{\sqrt{2\,\pi\,C}}, \qquad (10)$$

where the third inequality follows from the fact that the function $f(\alpha) = e^{-\alpha}\,\alpha^C$ attaints its maximum at $\alpha = C$ and the fourth inequality follows by noting that we have $C! \geq \sqrt{2\,\pi\,C}\,(C/e)^C$ by Stirling's approximation. The rest of the proof proceeds in two parts.

In the first part, we assume that $\sum_{j=1}^n q_j\,\lambda_j\,x_j^* > C$. We observe that for any $j = 1,\ldots,n$ with $x_j^* > 0$, we must have $f_j/q_j \geq \theta$. To see this, the optimal objective value of problem (8) is given by $\sum_{j=1}^n [f_j - \theta\,q_j]\,\lambda_j\,x_j^* + \theta\,C$ and if there exists a $j = 1,\ldots,n$ with $x_j^* > 0$ and $f_j/q_j < \theta$, then we can obtain a strictly better objective function value by decreasing the value of the decision variable x_j. In this case, letting j^* be the largest $j = 1,\ldots,n$ such that $x_j^* > 0$, the optimal objective value of problem (8) satisfies $Z^* = \sum_{j=1}^n f_j\,\lambda_j\,x_j^* - \theta\,[\sum_{j=1}^n q_j\,\lambda_j\,x_j^* - C] \geq \sum_{j=1}^n f_j\,\lambda_j\,x_j^* - \frac{f_{j^*}}{q_{j^*}}[\sum_{j=1}^n q_j\,\lambda_j\,x_j^* - C] = \sum_{j=1}^n q_j\,\big[\frac{f_j}{q_j} - \frac{f_{j^*}}{q_{j^*}}\big]\,\lambda_j\,x_j^* + \frac{f_{j^*}}{q_{j^*}}\,C \geq \frac{f_n}{q_n}\,C$, where the second inequality follows from the fact that for any $j = 1,\ldots,n$ with $x_j^* > 0$, we have $j \leq j^*$ so that $f_j/q_j \geq f_{j^*}/q_{j^*}$. Therefore, we obtain

$$\frac{\Pi(x^*)}{Z^*} = \frac{\sum_{j=1}^n f_j\,\lambda_j\,x_j^* - \theta\,\mathbb{E}\big\{[\mathsf{Pois}(\sum_{j=1}^n q_j\,\lambda_j\,x_j^*) - C]^+\big\}}{\sum_{j=1}^n f_j\,\lambda_j\,x_j^* - \theta\,[\sum_{j=1}^n q_j\,\lambda_j\,x_j^* - C]}$$
$$\geq \frac{\sum_{j=1}^n f_j\,\lambda_j\,x_j^* - \theta\,[\sum_{j=1}^n q_j\,\lambda_j\,x_j^* - C] - \frac{\theta}{\sqrt{2\,\pi C}}\,\sum_{j=1}^n q_j\,\lambda_j\,x_j^*}{\sum_{j=1}^n f_j\,\lambda_j\,x_j^* - \theta\,[\sum_{j=1}^n q_j\,\lambda_j\,x_j^* - C]}$$
$$\geq 1 - \frac{\frac{\theta}{\sqrt{2\,\pi C}}\,\sum_{j=1}^n q_j\,\lambda_j\,x_j^*}{\frac{f_n}{q_n}\,C} \geq 1 - \frac{\theta\,\sum_{j=1}^n q_j\,\lambda_j}{\sqrt{2\,\pi}\,\frac{f_n}{q_n}\,\sqrt{C}\,C}, \qquad (11)$$

where the first inequality follows from (10), the second inequality follows from the fact that $Z^* \geq \frac{f_n}{q_n}\,C$ and the third inequality follows from the fact that $x_j^* \in [0,1]$ for all $j = 1,\ldots,n$.

In the second part, we assume that $\sum_{j=1}^n q_j\,\lambda_j\,x_j^* \leq C$. Noting that $\frac{f_j}{q_j} \geq \frac{f_n}{q_n}$ for all $j = 1\ldots,n$, the optimal objective value of problem (8) satisfies $Z^* = \sum_{j=1}^n f_j\,\lambda_j\,x_j^* \geq \frac{f_n}{q_n}\,\sum_{j=1}^n q_j\,\lambda_j\,x_j^*$. We have

$$\frac{\Pi(x^*)}{Z^*} = \frac{\sum_{j=1}^n f_j\,\lambda_j\,x_j^* - \theta\,\mathbb{E}\big\{[\mathsf{Pois}(\sum_{j=1}^n q_j\,\lambda_j\,x_j^*) - C]^+\big\}}{\sum_{j=1}^n f_j\,\lambda_j\,x_j^*}$$
$$\geq \frac{\sum_{j=1}^n f_j\,\lambda_j\,x_j^* - \frac{\theta}{\sqrt{2\,\pi C}}\,\sum_{j=1}^n q_j\,\lambda_j\,x_j^*}{\sum_{j=1}^n f_j\,\lambda_j\,x_j^*} \geq 1 - \frac{\theta}{\sqrt{2\,\pi}\,\frac{f_n}{q_n}\,\sqrt{C}}, \qquad (12)$$

where the first inequality follows from (10) and the second inequality follows from the fact that $\sum_{j=1}^n f_j \, \lambda_j \, x_j^* \geq \frac{f_n}{q_n} \sum_{j=1}^n q_j \, \lambda_j \, x_j^*$. The result follows by combining (11) and (12). $\Box$

Since we have $\max_{p \in [0,1]^n} \Pi(p)/V_0(\bar{0}) \geq \Pi(x^*)/V_0(\bar{0})$, Proposition 5 also provides a performance bound for the acceptance probabilities that we obtain by solving problem (4). We observe that the performance bound in Proposition 5 gets weaker as the penalty cost increases and the degradation in the performance bound can be interpreted as the price of not taking the uncertainty into consideration in problem (8). If we accept all of the requests, then the total expected number of reservations that show up at the departure time is given by $\sum_{j=1}^n q_j \, \lambda_j$. Therefore, the ratio $\sum_{j=1}^n q_j \, \lambda_j / C$ can be visualized as the load factor on the flight leg. As long as the load factor is below one, the performance bound gets stronger as the capacity on the flight leg increases and it is independent of the demand for different fare classes. Similarly, if we increase the capacity on the flight leg while keeping the load factor constant, then the performance bound gets stronger and approaches one with a rate of $1/\sqrt{C}$. To get a feel for the performance bound in Proposition 5, we consider the case where the penalty cost is 800, the lowest fare is 100, all of the reservations show up with a probability of 0.8 and the capacity on the flight leg is 100. If the load factor is less than one, then the performance bound on the right side of (9) evaluates to 0.74. This is to say that the acceptance probabilities that we obtain by solving problem (4) or problem (8) provide at least 74% of the optimal net expected revenue. Finally, the performance bound gets weaker as the average numbers of requests for different fare classes increase and this degradation can be attributed to the fact that one should deal with the uncertainty due to the show up decisions of the reservations even if we have enough demand to fill the capacity with high probability.

Problem (8) has connections with the knapsack problem and this connection provides an interesting comparison between problems (4) and (8). In particular, if we define a new decision variable $S = \sum_{j=1}^n q_j \, \lambda_j \, x_j - C$, then problem (8) can be written as

$$\max \left\{ \sum_{j=1}^n f_j \, \lambda_j \, x_j - \theta \, [S]^+ \; : \; \sum_{j=1}^n q_j \, \lambda_j \, x_j - S = C, \; x \in [0,1]^n, \; S \in \Re \right\}. \tag{13}$$

For any fixed value of the decision variable S, the problem above is a knapsack problem. The items are indexed by $j = 1, \ldots, n$. The capacity of the knapsack is $C + S$, the utility of item j is $f_j \, \lambda_j$ and the capacity consumption of item j is $q_j \, \lambda_j$. Therefore, we can solve the problem above by filling the knapsack with the items in the increasing order of utility to capacity consumption ratios. Noting that the utility to capacity consumption ratio of item j is $[f_j \, \lambda_j]/[q_j \, \lambda_j] = f_j/q_j$ and $f_1/q_1 > f_2/q_2 > \ldots > f_n/q_n$, this is precisely the same ordering prescribed by Proposition 2. In other words, problems (4) and (8) prioritize the fare classes in exactly the same fashion, but problem (4) uses the probability distributions of the requests when choosing the last fare class for which we still accept the requests, whereas problem (8) carries out a simple deterministic computation. The knapsack problem in (13) is deterministic. The paper by van Slyke and Young (2000) studies stochastic knapsack problems, where the items arrive randomly over time and one needs to decide whether to put each item into the knapsack on the fly. The authors demonstrate that the optimal policies in stochastic knapsack problems can be nonintuitive and show the connections of stochastic knapsack problems to revenue management. However, their work does not consider overbooking, since it does not allow exceeding the capacity of the knapsack.

6 Practical Extensions

In this section, we extend our approach in two practically useful directions by incorporating general convex penalty cost functions and demand processes that are more variable than the Poisson process.

6.1 General Convex Penalty Costs

Our development in earlier sections assumes that if we cannot accommodate a reservation that shows up at the departure time, then we incur the same penalty cost of θ for each reservation that we deny boarding. In some practical settings, it is possible that we incur a progressively higher penalty cost for each denied reservation. For example, if we attempt to convince the passengers to voluntarily give up their seats by offering them incentives such as travel vouchers and alternative accommodations, then each additional passenger that we attempt to convince may require raising the incentives, resulting in a progressively higher penalty cost. In this section, we begin by generalizing our development to capture this kind of a progressively higher penalty cost for each denied reservation.

We let $\Theta(\cdot) : \Re \to \Re$ be an increasing convex function with bounded directional derivatives and assume that if there are n reservations that show up at the departure time, then we incur a penalty cost of $\Theta(n)$. The convexity of $\Theta(\cdot)$ is intended to capture the fact that the penalty cost for each denied reservation is progressively higher. It is natural to impose the condition that $\Theta(n) = 0$ whenever $n \leq C$, since we do not incur any penalty costs when the number of reservations that show up at the departure time does not exceed the capacity on the flight leg. In this case, it is possible to check that all of our results in earlier sections go through with minor modifications. In particular, if we use the acceptance probabilities p, then the net expected revenue that we obtain can be written as

$$\Pi(p) = \sum_{j=1}^{n} f_j \, \lambda_j \, p_j - \mathbb{E}\big\{\Theta(\mathsf{Pois}(\textstyle\sum_{j=1}^{n} q_j \, \lambda_j \, p_j))\big\}, \tag{14}$$

in which case, we can compute the best set of acceptance probabilities by solving problem (4), but we use the net expected revenue function in (14) instead of the one in (3). Lemma 6 in Appendix A shows that $\mathbb{E}\big\{\Theta(\mathsf{Pois}(\alpha))\big\}$ is a differentiable and convex function of α and its derivative with respect to α is given by $\mathbb{E}\big\{\Theta(\mathsf{Pois}(\alpha)+1) - \Theta(\mathsf{Pois}(\alpha))\big\}$. Therefore, differentiating $\Pi(p)$ with respect to p_j and using Lemma 6 and the chain rule, we obtain

$$\begin{aligned}\frac{\partial \Pi(p)}{\partial p_j} &= f_j \, \lambda_j - q_j \, \lambda_j \, \mathbb{E}\big\{\Theta(\mathsf{Pois}(\textstyle\sum_{i=1}^{n} q_i \, \lambda_i \, p_i) + 1) - \Theta(\mathsf{Pois}(\textstyle\sum_{i=1}^{n} q_i \, \lambda_i \, p_i))\big\} \\ &= q_j \, \lambda_j \left\{ \frac{f_j}{q_j} - \mathbb{E}\big\{\Theta(\mathsf{Pois}(\textstyle\sum_{i=1}^{n} q_i \, \lambda_i \, p_i) + 1) - \Theta(\mathsf{Pois}(\textstyle\sum_{i=1}^{n} q_i \, \lambda_i \, p_i))\big\} \right\}. \end{aligned} \tag{15}$$

Using (15), it is possible to check that Proposition 2 continues to hold when we use a general convex penalty cost function $\Theta(\cdot)$. Similarly, we can show that Proposition 3 continues to hold as long as we replace the expression $\theta \, \mathbb{P}\big\{\mathsf{Pois}(\sum_{i=1}^{j-1} q_i \, \lambda_i) \geq C\big\}$ in (6) with $\mathbb{E}\big\{\Theta(\mathsf{Pois}(\sum_{i=1}^{j-1} q_i \, \lambda_i) + 1) - \Theta(\mathsf{Pois}(\sum_{i=1}^{j-1} q_i \, \lambda_i))\big\}$. To develop a performance bound, we use the deterministic approximation

$$\max_{x \in [0,1]^n} \; \sum_{j=1}^{n} f_j \, \lambda_j \, x_j - \Theta(\textstyle\sum_{j=1}^{n} q_j \, \lambda_j \, x_j) \tag{16}$$

instead of the one in (8). Letting Z^* be the optimal objective value of problem (16), Lemma 4 continues to hold. Finally, it is possible to check that the performance bound in Proposition 5 continues to hold once we replace θ in (9) by the bound on the directional derivative of $\Theta(\cdot)$. Therefore, all of our results hold when we incur a progressively higher penalty cost for each denied reservation. However, it is worthwhile to emphasize that if we have the constant penalty cost of θ for each denied reservation, then the expression in (6) has a close connection with the single fare class overbooking model in Section 4.2.1.2 in Talluri and van Ryzin (2005), but this connection becomes less apparent when we use a general convex penalty cost function $\Theta(\cdot)$.

6.2 Random Arrival Rates

Another practical extension for our approach involves allowing more variability in the numbers of requests for different fare classes. Our development in earlier sections assumes that the total number of requests for fare class j has a Poisson distribution with mean λ_j. The coefficient of variation for this distribution is $1/\sqrt{\lambda_j}$, which can be small, especially when the mean demand for fare class j is large. A natural question is whether we can incorporate more variability into the demand distributions in a tractable fashion. One possible approach for incorporating more variability is to assume that the arrival rates for the requests for different fare classes are themselves random variables. In particular, we assume that $\lambda = (\lambda_1, \ldots, \lambda_n)$ is a random variable, taking value $\lambda^\ell = (\lambda_1^\ell, \ldots, \lambda_n^\ell)$ with probability ρ^ℓ. We need to choose the acceptance probabilities before we get to know the arrival rates for the requests for different fare classes. In this case, the sources of uncertainty become the arrival rates for the requests for different fare classes, the arrivals of the requests themselves and the show up decisions of the reservations. Therefore, we need to modify the net expected revenue function in (3) as

$$\Pi(p) = \sum_{\ell=1}^{L} \rho^\ell \left\{ \sum_{j=1}^{n} f_j \, \lambda_j^\ell \, p_j - \theta \, \mathbb{E}\big\{[\mathsf{Pois}(\textstyle\sum_{j=1}^n q_j \, \lambda_j^\ell \, p_j) - C]^+\big\} \right\}, \tag{17}$$

where L is the number of possible realizations for the random variable λ. Noting the discussion that follows Lemma 1, the expression in the curly brackets in (17) is a concave function of p so that the net expected revenue in (17) is also a concave function of p. Therefore, the problem of maximizing the net expected revenue continues to be a convex optimization problem, but Propositions 2 and 3 do not necessarily hold when we use the net expected revenue function in (17).

To get a feel for the complications that may arise when the arrival rates are random variables, we consider a problem instance where there are two fare classes and the arrival rates $\lambda = (\lambda_1, \lambda_2)$ can take two possible values $\lambda^1 = (\lambda_1^1, \lambda_2^1) = (1, 10)$ and $\lambda^2 = (\lambda_1^2, \lambda_2^2) = (10, 1)$ with equal probabilities $\rho^1 = \rho^2 = 0.5$. The fares and show up probabilities are $(f_1, f_2) = (2, 1)$ and $(q_1, q_2) = (1, 1)$. The capacity on the flight leg is $C = 5$. The penalty cost is $\theta = 3$. For this problem instance, one can verify that $(p_1^*, p_2^*) \approx (0.56, 0.31)$ maximizes the net expected revenue function in (17) and yields a net expected revenue of about 5.24. On the other hand, if Proposition 2 held, then the optimal acceptance probabilities would have the form $(p_1^*, p_2^*) = (p_1, 0)$ for some $p_1 \in [0, 1]$ or $(p_1^*, p_2^*) = (1, p_2)$ for some $p_2 \in [0, 1]$. In Figure 8, the thin data series plot $\Pi((p_1, 0))$ for all $p_1 \in [0, 1]$ and the dashed data series plot $\Pi((1, p_2))$ for all $p_2 \in [0, 1]$ and both of these data series lie below the optimal net expected revenue of 5.24 plotted by the thick data series, which indicates that the optimal acceptance probabilities cannot have the form given in Proposition 2. In this problem instance, the possible realizations of the arrival rates are such that the arrival rate for the requests for one fare class is large whenever the arrival rate for the requests for the other fare class is small. Therefore, it becomes optimal to use fractional acceptance probabilities for both fare classes to protect against the uncertainty in the arrival rates, rather than deterministically accepting or rejecting all of the requests for a certain fare class.

Despite the fact that the structural results in Propositions 2 and 3 do not necessarily hold, it is not difficult to solve the problem $\max_{p \in [0,1]^n} \Pi(p)$ when we use the net expected revenue function in (17). Since the net expected revenue in (17) is a concave function of p, one possible approach for solving the problem $\max_{p \in [0,1]^n} \Pi(p)$ is to use a standard cutting plane method. In Appendix B, we describe a cutting plane method that generates a sequence of solutions $\{p^k\}_k$ and Theorem 7.7 in Ruszczynski (2006) shows that this sequence of solutions satisfy $\lim_{k \to \infty} \Pi(p^k) = \Pi(p^*)$, where p^* is an optimal solution to the problem $\max_{p \in [0,1]^n} \Pi(p)$. The cutting plane method in Appendix B is simple to implement as it requires solving a linear program at each iteration. Another possible approach for solving the problem $\max_{p \in [0,1]^n} \Pi(p)$ is based on the observation that the function $F(\alpha) = \mathbb{E}\big\{[\mathsf{Pois}(\alpha) - C]^+\big\}$ is a convex scalar function of α, in which case, we can construct a piecewise linear convex approximation to $F(\cdot)$ by computing $F(\cdot)$ at a finite number of points and interpolating the value of the function between the points. Letting $\hat{F}(\cdot)$ be a piecewise linear convex approximation to $F(\cdot)$, we can find an approximate maximizer of the net expected revenue function in (17) by solving the problem

$$\max_{p \in [0,1]^n} \quad \sum_{\ell=1}^{L} \rho^\ell \left\{ \sum_{j=1}^{n} f_j \, \lambda_j^\ell \, p_j - \theta \, \hat{F}\big(\textstyle\sum_{j=1}^n q_j \, \lambda_j^\ell \, p_j\big) \right\}. \tag{18}$$

Since $\hat{F}(\cdot)$ is a piecewise linear convex function, we can formulate the problem above as a linear program. Furthermore, noting that $p_j \in [0,1]$, the relevant domain of $\hat{F}(\cdot)$ is between zero and $U = \max\big\{ \sum_{j=1}^n q_j \, \lambda_j^\ell \, : \, \ell = 1, \ldots, L \big\}$ and we can obtain an accurate approximation by computing $F(\cdot)$ at a large number of points over the interval $[0, U]$. A natural idea is to begin with a relatively small number of points to obtain an approximation $\hat{F}(\cdot)$, solve problem (18) with this approximation to obtain a solution and refine the approximation $\hat{F}(\cdot)$ around this solution. The cutting plane method in Appendix B requires iteratively solving a sequence of linear programs and it generates an optimal solution in the limit, whereas working with problem (18) allows directly formulating and solving this problem as a linear program, but the accuracy of this approach is limited by the accuracy of the approximation $\hat{F}(\cdot)$. In our computational experiments, we use the cutting plane method.

Random arrival rates can be quite useful for incorporating more variability into the numbers of requests for different fare classes. For example, if the number of requests for a fare class has Poisson distribution with mean 100, then the coefficient of variation for this distribution is 0.1. On the other hand, if the number of requests has Poisson distribution with mean 50 or 150, each case occurring with a probability of 0.5, then the coefficient of variation for this distribution is 0.51. Finally, we emphasize that the same approach can be used to model the show up probabilities $q = (q_1, \ldots, q_n)$ as random variables, which can be useful to capture show up decisions that are more variable than those provided by the binomial distribution.

7 Computational Experiments

In this section, we test the performance of the acceptance probabilities that are obtained by solving problem (4).

7.1 Experimental Setup

In our computational experiments, we consider test problems with eight fare classes. Normalizing the length of the selling horizon to one and counting from the lowest to the highest fare class, the arrival rates for the requests for the eight fare classes are $(20, 16, 16, 15, 6, 15, 6, 6)$ so that the total expected number of requests is 100. We note that the demand is somewhat tilted towards the lower fare classes, which is in parallel with what is usually observed in practice. The fare associated with the lowest fare class is fixed at 100 in all of our test problems. The fare associated with the highest fare class is κ times the fare associated with the lowest fare class. The fares for other fare classes are evenly distributed over the interval $[100, \kappa\, 100]$. For the four fare classes with the lowest fares, the probability that a reservation shows up at the departure time is q^L. The intensity functions for the arrivals of the requests for these fare classes are calibrated in such a fashion that the requests occur more frequently early in the selling horizon. For the four fare classes with the highest fares, the probability that a reservation shows up at the departure time is q^H and the requests for these fare classes tend to occur more frequently later in the selling horizon. The penalty cost of denying boarding to a reservation is γ times the fare associated with the highest fare class. Noting that $\sum_{j=1}^n q_j\, \lambda_j$ corresponds to the expected number of reservations that show up at the departure time when we accept all of the requests, we set the capacity on the flight leg as $\lceil \sum_{j=1}^n q_j\, \lambda_j / \rho \rceil$, where $\lceil \cdot \rceil$ denotes the round up function and ρ is a parameter that we vary. Therefore, small values for ρ provide test problems with plenty of capacity and allow accepting a large portion of the requests without running over the capacity, whereas large values for ρ provide test problems with tight capacity and require careful planning as to which requests to accept. This setup yields capacities on the flight leg that range between 40 and 90. We label our test problems by the tuple $(\kappa, q^L, q^H, \gamma, \rho)$ and use $\kappa \in \{4, 8\}$, $q^L \in \{0.7, 0.9\}$, $q^H \in \{0.7, 0.9\}$, $\gamma \in \{2, 4\}$ and $\rho \in \{1.0, 1.2, 1.6\}$, in which case, we obtain 48 test problems.

7.2 Benchmark Strategies

We compare the performances of the following three benchmark strategies.

Optimal Policy (OPT) This benchmark strategy corresponds to a good approximation to the optimal policy that we obtain by solving the Hamilton Jacobi Bellman equation in (2). We note that q^L and q^H are the two possible values for the show up probabilities and the only feature of the reservations that affects the penalty cost at the departure time is their show up probabilities. Therefore, we can solve the Hamilton Jacobi Bellman equation in (2) by using a two dimensional state variable that keeps track of the total numbers of reservations with the two different show up probabilities. To deal with the continuous time, OPT discretizes the selling horizon with time steps $\Delta = \tau / 1,000$ and uses the discrete time version of the Hamilton Jacobi Bellman equation in (1).

Acceptance Probabilities (APR) This benchmark strategy corresponds to the acceptance probabilities that we obtain by solving problem (4). Our implementation of APR divides the selling horizon into ten equal segments and resolves problem (4) at the beginning of each segment to refresh the acceptance probabilities. Our goal in periodically resolving problem (4) is to capture the industry practice, where the policy parameters are periodically refreshed as the state of the reservations evolves.

One concern when resolving problem (4) at the beginning of each segment is that the reservations that show up for a particular fare class are generated by a mixture of two sets of requests, where the first set of requests are those that have already been accepted and the second set of requests are those that are accepted in the future. In particular, if we are at the beginning of segment s, the number of requests for fare class j that we have already accepted is given by r_j and we use the acceptance probability p_j to accept the future requests, then the number of reservations for fare class j that show up at the departure time is given by $\mathsf{Pois}(q_j\, \lambda_j\, p_j) + \mathsf{Binom}(r_j, q_j)$, where λ_j is understood as the expected number of arrivals for the requests for fare class j in the future, which is given by $\lambda_j = \int_{(s-1)\,\frac{\tau}{10}}^{\tau} \Lambda_j(t)\, dt$. In this case, the net expected revenue function at the beginning of segment s can be written as

$$\Pi(p) = \sum_{j=1}^{n} f_j\, \lambda_j\, p_j - \theta\, \mathbb{E}\big\{[\mathsf{Pois}(\textstyle\sum_{j=1}^n q_j\, \lambda_j\, p_j) + \sum_{j=1}^n \mathsf{Binom}(r_j, q_j) - C]^+\big\}, \tag{19}$$

where the random variables $\mathsf{Pois}(\sum_{j=1}^n q_j\, \lambda_j\, p_j)$ and $\mathsf{Binom}(r_j, q_j)$ for all $j = 1, \ldots, n$ are independent of each other. The maximizer of the net expected revenue function in (19) continues to satisfy Propositions 2 and 3 so that we can use these propositions to maximize the net expected revenue. To see these results, we condition on $\sum_{j=1}^n \mathsf{Binom}(r_j, q_j)$ and write the conditional version of the net expected revenue function in (19) as $\Pi(p \,|\, B) = \sum_{j=1}^n f_j\, \lambda_j\, p_j - \theta\, \mathbb{E}\big\{[\mathsf{Pois}(\sum_{j=1}^n q_j\, \lambda_j\, p_j) + B - C]^+\big\}$ so that $\Pi(p) = \mathbb{E}\big\{\Pi(p \,|\, \sum_{j=1}^n \mathsf{Binom}(r_j, q_j))\big\}$. We observe that $\Pi(p \,|\, B)$ has the same form as the net expected revenue function in (3) except for the fact that $\Pi(p \,|\, B)$ replaces C in (3) by $C - B$. Therefore, the derivative of $\Pi(p \,|\, B)$ with respect to p_j can be obtained by replacing C in (5) by $C - B$, which implies that

$$\begin{aligned}&\frac{\partial \Pi(p \,|\, \sum_{i=1}^n \mathsf{Binom}(r_i, q_i))}{\partial p_j} \\ &\qquad = f_j\, \lambda_j - \theta\, q_j\, \lambda_j\, \mathbb{P}\big\{\mathsf{Pois}(\textstyle\sum_{i=1}^n q_i\, \lambda_i\, p_i) \geq C - \sum_{i=1}^n \mathsf{Binom}(r_i, q_i) \,|\, \sum_{i=1}^n \mathsf{Binom}(r_i, q_i)\big\}.\end{aligned} \tag{20}$$

Taking expectations in (20) and noting that the expectation of the conditional probability on the right side of (20) is $\mathbb{P}\big\{\mathsf{Pois}(\sum_{i=1}^n q_i\, \lambda_i\, p_i) + \sum_{i=1}^n \mathsf{Binom}(r_i, q_i) \geq C\big\}$, we obtain

$$\begin{aligned}\frac{\partial \Pi(p)}{\partial p_j} &= \frac{\partial \mathbb{E}\big\{\Pi(p \,|\, \sum_{i=1}^n \mathsf{Binom}(r_i, q_i))\big\}}{\partial p_j} = \mathbb{E}\left\{\frac{\partial \Pi(p \,|\, \sum_{i=1}^n \mathsf{Binom}(r_i, q_i))}{\partial p_j}\right\} \\ &= f_j\, \lambda_j - \theta\, q_j\, \lambda_j\, \mathbb{P}\big\{\mathsf{Pois}(\textstyle\sum_{i=1}^n q_i\, \lambda_i\, p_i) + \sum_{i=1}^n \mathsf{Binom}(r_i, q_i) \geq C\big\},\end{aligned} \tag{21}$$

where interchanging the order of the expectation and the derivative in the second equality is justified by noting that the random variable $\sum_{i=1}^n \mathsf{Binom}(r_i, q_i)$ has a finite number of realizations. Using (21), it is possible to check that Proposition 2 continues to hold for the maximizer of the net expected revenue function in (19). Similarly, we can show that Proposition 3 continues to hold as long as we replace the expression $\mathbb{P}\big\{\mathsf{Pois}(\sum_{i=1}^{j-1} q_i\, \lambda_i) \geq C\big\}$ in (6) with $\mathbb{P}\big\{\mathsf{Pois}(\sum_{i=1}^{j-1} q_i\, \lambda_i) + \sum_{i=1}^n \mathsf{Binom}(r_i, q_i) \geq C\big\}$.

Deterministic Approximation (DET) Similar to APR, this benchmark strategy uses acceptance probabilities, but these acceptance probabilities are obtained by solving the deterministic approximation in (8). DET may perform poorly as it ignores the uncertainty and our goal in using DET is to demonstrate how much additional benefit we obtain by using a benchmark strategy that addresses the probabilistic nature of the problem. Similar to APR, we divide the selling horizon into ten equal segments and refresh the acceptance probabilities for DET at the beginning of each segment.

7.3 Computational Results with Known Arrival Rates

We give our main computational results in Tables 1 and 2. In particular, Tables 1 and 2 respectively focus on the test problems where the ratio between the highest and lowest fares is four and eight. The first column in these tables shows the problem characteristics by using the tuple $(\kappa, q^L, q^H, \gamma, \rho)$. The second, third and fourth columns show the net expected revenues obtained by OPT, APR and DET. We estimate these net expected revenues by simulating the arrivals of the requests and the show up decisions over 1,000 sample paths. We use common random numbers when simulating the performances of the different benchmark strategies. The fifth and sixth columns show the percent gaps between the net expected revenues obtained by OPT and the remaining two benchmark strategies.

The results indicate that APR performs remarkably well. The average performance gap between APR and OPT is 1.57% and the net expected revenues obtained by APR deviate from those obtained by OPT by at most 4.57%. The performance of DET can be poor and the net expected revenues obtained by DET can lag behind those obtained by OPT by up to 17.08%. The average performance gap between DET and OPT is 8.47%. Overall, the ability of APR to take the uncertainty into consideration pays off and APR can provide significant improvements over DET. There are a few crucial problem parameters that affect the performance of APR. In particular, as ρ increases and the capacity on the flight leg gets tighter, the performance gap between APR and OPT increases. For the test problems with a load factor of one, the performance gap between APR and OPT is usually a fraction of a percent. The performance gap can increase up to 4.57% as the load factor increases to 1.6. If we only consider the test problems with $\rho = 1.0$, then the average gap between the net expected revenues obtained by APR and OPT is only 0.81%, whereas if we only consider the test problems with $\rho = 1.6$, then the same average gap is 2.51%. On the other hand, as γ increases and it becomes more costly to deny boarding to a reservation, the gap between the net expected revenues obtained by APR and OPT increases. For the test problems with $\gamma = 2$, the average performance gap between APR and OPT is 1.09%, whereas for the test problems with $\gamma = 4$, the average performance gap between APR and OPT is 2.04%. To get a feel for how the performance gaps vary with γ and ρ, Table 3 shows the performance gaps between the different benchmark strategies averaged over the test problems with a particular value for γ and ρ. The first column in this table shows the problem characteristics over which we average the performance gap. The second column shows the performance gap between OPT and APR, the third column shows the performance gap between OPT and DET and the fourth column shows the performance gap between APR and DET. For example, the top left entry in Table 3 shows the performance gap between OPT and APR averaged over all of the test problems with $(\gamma, \rho) = (2, 1.0)$. The results indicate that the performance gaps between OPT and APR increase when we increase γ or ρ, keeping the other parameter constant. Similarly, the performance gaps between OPT and DET increase when we increase γ, keeping ρ constant. We observe a similar trend for the performance gaps between APR and DET. Overall, the average performance gap between OPT and APR does not exceed 3.19%, whereas the average performance gap between OPT and DET can be as high as 14.27%.

Since our model is intended to maximize the net expected revenue, we use this performance measure when comparing the different benchmark strategies in Tables 1 and 2. In practice, however, there may be other performance measures of interest related to service levels and seat utilizations. In Tables 4 and 5, we provide additional performance measures for OPT, APR and DET. The first column in these tables shows the problem characteristics. The second, third and fourth columns show the expected revenues obtained, whereas the fifth, sixth and seventh columns show the expected penalty costs incurred by OPT, APR and DET. The difference between the expected revenues and the expected penalty costs in Tables 4 and 5 yields the net expected revenues in Tables 1 and 2. The eighth, ninth and tenth columns show what percent of the reservations that show up at the departure time are denied boarding by OPT, APR and DET. Finally, the eleventh, twelfth and thirteenth columns show what percent of the seats are occupied when the flight leg departs under the policies used by OPT, APR and DET.

Comparing OPT and APR in Tables 4 and 5, we observe that APR obtains lower expected revenues and provides lower seat occupancy than OPT, indicating that APR tends to be less aggressive than OPT when accepting the requests. We do not see a clear ordering between OPT and APR in terms of expected penalty costs and denied reservations. For the test problems with low load factors, while OPT obtains higher expected revenues, APR incurs lower expected penalty costs and this balances the performances of the two benchmark strategies. The expected revenues, expected penalty costs, denied reservations and seat occupancies for DET are significantly higher than those for OPT and APR, indicating that DET is overly aggressive when accepting the requests. The denied reservations for DET can be larger than those for OPT and APR by an order of magnitude. Furthermore, since DET assumes that the show up decisions take on their expected values, the denied reservations for DET are not affected by the penalty cost. In contrast, as γ increases and it becomes more costly to deny boarding to a reservation, we observe a decrease in the denied reservations for both OPT and APR. Thus, DET appears to struggle in finding the right balance between overbooking and letting some of the seats on the flight leg depart unoccupied.

A possible concern for APR is that this benchmark strategy provides a randomized policy and revenue managers may not be comfortable with using such a randomized policy in practice. One way to obtain a deterministic policy for APR is to solve problem (4) after imposing the constraints $p_j \in \{0, 1\}$ for all $j = 1, \ldots, n$ on the decision variables. It turns out that imposing such binary constraints on the decision variables in problem (4) drastically changes the structure of the optimal solution. In particular, Proposition 2 shows that when there are no binary constraints on the decision variables, the optimal solution to problem (4) is such that if we accept the requests for fare class j, then we must also accept the requests for fare class $j-1$. In contrast, when there are binary constraints on the decision variables, it is possible to have an optimal solution to problem (4) where we accept the requests for fare class j, but not the requests fare class $j-1$. In Appendix C, we demonstrate this situation by studying a problem instance with three fare classes such that if we impose binary constraints on the decision variables, then the optimal solution to problem (4) is $(p_1^*, p_2^*, p_3^*) = (0, 1, 1)$. Given that the optimal solution to problem (4) does not have an intuitive structure when we impose binary constraints on the decision variables, we resort to a heuristic modification of APR to obtain a deterministic policy for this benchmark strategy. In particular, this version of APR solves problem (4) without binary constraints to obtain the acceptance probabilities, but it rounds the acceptance probabilities to the nearest integer when deciding which requests to accept. We give the performance of this simple modification of APR in Table 6. The first column in Table 6 shows the problem characteristics. The second column shows the net expected revenue obtained by the original version of APR and it is identical to the third column in Tables 1 and 2, whereas the third column shows the net expected revenue obtained by the heuristic modification of APR with a deterministic policy. The fourth column shows the percent gap between the net expected revenues obtained by the two versions. Table 6 indicates that the two versions of APR obtain essentially the same net expected revenues. Thus, we can indeed use a simple modification of APR to come up with a deterministic policy without any noticeable performance loss.

7.4 Computational Experiments with Random Arrival Rates

In this section, we provide computational experiments on test problems with random arrival rates. As we describe in Section 6.2, using a Poisson process to model the arrivals of the requests may not provide adequate variability for some practical applications and it may be desirable to incorporate more variable demand processes by using arrival rates that are themselves random variables. In this case, we need to work with the net expected revenue function in (17). The experimental setup in this section differs from the one in Section 7.1 only in the arrival rates for the requests for different fare classes. We use the arrival rates in Section 7.1 as a base case and assume that the arrival rates can deviate from the base case by $\mp D\%$. Thus, the arrival rate for the requests for each fare class can take the base case value, $D\%$ more than the base case value and $D\%$ less than the base case value, each occurring with equal probability. We label our test problems by the tuple $(D, q^L, q^H, \gamma, \rho)$, where q^L, q^H, γ and ρ are as in Section 7.1. We use $D \in \{10, 20, 40\}$, $q^L = \{0.7, 0.9\}$, $q^H \in \{0.7, 0.9\}$, $\gamma \in \{2, 4\}$ and $\rho \in \{1.0, 1.2, 1.6\}$, in which case, we obtain 72 test problems. In all of our test problems, we have eight fare classes. The fare associated with the highest fare class is four times the fare associated with the lowest fare class and the fares for the other fare classes are evenly distributed between the highest and lowest fares.

As benchmark strategies, we work with APR and DET, where APR uses the net expected revenue function in (17) and DET solves problem (8) after replacing the arrival rates for the requests by their expected values. When using APR, we maximize the net expected revenue function in (17) by using the cutting plane method that we describe in Appendix B. We use $\epsilon = 10^{-4}$ in the stopping condition of the cutting plane method. In all of our test problems, the cutting plane method never takes more than half a second. Similar to our implementation in Section 7.2, APR and DET periodically refresh the acceptance probabilities over the selling horizon.

The net expected revenue function in (17) is constructed under the assumption that the arrival rates are chosen according to the probability distribution $\{\rho^\ell : \ell = 1, \ldots, L\}$ at the beginning of the selling horizon, but they remain unknown to the decision maker. In contrast, the Hamilton Jacobi Bellman equation in (2) is constructed under the assumption that the arrival rates are known to the decision maker. We use the Hamilton Jacobi Bellman equation in (2) to develop an upper bound on the net expected revenue that can be obtained under the assumption that the arrival rates are unknown to the decision maker. For each possible set of arrival rates, we solve the Hamilton Jacobi Bellman equation in (2) to compute the net expected revenue under a particular set of arrival rates. In this case, if we average these net expected revenues according to the weights $\{\rho^\ell : \ell = 1, \ldots, L\}$, then we obtain the net expected revenue under the assumption that the arrival rates are chosen according to the probability distribution $\{\rho^\ell : \ell = 1, \ldots, L\}$ at the beginning of the selling horizon and they are known to the decision maker. The net expected revenue that we obtain in this fashion is clearly an upper bound on the net expected revenue that we obtain under the assumption that the arrival rates are chosen according to the probability distribution $\{\rho^\ell : \ell = 1, \ldots, L\}$ at the beginning of the selling horizon, but they remain unknown to the decision maker. If the performance of APR turns out to be close to the upper bound, then we can safely state that APR performs well when the arrival rates are random. We caution the reader that if the performance of APR turns out to be distant from the upper bound, then it is not possible to say whether it is because the upper bound is loose or the performance of APR is not satisfactory. Therefore, the upper bound is useful in identifying when APR performs well, but not necessarily in identifying when APR performs poorly.

We give our computational results in Tables 7, 8 and 9. In particular, Tables 7, 8 and 9 respectively focus on the test problems with $D = 10$, $D = 20$ and $D = 40$. The first column in these tables shows the problem characteristics by using the tuple $(D, q^L, q^H, \gamma, \rho)$. The second and third columns show the net expected revenues obtained by APR and DET. The fourth column shows the percent gap between the net expected revenues obtained by APR and DET. The fifth column shows the upper bound on the optimal net expected revenue that we obtain by using the Hamilton Jacobi Bellman equation in (2).

Our results indicate that APR continues to provide significant improvements over DET. The average performance gap between APR and DET is 7.55% and there are test problems where the performance gap between the two benchmark strategies can exceed 20%. We observe two problem parameters that affect the performance gaps between APR and DET. Similar to our observations in Tables 1 and 2, as γ increases and it becomes more costly to deny boarding to a reservation, the performance gap between APR and DET increases. For the test problems with $\gamma = 2$, the average performance gap between APR and DET is 3.63%, whereas for the test problems with $\gamma = 4$, the same average gap is 11.49%. On the other hand, as D increases and the arrival rates become more variable, we observe an increase in the performance gaps between APR and DET. If we focus only on the test problems with $D = 10$, then the performance gap between APR and DET is 4.96%. The same average gap increases respectively to 6.35% and 11.36% when we focus only on the test problems with $D = 20$ and $D = 40$. Overall, our results indicate that APR handles the uncertainty in the arrivals of the requests and the uncertainty in the arrival rates noticeably better than DET. Finally, to get a feel for how the performance gaps vary with D and γ, Table 10 shows the performance gaps between APR and DET averaged over the test problems with a particular value for D and γ. Table 10 indicates that the performance gaps between APR and DET increase when we increase γ or D, keeping the other parameter constant. The average performance gap between APR and DET can be as high as 17.37% for the test problems with large variability in the arrival rates and large penalty cost.

8 Conclusions

In this paper, we developed a tractable model for making overbooking and capacity control decisions over a single flight leg. The fundamental idea is to assume that the requests for each fare class are accepted with a fixed acceptance probability and to find a set of acceptance probabilities that maximize the net expected revenue. We showed that the optimal acceptance probabilities can be computed in a tractable fashion. Furthermore, the optimal acceptance probabilities randomize the acceptance decisions for at most one fare class and provide an open loop policy with a computable performance bound. Computational experiments indicated that our policies perform quite well. An important avenue for future investigation is to try to incorporate the customer choice process into our model, where each arriving customer observes the subset of fare classes that are left open for purchase and makes a choice among them. In this case, the challenge would be to find the probability with which each subset of fare classes should be left open for purchase.

References

Aydin, N., Birbil, S. I., Frenk, J. B. G. and Noyan, N. (2010), Single-leg airline revenue management with overbooking, Technical report, Sabanci University, Istanbul, Turkey.

Ball, M. O. and Queyranne, M. (2009), 'Toward robust revenue management: Competitive analysis of online booking', *Operations Research* **57**(4), 950–963.

Beckmann, M. J. (1958), 'Decision and team problems in airline reservations', *Econometrica* **26**(1), 134–145.

Belobaba, P. P. (1987), Air Travel Demand and Airline Seat Inventory Control, PhD thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, Cambridge, MA.

Belobaba, P. P. (2006), 'Flight overbooking: Models and practice', Unpublished lecture notes, Massachusetts Institute of Technology, Cambridge, MA.

Bertsimas, D. and Popescu, I. (2003), 'Revenue management in a dynamic network environment', *Transportation Science* **37**, 257–277.

Boyd, S. and Vandenberghe, L. (2005), *Convex Optimization*, Cambridge University Press, Cambridge, UK.

Bremaud, P. (1980), *Point Processes and Queues: Martingale Dynamics*, Springer-Verlang, New York.

Brumelle, S. L. and McGill, J. I. (1993), 'Airline seat allocation with multiple nested fare classes', *Operations Research* **41**, 127–137.

Chatwin, R. E. (1998), 'Multiperiod airline overbooking with a single fare class', *Operations Research* **46**(6), 805–819.

Chatwin, R. E. (1999), 'Continuous-time airline overbooking with time-dependent fares and refunds', *Transportation Science* **33**(2), 182–191.

Cooper, W. L. (2002), 'Asymptotic behavior of an allocation policy for revenue management', *Operations Research* **50**(4), 720–727.

Cormen, T. H., Leiserson, C. E., Rivest, R. L. and Stein, C. (2009), *Introduction to Algorithms*, The MIT Press, Cambridge. MA.

Coughlan, J. (1999), 'Airline overbooking in the multi-class case', *The Journal of the Operational Research Society* **50**(11), 1098–1103.

Erdelyi, A. and Topaloglu, H. (2009), 'Separable approximations for joint capacity control and overbooking decisions in network revenue management', *Journal of Revenue and Pricing Management* **8**(1), 3–20.

Erdelyi, A. and Topaloglu, H. (2010), 'A dynamic programming decomposition method for making overbooking decisions over an airline network', *INFORMS Journal on Computing* **22**(3), 443–456.

Gallego, G. and van Ryzin, G. (1994), 'Optimal dynamic pricing of inventories with stochastic demand over finite horizons', *Management Science* **40**(8), 999–1020.

Gallego, G. and van Ryzin, G. (1997), 'A multiproduct dynamic pricing problem and its applications to yield management', *Operations Research* **45**(1), 24–41.

Karaesmen, I. and van Ryzin, G. (2004*a*), Coordinating overbooking and capacity control decisions on a network, Technical report, Columbia Business School.

Karaesmen, I. and van Ryzin, G. (2004*b*), 'Overbooking with substitutable inventory classes', *Operations Research* **52**(1), 83–104.

Kunnumkal, S. and Topaloglu, H. (2008), 'A tractable revenue management model for capacity allocation and overbooking over an airline network', *Flexible Services and Manufacturing Journal* **20**(3), 125–147.

Lan, Y., Ball, M. O. and Karaesmen, I. Z. (2011), 'Regret in overbooking and fare-class allocation for single leg', *Manufacting & Service Operations Management* **13**(2), 194–208.

Levi, R. and Radovanovic, A. (2010), 'Provably near-optimal LP-based policies for revenue management in systems with reusable resources', *Operations Research* **58**(2), 503–507.

Littlewood, K. (1972), Forecasting and control of passengers, *in* 'Proceedings of the 12th AGIFORS', New York, pp. 95–117.

Rothstein, M. (1971), 'An airline overbooking model', *Transportation Science* **5**(2), 180–192.

Rothstein, M. (1974), 'Hotel overbooking as a Markovian sequential decision process', *Decision Sciences* **5**(3), 389–404.

Ruszczynski, A. (2006), *Nonlinear Optimization*, Princeton University Press, Princeton, New Jersey.

Subramanian, J., Stidham, S. and Lautenbacher, C. J. (1999), 'Airline yield management with overbooking, cancellations and no-shows', *Transportation Science* **33**(2), 147–167.

Talluri, K. T. and van Ryzin, G. J. (2005), *The Theory and Practice of Revenue Management*, Springer, New York, NY.

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

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

Thompson, H. R. (1961), 'Statistical problems in airline reservation control', *Journal of the Operational Research Society* **12**(3), 167–185.

Topaloglu, H. (2009), 'On the asymptotic optimality of the randomized linear program for network revenue management', *European Journal of Operational Research* **197**, 884–896.

van Slyke, R. and Young, Y. (2000), ‘Finite horizon stochastic knapsacks with applications to yield management’, *Operations Research* **48**(1), 155–172.

Williamson, D. P. and Shmoys, D. B. (2011), *The Design of Approximation Algorithms*, Cambridge University Press, New York.

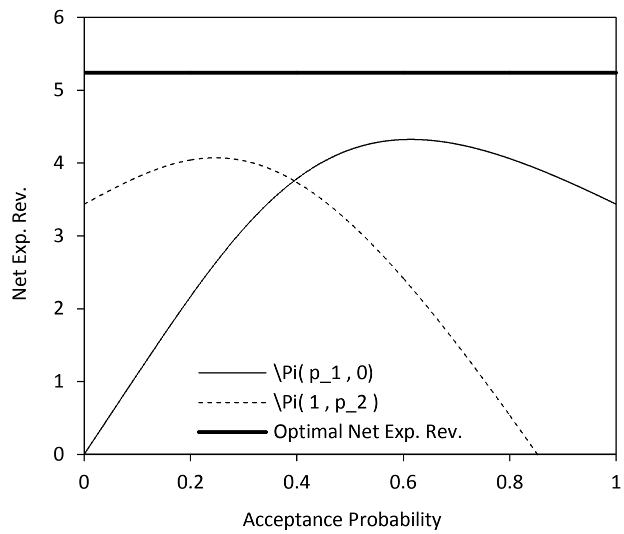

Figure 1: Net expected revenue function $\Pi(p_1, 0)$ for all $p_1 \in [0, 1]$ and $\Pi(1, p_2)$ for all $p_2 \in [0, 1]$.

Test Problem	Net Expected Revenue			OPT vs.	
$(\kappa, q^L, q^H, \gamma, \rho)$	OPT	APR	DET	APR	DET
$(4, 0.7, 0.7, 2, 1.0)$	20,437	20,368	19,936	0.34	2.45
$(4, 0.7, 0.7, 2, 1.2)$	18,719	18,641	18,150	0.42	3.04
$(4, 0.7, 0.7, 2, 1.6)$	15,604	15,392	15,161	1.35	2.83
$(4, 0.7, 0.7, 4, 1.0)$	20,215	20,062	18,604	0.75	7.97
$(4, 0.7, 0.7, 4, 1.2)$	18,398	18,196	16,551	1.10	10.04
$(4, 0.7, 0.7, 4, 1.6)$	15,222	14,816	13,546	2.67	11.01
$(4, 0.7, 0.9, 2, 1.0)$	20,434	20,264	19,831	0.83	2.95
$(4, 0.7, 0.9, 2, 1.2)$	18,604	18,276	17,894	1.77	3.81
$(4, 0.7, 0.9, 2, 1.6)$	15,282	14,861	14,567	2.75	4.68
$(4, 0.7, 0.9, 4, 1.0)$	20,233	19,896	18,367	1.66	9.22
$(4, 0.7, 0.9, 4, 1.2)$	18,320	17,760	16,138	3.05	11.91
$(4, 0.7, 0.9, 4, 1.6)$	15,008	14,346	12,813	4.41	14.62
$(4, 0.9, 0.7, 2, 1.0)$	20,054	19,950	19,350	0.52	3.51
$(4, 0.9, 0.7, 2, 1.2)$	18,461	18,343	17,765	0.64	3.77
$(4, 0.9, 0.7, 2, 1.6)$	15,726	15,558	15,122	1.07	3.84
$(4, 0.9, 0.7, 4, 1.0)$	19,909	19,629	18,060	1.41	9.29
$(4, 0.9, 0.7, 4, 1.2)$	18,279	17,999	16,315	1.53	10.75
$(4, 0.9, 0.7, 4, 1.6)$	15,481	15,112	13,589	2.38	12.22
$(4, 0.9, 0.9, 2, 1.0)$	20,685	20,486	20,024	0.96	3.20
$(4, 0.9, 0.9, 2, 1.2)$	18,928	18,616	18,097	1.64	4.39
$(4, 0.9, 0.9, 2, 1.6)$	15,999	15,537	15,191	2.89	5.05
$(4, 0.9, 0.9, 4, 1.0)$	20,570	20,226	18,752	1.67	8.84
$(4, 0.9, 0.9, 4, 1.2)$	18,785	18,260	16,513	2.80	12.10
$(4, 0.9, 0.9, 4, 1.6)$	15,830	15,107	13,474	4.57	14.88

Table 1: Computational results for the test problems with $\kappa = 4$.

Test Problem $(\kappa, q^L, q^H, \gamma, \rho)$	Net Expected Revenue OPT	APR	DET	OPT vs. APR	DET
$(8, 0.7, 0.7, 2, 1.0)$	35,945	35,918	34,270	0.08	4.66
$(8, 0.7, 0.7, 2, 1.2)$	33,719	33,609	32,013	0.33	5.06
$(8, 0.7, 0.7, 2, 1.6)$	28,944	28,622	27,788	1.11	4.00
$(8, 0.7, 0.7, 4, 1.0)$	35,736	35,566	31,606	0.48	11.56
$(8, 0.7, 0.7, 4, 1.2)$	33,321	33,082	28,815	0.72	13.52
$(8, 0.7, 0.7, 4, 1.6)$	28,376	27,736	24,557	2.25	13.46
$(8, 0.7, 0.9, 2, 1.0)$	35,883	35,763	34,054	0.33	5.10
$(8, 0.7, 0.9, 2, 1.2)$	33,452	33,064	31,621	1.16	5.48
$(8, 0.7, 0.9, 2, 1.6)$	28,352	27,756	26,790	2.10	5.51
$(8, 0.7, 0.9, 4, 1.0)$	35,701	35,336	31,126	1.02	12.82
$(8, 0.7, 0.9, 4, 1.2)$	33,101	32,381	28,107	2.18	15.09
$(8, 0.7, 0.9, 4, 1.6)$	27,879	26,909	23,283	3.48	16.49
$(8, 0.9, 0.7, 2, 1.0)$	35,616	35,448	33,598	0.47	5.66
$(8, 0.9, 0.7, 2, 1.2)$	33,438	33,303	31,602	0.41	5.49
$(8, 0.9, 0.7, 2, 1.6)$	29,222	28,959	27,746	0.90	5.05
$(8, 0.9, 0.7, 4, 1.0)$	35,409	35,040	31,017	1.04	12.40
$(8, 0.9, 0.7, 4, 1.2)$	33,149	32,799	28,701	1.05	13.42
$(8, 0.9, 0.7, 4, 1.6)$	28,830	28,269	24,680	1.95	14.39
$(8, 0.9, 0.9, 2, 1.0)$	36,239	36,089	34,422	0.41	5.02
$(8, 0.9, 0.9, 2, 1.2)$	34,060	33,649	31,966	1.21	6.15
$(8, 0.9, 0.9, 2, 1.6)$	29,635	28,915	27,793	2.43	6.22
$(8, 0.9, 0.9, 4, 1.0)$	36,120	35,775	31,878	0.95	11.74
$(8, 0.9, 0.9, 4, 1.2)$	33,863	33,172	28,798	2.04	14.96
$(8, 0.9, 0.9, 4, 1.6)$	29,377	28,247	24,359	3.85	17.08

Table 2: Computational results for the test problems with $\kappa = 8$.

(γ, ρ)	OPT vs. APR	OPT vs. DET	APR vs. DET
$(2, 1.0)$	0.49	4.07	3.59
$(2, 1.2)$	0.95	4.65	3.74
$(2, 1.6)$	1.83	4.65	2.87
$(4, 1.0)$	1.12	10.48	9.46
$(4, 1.2)$	1.81	12.72	11.11
$(4, 1.6)$	3.19	14.27	11.44

Table 3: Performance gaps between OPT, APR and DET averaged over the test problems with a particular value for (γ, ρ).

Test Problem	Expected Revenue			Expected Cost			Denied Boarding			Seat Occupancy		
$(\kappa, q^L, q^H, \gamma, \rho)$	OPT	APR	DET	OPT	APR	DET	OPT	APR	DET	OPT	APR	DET
$(4, 0.7, 0.7, 2, 1.0)$	20,975	20,697	21,268	538	329	1,332	0.91	0.55	2.19	92.27	88.73	94.18
$(4, 0.7, 0.7, 2, 1.2)$	19,400	19,196	19,749	681	555	1,599	1.36	1.09	3.07	92.98	90.01	95.12
$(4, 0.7, 0.7, 2, 1.6)$	16,368	16,324	16,777	764	932	1,615	2.03	2.37	4.07	93.67	90.80	94.70
$(4, 0.7, 0.7, 4, 1.0)$	20,624	20,273	21,268	410	211	2,664	0.35	0.18	2.19	90.09	85.22	94.18
$(4, 0.7, 0.7, 4, 1.2)$	18,942	18,623	19,749	544	427	3,198	0.55	0.43	3.07	90.52	85.90	95.12
$(4, 0.7, 0.7, 4, 1.6)$	15,777	15,525	16,777	555	709	3,230	0.75	0.92	4.07	90.55	85.52	94.70
$(4, 0.7, 0.9, 2, 1.0)$	20,857	20,656	21,295	423	392	1,464	0.66	0.60	2.19	93.07	89.35	94.44
$(4, 0.7, 0.9, 2, 1.2)$	19,123	18,994	19,651	519	718	1,757	0.95	1.27	3.06	94.12	90.70	95.28
$(4, 0.7, 0.9, 2, 1.6)$	15,860	15,669	16,321	578	808	1,754	1.41	1.89	3.96	95.05	91.29	94.76
$(4, 0.7, 0.9, 4, 1.0)$	20,559	20,204	21,295	326	307	2,928	0.26	0.24	2.19	91.43	86.01	94.44
$(4, 0.7, 0.9, 4, 1.2)$	18,757	18,360	19,651	437	600	3,514	0.41	0.54	3.06	92.23	86.70	95.28
$(4, 0.7, 0.9, 4, 1.6)$	15,421	14,724	16,321	413	378	3,507	0.51	0.45	3.96	92.86	85.67	94.76
$(4, 0.9, 0.7, 2, 1.0)$	20,463	20,141	20,641	410	191	1,290	0.64	0.29	1.95	93.62	88.79	95.54
$(4, 0.9, 0.7, 2, 1.2)$	18,966	18,692	19,216	505	349	1,450	0.93	0.63	2.57	93.95	89.73	95.90
$(4, 0.9, 0.7, 2, 1.6)$	16,290	16,152	16,655	564	594	1,533	1.37	1.38	3.52	93.85	89.88	95.40
$(4, 0.9, 0.7, 4, 1.0)$	20,213	19,753	20,641	304	125	2,581	0.24	0.10	1.95	91.92	85.27	95.54
$(4, 0.9, 0.7, 4, 1.2)$	18,628	18,220	19,216	349	221	2,901	0.32	0.20	2.57	92.09	85.66	95.90
$(4, 0.9, 0.7, 4, 1.6)$	15,870	15,551	16,655	389	438	3,066	0.48	0.52	3.52	91.41	84.93	95.40
$(4, 0.9, 0.9, 2, 1.0)$	20,976	20,692	21,296	291	206	1,272	0.40	0.27	1.65	93.77	88.83	94.82
$(4, 0.9, 0.9, 2, 1.2)$	19,255	19,044	19,681	327	427	1,584	0.53	0.67	2.43	94.76	90.05	95.93
$(4, 0.9, 0.9, 2, 1.6)$	16,371	16,278	16,908	372	741	1,717	0.80	1.49	3.40	95.25	90.43	95.55
$(4, 0.9, 0.9, 4, 1.0)$	20,805	20,343	21,296	235	117	2,544	0.16	0.08	1.65	92.84	85.78	94.82
$(4, 0.9, 0.9, 4, 1.2)$	19,027	18,567	19,681	242	307	3,168	0.20	0.25	2.43	93.56	86.43	95.93
$(4, 0.9, 0.9, 4, 1.6)$	16,137	15,636	16,908	307	530	3,434	0.33	0.54	3.40	93.86	86.01	95.55

Table 4: Detailed performance measures for OPT, APR and DET for the test problems with $\kappa = 4$.

Test Problem	Expected Revenue			Expected Cost			Denied Boarding			Seat Occupancy		
$(\kappa, q^L, q^H, \gamma, \rho)$	OPT	APR	DET	OPT	APR	DET	OPT	APR	DET	OPT	APR	DET
$(8, 0.7, 0.7, 2, 1.0)$	36,534	36,179	36,934	589	261	2,664	0.50	0.22	2.19	89.86	85.56	94.18
$(8, 0.7, 0.7, 2, 1.2)$	34,689	34,271	35,211	970	662	3,198	0.97	0.66	3.07	91.23	87.43	95.12
$(8, 0.7, 0.7, 2, 1.6)$	30,167	29,988	31,018	1,222	1,366	3,230	1.63	1.76	4.07	92.38	88.90	94.70
$(8, 0.7, 0.7, 4, 1.0)$	36,191	35,755	36,934	454	189	5,328	0.20	0.08	2.19	87.70	82.55	94.18
$(8, 0.7, 0.7, 4, 1.2)$	34,070	33,562	35,211	749	480	6,397	0.38	0.24	3.07	88.95	83.70	95.12
$(8, 0.7, 0.7, 4, 1.6)$	29,256	28,789	31,018	880	1,053	6,461	0.60	0.69	4.07	89.40	83.83	94.70
$(8, 0.7, 0.9, 2, 1.0)$	36,384	36,113	36,982	501	350	2,928	0.39	0.27	2.19	90.88	86.36	94.44
$(8, 0.7, 0.9, 2, 1.2)$	34,244	33,960	35,134	792	896	3,514	0.73	0.80	3.06	92.52	88.31	95.28
$(8, 0.7, 0.9, 2, 1.6)$	29,291	28,838	30,297	939	1,082	3,507	1.15	1.28	3.96	94.17	89.46	94.76
$(8, 0.7, 0.9, 4, 1.0)$	36,066	35,662	36,982	365	326	5,856	0.14	0.13	2.19	89.19	83.57	94.44
$(8, 0.7, 0.9, 4, 1.2)$	33,741	33,162	35,134	640	781	7,027	0.30	0.35	3.06	90.81	84.65	95.28
$(8, 0.7, 0.9, 4, 1.6)$	28,587	27,508	30,297	707	598	7,014	0.44	0.36	3.96	91.94	84.53	94.76
$(8, 0.9, 0.7, 2, 1.0)$	36,075	35,621	36,179	459	173	2,581	0.36	0.13	1.95	91.49	86.04	95.54
$(8, 0.9, 0.7, 2, 1.2)$	34,160	33,674	34,503	722	371	2,901	0.67	0.34	2.57	92.55	87.39	95.90
$(8, 0.9, 0.7, 2, 1.6)$	30,144	29,813	30,812	922	854	3,066	1.11	1.00	3.52	92.76	88.02	95.40
$(8, 0.9, 0.7, 4, 1.0)$	35,806	35,165	36,179	397	125	5,162	0.16	0.05	1.95	89.96	83.04	95.54
$(8, 0.9, 0.7, 4, 1.2)$	33,700	33,046	34,503	550	246	5,802	0.26	0.11	2.57	90.70	83.54	95.90
$(8, 0.9, 0.7, 4, 1.6)$	29,477	28,902	30,812	646	634	6,131	0.40	0.38	3.52	90.30	83.28	95.40
$(8, 0.9, 0.9, 2, 1.0)$	36,593	36,233	36,966	354	144	2,544	0.24	0.10	1.65	92.21	86.05	94.82
$(8, 0.9, 0.9, 2, 1.2)$	34,522	34,131	35,134	462	482	3,168	0.38	0.38	2.43	93.34	87.75	95.93
$(8, 0.9, 0.9, 2, 1.6)$	30,235	29,994	31,226	600	1,078	3,434	0.64	1.09	3.40	94.28	88.74	95.55
$(8, 0.9, 0.9, 4, 1.0)$	36,405	35,877	36,966	285	102	5,088	0.10	0.03	1.65	91.21	83.32	94.82
$(8, 0.9, 0.9, 4, 1.2)$	34,225	33,536	35,134	362	365	6,336	0.15	0.15	2.43	92.25	84.52	95.93
$(8, 0.9, 0.9, 4, 1.6)$	29,876	29,027	31,226	499	781	6,867	0.27	0.40	3.40	92.96	84.54	95.55

Table 5: Detailed performance measures for OPT, APR and DET for the test problems with $\kappa = 8$.

Test Problem $(\kappa, q^L, q^H, \gamma, \rho)$	Net Expected Revenue Rand.	Net Expected Revenue Deter.	Rand. vs. Deter.
$(4, 0.7, 0.7, 2, 1.0)$	20,368	20,364	0.02
$(4, 0.7, 0.7, 2, 1.2)$	18,641	18,650	-0.05
$(4, 0.7, 0.7, 2, 1.6)$	15,392	15,391	0.01
$(4, 0.7, 0.7, 4, 1.0)$	20,062	20,025	0.18
$(4, 0.7, 0.7, 4, 1.2)$	18,196	18,189	0.04
$(4, 0.7, 0.7, 4, 1.6)$	14,816	14,811	0.03
$(4, 0.7, 0.9, 2, 1.0)$	20,264	20,266	-0.01
$(4, 0.7, 0.9, 2, 1.2)$	18,276	18,281	-0.03
$(4, 0.7, 0.9, 2, 1.6)$	14,861	14,861	0.00
$(4, 0.7, 0.9, 4, 1.0)$	19,896	19,872	0.12
$(4, 0.7, 0.9, 4, 1.2)$	17,760	17,765	-0.03
$(4, 0.7, 0.9, 4, 1.6)$	14,346	14,329	0.12
$(4, 0.9, 0.7, 2, 1.0)$	19,950	19,926	0.12
$(4, 0.9, 0.7, 2, 1.2)$	18,343	18,354	-0.06
$(4, 0.9, 0.7, 2, 1.6)$	15,558	15,551	0.05
$(4, 0.9, 0.7, 4, 1.0)$	19,629	19,631	-0.01
$(4, 0.9, 0.7, 4, 1.2)$	17,999	17,981	0.10
$(4, 0.9, 0.7, 4, 1.6)$	15,112	15,076	0.24
$(4, 0.9, 0.9, 2, 1.0)$	20,486	20,460	0.13
$(4, 0.9, 0.9, 2, 1.2)$	18,616	18,635	-0.10
$(4, 0.9, 0.9, 2, 1.6)$	15,537	15,536	0.00
$(4, 0.9, 0.9, 4, 1.0)$	20,226	20,169	0.28
$(4, 0.9, 0.9, 4, 1.2)$	18,260	18,256	0.02
$(4, 0.9, 0.9, 4, 1.6)$	15,107	15,093	0.09

Test Problem $(\kappa, q^L, q^H, \gamma, \rho)$	Net Expected Revenue Rand.	Net Expected Revenue Deter.	Rand. vs. Deter.
$(8, 0.7, 0.7, 2, 1.0)$	35,918	35,905	0.04
$(8, 0.7, 0.7, 2, 1.2)$	33,609	33,642	-0.10
$(8, 0.7, 0.7, 2, 1.6)$	28,622	28,584	0.13
$(8, 0.7, 0.7, 4, 1.0)$	35,566	35,538	0.08
$(8, 0.7, 0.7, 4, 1.2)$	33,082	33,053	0.09
$(8, 0.7, 0.7, 4, 1.6)$	27,736	27,700	0.13
$(8, 0.7, 0.9, 2, 1.0)$	35,763	35,746	0.05
$(8, 0.7, 0.9, 2, 1.2)$	33,064	33,040	0.07
$(8, 0.7, 0.9, 2, 1.6)$	27,756	27,739	0.06
$(8, 0.7, 0.9, 4, 1.0)$	35,336	35,338	-0.01
$(8, 0.7, 0.9, 4, 1.2)$	32,381	32,409	-0.09
$(8, 0.7, 0.9, 4, 1.6)$	26,909	26,928	-0.07
$(8, 0.9, 0.7, 2, 1.0)$	35,448	35,472	-0.07
$(8, 0.9, 0.7, 2, 1.2)$	33,303	33,246	0.17
$(8, 0.9, 0.7, 2, 1.6)$	28,959	28,928	0.11
$(8, 0.9, 0.7, 4, 1.0)$	35,040	35,077	-0.11
$(8, 0.9, 0.7, 4, 1.2)$	32,799	32,779	0.06
$(8, 0.9, 0.7, 4, 1.6)$	28,269	28,231	0.13
$(8, 0.9, 0.9, 2, 1.0)$	36,089	36,058	0.09
$(8, 0.9, 0.9, 2, 1.2)$	33,649	33,648	0.00
$(8, 0.9, 0.9, 2, 1.6)$	28,915	28,894	0.07
$(8, 0.9, 0.9, 4, 1.0)$	35,775	35,764	0.03
$(8, 0.9, 0.9, 4, 1.2)$	33,172	33,165	0.02
$(8, 0.9, 0.9, 4, 1.6)$	28,247	28,175	0.25

Table 6: Net expected revenues obtained by APR by using a randomized and a deterministic policy.

Test Problem $(D, q^L, q^H, \gamma, \rho)$	Net Expected Revenue APR	Net Expected Revenue DET	APR vs. DET	Upper Bound
$(10, 0.7, 0.7, 2, 1.0)$	20,138	19,653	2.41	20,296
$(10, 0.7, 0.7, 2, 1.2)$	18,459	18,039	2.27	18,611
$(10, 0.7, 0.7, 2, 1.6)$	15,337	15,087	1.63	15,516
$(10, 0.7, 0.7, 4, 1.0)$	19,748	18,345	7.11	20,031
$(10, 0.7, 0.7, 4, 1.2)$	17,966	16,580	7.72	18,293
$(10, 0.7, 0.7, 4, 1.6)$	14,672	13,617	7.18	15,090
$(10, 0.7, 0.9, 2, 1.0)$	20,422	19,921	2.45	20,674
$(10, 0.7, 0.9, 2, 1.2)$	18,873	18,303	3.02	19,126
$(10, 0.7, 0.9, 2, 1.6)$	16,029	15,650	2.36	16,240
$(10, 0.7, 0.9, 4, 1.0)$	20,091	18,827	6.29	20,542
$(10, 0.7, 0.9, 4, 1.2)$	18,419	17,043	7.47	18,952
$(10, 0.7, 0.9, 4, 1.6)$	15,492	14,297	7.72	15,991
$(10, 0.9, 0.7, 2, 1.0)$	20,039	19,537	2.50	20,263
$(10, 0.9, 0.7, 2, 1.2)$	18,210	17,816	2.16	18,470
$(10, 0.9, 0.7, 2, 1.6)$	14,834	14,531	2.04	15,219
$(10, 0.9, 0.7, 4, 1.0)$	19,576	18,069	7.70	20,066
$(10, 0.9, 0.7, 4, 1.2)$	17,585	16,224	7.74	18,248
$(10, 0.9, 0.7, 4, 1.6)$	14,133	12,945	8.41	14,940
$(10, 0.9, 0.9, 2, 1.0)$	20,240	19,735	2.49	20,551
$(10, 0.9, 0.9, 2, 1.2)$	18,491	17,977	2.78	18,852
$(10, 0.9, 0.9, 2, 1.6)$	15,555	15,154	2.58	15,958
$(10, 0.9, 0.9, 4, 1.0)$	19,828	18,484	6.78	20,439
$(10, 0.9, 0.9, 4, 1.2)$	17,925	16,523	7.82	18,720
$(10, 0.9, 0.9, 4, 1.6)$	14,867	13,617	8.40	15,798

Table 7: Computational results for the test problems with $D = 10$.

Test Problem $(D, q^L, q^H, \gamma, \rho)$	Net Expected Revenue APR	DET	APR vs. DET	Upper Bound
$(20, 0.7, 0.7, 2, 1.0)$	19,836	19,176	3.33	20,122
$(20, 0.7, 0.7, 2, 1.2)$	18,180	17,674	2.78	18,554
$(20, 0.7, 0.7, 2, 1.6)$	15,138	14,835	2.00	15,502
$(20, 0.7, 0.7, 4, 1.0)$	19,427	17,693	8.93	19,934
$(20, 0.7, 0.7, 4, 1.2)$	17,671	16,043	9.21	18,302
$(20, 0.7, 0.7, 4, 1.6)$	14,426	13,212	8.41	15,154
$(20, 0.7, 0.9, 2, 1.0)$	20,136	19,346	3.92	20,462
$(20, 0.7, 0.9, 2, 1.2)$	18,564	17,876	3.71	19,014
$(20, 0.7, 0.9, 2, 1.6)$	15,774	15,359	2.63	16,208
$(20, 0.7, 0.9, 4, 1.0)$	19,793	18,058	8.77	20,342
$(20, 0.7, 0.9, 4, 1.2)$	18,139	16,418	9.49	18,862
$(20, 0.7, 0.9, 4, 1.6)$	15,234	13,845	9.12	15,997
$(20, 0.9, 0.7, 2, 1.0)$	19,703	19,036	3.38	20,076
$(20, 0.9, 0.7, 2, 1.2)$	17,860	17,391	2.62	18,419
$(20, 0.9, 0.7, 2, 1.6)$	14,560	14,208	2.42	15,183
$(20, 0.9, 0.7, 4, 1.0)$	19,159	17,342	9.48	19,891
$(20, 0.9, 0.7, 4, 1.2)$	17,191	15,541	9.60	18,158
$(20, 0.9, 0.7, 4, 1.6)$	13,872	12,385	10.72	14,914
$(20, 0.9, 0.9, 2, 1.0)$	19,916	19,090	4.15	20,334
$(20, 0.9, 0.9, 2, 1.2)$	18,153	17,521	3.48	18,755
$(20, 0.9, 0.9, 2, 1.6)$	15,263	14,766	3.25	15,916
$(20, 0.9, 0.9, 4, 1.0)$	19,463	17,504	10.06	20,236
$(20, 0.9, 0.9, 4, 1.2)$	17,581	15,801	10.13	18,624
$(20, 0.9, 0.9, 4, 1.6)$	14,506	12,938	10.81	15,747

Table 8: Computational results for the test problems with $D = 20$.

Test Problem $(D, q^L, q^H, \gamma, \rho)$	Net Expected Revenue APR	DET	APR vs. DET	Upper Bound
$(40, 0.7, 0.7, 2, 1.0)$	18,480	17,535	5.11	19,058
$(40, 0.7, 0.7, 2, 1.2)$	16,869	16,193	4.01	17,665
$(40, 0.7, 0.7, 2, 1.6)$	13,895	13,504	2.81	15,003
$(40, 0.7, 0.7, 4, 1.0)$	18,000	15,460	14.11	18,846
$(40, 0.7, 0.7, 4, 1.2)$	16,203	13,917	14.11	17,382
$(40, 0.7, 0.7, 4, 1.6)$	12,987	11,091	14.60	14,614
$(40, 0.7, 0.9, 2, 1.0)$	18,924	17,629	6.85	19,508
$(40, 0.7, 0.9, 2, 1.2)$	17,342	16,302	6.00	18,185
$(40, 0.7, 0.9, 2, 1.6)$	14,623	13,980	4.39	15,761
$(40, 0.7, 0.9, 4, 1.0)$	18,564	15,778	15.01	19,369
$(40, 0.7, 0.9, 4, 1.2)$	16,882	14,174	16.04	18,008
$(40, 0.7, 0.9, 4, 1.6)$	13,989	11,692	16.42	15,520
$(40, 0.9, 0.7, 2, 1.0)$	18,179	17,143	5.70	18,921
$(40, 0.9, 0.7, 2, 1.2)$	16,348	15,663	4.19	17,398
$(40, 0.9, 0.7, 2, 1.6)$	13,348	12,648	5.24	14,654
$(40, 0.9, 0.7, 4, 1.0)$	17,602	14,493	17.66	18,741
$(40, 0.9, 0.7, 4, 1.2)$	15,520	12,872	17.06	17,181
$(40, 0.9, 0.7, 4, 1.6)$	12,676	9,790	22.77	14,404
$(40, 0.9, 0.9, 2, 1.0)$	18,567	17,169	7.53	19,292
$(40, 0.9, 0.9, 2, 1.2)$	16,781	15,667	6.64	17,835
$(40, 0.9, 0.9, 2, 1.6)$	13,907	13,113	5.71	15,415
$(40, 0.9, 0.9, 4, 1.0)$	18,079	14,696	18.72	19,191
$(40, 0.9, 0.9, 4, 1.2)$	16,130	12,912	19.95	17,724
$(40, 0.9, 0.9, 4, 1.6)$	13,054	10,185	21.98	15,263

Table 9: Computational results for the test problems with $D = 40$.

(D, γ)	APR vs. DET
$(10, 2)$	2.39
$(10, 4)$	7.53
$(20, 2)$	3.14
$(20, 4)$	9.56
$(40, 2)$	5.35
$(40, 4)$	17.37

Table 10: Performance gaps between APR and DET averaged over the test problems with a particular value for (D, γ).

A Appendix: Omitted Proofs

The proof of Lemma 1 is based on a more general result given in Lemma 6 below, which becomes useful several times in the paper. Aydin et al. (2010) show the result in Lemma 6 by using a Bernoulli selection argument. We provide an alternative proof for completeness.

Lemma 6 *If* $F(\cdot) : \Re \to \Re$ *is a convex function with bounded directional derivatives, then the expectation* $\mathbb{E}\big\{F(\mathsf{Pois}(\alpha))\big\}$ *is a differentiable and convex function of* α *and the derivative of* $\mathbb{E}\big\{F(\mathsf{Pois}(\alpha))\big\}$ *with respect to* α *is given by* $\mathbb{E}\big\{F(\mathsf{Pois}(\alpha)+1) - F(\mathsf{Pois}(\alpha))\big\}$.

Proof We note that for any $\epsilon > 0$, $\mathsf{Pois}(\alpha+\epsilon)$ has the same distribution as $\mathsf{Pois}(\alpha) + \mathsf{Pois}(\epsilon)$, where $\mathsf{Pois}(\alpha)$ and $\mathsf{Pois}(\epsilon)$ are independent Poisson random variables respectively with means α and ϵ. In this case, if we use B to denote the bound on the directional derivatives of $F(\cdot)$ and condition on $\mathsf{Pois}(\alpha)$, then we obtain

$$\begin{aligned}\mathbb{E}\big\{F(\mathsf{Pois}(\alpha+\epsilon)) - F(\mathsf{Pois}(\alpha)) \,|\, \mathsf{Pois}(\alpha) = l\big\} &= \mathbb{E}\big\{F(l + \mathsf{Pois}(\epsilon)) - F(l)\big\} \\ = \sum_{k=1}^{\infty} [F(l+k) - F(l)] \, \frac{e^{-\epsilon}\,\epsilon^k}{k!} &\le [F(l+1) - F(l)]\, e^{-\epsilon}\,\epsilon + \sum_{k=2}^{\infty} Bk \, \frac{e^{-\epsilon}\,\epsilon^k}{k!} \\ &= [F(l+1) - F(l)]\, e^{-\epsilon}\,\epsilon + B\epsilon\,(1 - e^{-\epsilon}),\end{aligned}$$

where the last equality follows from the fact that $\sum_{k=2}^{\infty} k\, \frac{e^{-\epsilon}\,\epsilon^k}{k!} = \epsilon \sum_{k=1}^{\infty} \frac{e^{-\epsilon}\,\epsilon^k}{k!} = \epsilon\,(1 - e^{-\epsilon})$. Dividing the chain of inequalities above by ϵ and taking expectations, we obtain $\mathbb{E}\big\{F(\mathsf{Pois}(\alpha+\epsilon)) - F(\mathsf{Pois}(\alpha))\big\}/\epsilon \le \mathbb{E}\big\{F(\mathsf{Pois}(\alpha)+1) - F(\mathsf{Pois}(\alpha))\big\}\, e^{-\epsilon} + B\,(1 - e^{-\epsilon})$. Following the same argument, it is possible to show that $\mathbb{E}\big\{F(\mathsf{Pois}(\alpha+\epsilon)) - F(\mathsf{Pois}(\alpha))\big\}/\epsilon \ge \mathbb{E}\big\{F(\mathsf{Pois}(\alpha)+1) - F(\mathsf{Pois}(\alpha))\big\}\, e^{-\epsilon} - B\,(1 - e^{-\epsilon})$. In this case, since $\lim_{\epsilon \downarrow 0} B\,(1 - e^{-\epsilon}) = 0$, we obtain

$$\lim_{\epsilon \downarrow 0} \frac{\mathbb{E}\big\{F(\mathsf{Pois}(\alpha+\epsilon)) - F(\mathsf{Pois}(\alpha))\big\}}{\epsilon} = \mathbb{E}\big\{F(\mathsf{Pois}(\alpha)+1) - F(\mathsf{Pois}(\alpha))\big\}.$$

One can use a similar argument to show that the limit on the left side above as $\epsilon \uparrow 0$ is equal to the expression on the right side. This establishes that $\mathbb{E}\big\{F(\mathsf{Pois}(\alpha))\big\}$ is a differentiable function of α and its derivative with respect to α is given by $\mathbb{E}\big\{F(\mathsf{Pois}(\alpha)+1) - F(\mathsf{Pois}(\alpha))\big\}$. To see that $\mathbb{E}\big\{F(\mathsf{Pois}(\alpha)\big\}$ is a convex function of α, for any $\epsilon > 0$, the convexity of $F(\cdot)$ implies that

$$\begin{aligned}F(\mathsf{Pois}(\alpha+\epsilon)+1) - F(\mathsf{Pois}(\alpha+\epsilon)) &= F(\mathsf{Pois}(\alpha) + \mathsf{Pois}(\epsilon) + 1) - F(\mathsf{Pois}(\alpha) + \mathsf{Pois}(\epsilon)) \\ &\ge F(\mathsf{Pois}(\alpha)+1) - F(\mathsf{Pois}(\alpha)).\end{aligned}$$

Taking expectations, we observe that the derivative of $\mathbb{E}\big\{F(\mathsf{Pois}(\alpha))\big\}$ with respect to α is an increasing function of α and this completes the proof. □

Proof of Lemma 1 The differentiability and convexity of $\mathbb{E}\big\{[\mathsf{Pois}(\alpha) - C]^+\big\}$ with respect to α directly follows by using Lemma 6 with $F(\alpha) = [\alpha - C]^+$. To be able to obtain the derivative of $\mathbb{E}\big\{[\mathsf{Pois}(\alpha) - C]^+\big\}$ with respect to α, we apply the result in Lemma 6 to observe that this derivative is $\mathbb{E}\{[\mathsf{Pois}(\alpha)+1-C]^+ - [\mathsf{Pois}(\alpha)-C]^+\} = \mathbb{P}\{\mathsf{Pois}(\alpha) \geq C\}$, where the equality follows from the fact that the random variable in the last expectation takes value one if $\mathsf{Pois}(\alpha) \geq C$ and zero otherwise. □

Proof of Lemma 4 Given that we use the optimal policy, we let R_j^* be the number of accepted requests for fare class j and S_j^* be the number of reservations for fare class j that show up at the departure time. Naturally, both R_j^* and S_j^* are random variables. Since the number of requests for fare class j is a Poisson random variable with mean λ_j and the number of accepted requests for fare class j cannot exceed the number of requests, we have $\mathbb{E}\{R_j^*\} \leq \lambda_j$. Furthermore, since the show up decisions of different reservations are independent of each other and each reservation for fare class j shows up with probability q_j, we have $\mathbb{E}\{S_j^* \mid R_j^*\} = q_j\, R_j^*$, in which case, we obtain $\mathbb{E}\{S_j^*\} = q_j\, \mathbb{E}\{R_j^*\}$. We let $\hat{x}_j = \mathbb{E}\{R_j^*\}/\lambda_j$ for all $j = 1, \ldots, n$ so that $\hat{x}_j \in [0,1]$ and $\mathbb{E}\{S_j^*\} = q_j\, \lambda_j\, \hat{x}_j$. Noting that the optimal policy accepts R_j^* requests for fare class j and S_j^* reservations for fare class j show up at the departure time under the optimal policy, the net expected revenue obtained by the optimal policy can be written as $V_0(\bar{0}) = \sum_{j=1}^n f_j\, \mathbb{E}\{R_j^*\} - \theta\, \mathbb{E}\{[\sum_{j=1}^n S_j^* - C]^+\}$ and $V_0(\bar{0})$ satisfies

$$V_0(\bar{0}) = \sum_{j=1}^n f_j\, \mathbb{E}\{R_j^*\} - \theta\, \mathbb{E}\{[\textstyle\sum_{j=1}^n S_j^* - C]^+\} \leq \displaystyle\sum_{j=1}^n f_j\, \mathbb{E}\{R_j^*\} - \theta\, [\textstyle\sum_{j=1}^n \mathbb{E}\{S_j^*\} - C]^+$$
$$= \sum_{j=1}^n f_j\, \lambda_j\, \hat{x}_j^* - \theta\, [\textstyle\sum_{j=1}^n q_j\, \lambda_j\, \hat{x}_j^* - C]^+ \leq Z^*,$$

where the first inequality uses Jensen's inequality and the second inequality follows from the fact that $\hat{x} = (\hat{x}_1, \ldots, \hat{x}_n) \in [0,1]^n$ is a feasible but not necessarily an optimal solution to problem (8). □

B Appendix: Cutting Plane Method

In this section, we describe a standard cutting plane method that can be used to maximize the net expected revenue function in (17). This cutting plane method is based on representing the function $\sum_{\ell=1}^L \rho^\ell\, \mathbb{E}\{[\mathsf{Pois}(\sum_{j=1}^n q_j\, \lambda_j^\ell\, p_j) - C]^+\}$ by a series of cutting plane approximations.

Step 1 We initialize a guess at the optimal solution by choosing $p^1 = (p_1^1, \ldots, p_n^1)$ such that $p^1 \in [0,1]^n$. We initialize a guess at $\sum_{\ell=1}^L \rho^\ell\, \mathbb{E}\{[\mathsf{Pois}(\sum_{j=1}^n q_j\, \lambda_j^\ell\, p_j) - C]^+\}$ computed at the optimal solution by letting $y^1 = -\infty$. We initialize the iteration counter by letting $k = 1$.

Step 2 If we have $\sum_{\ell=1}^L \rho^\ell\, \mathbb{E}\{[\mathsf{Pois}(\sum_{j=1}^n q_j\, \lambda_j^\ell\, p_j^k) - C]^+\} = y^k$, then p^k is a maximizer of the net expected revenue function given in (17) and we stop. Otherwise, Lemma 1 implies that $\mathbb{E}\{[\mathsf{Pois}(\sum_{i=1}^n q_i\, \lambda_i^\ell\, p_i) - C]^+\}$ is a convex function of p and the derivative of this function with respect to p_j is given by the expression $q_j\, \lambda_j^\ell\, \mathbb{P}\{\mathsf{Pois}(\sum_{i=1}^n q_i\, \lambda_i^\ell\, p_i) \geq C\}$. Therefore, the function $\sum_{\ell=1}^L \rho^\ell\, \mathbb{E}\{[\mathsf{Pois}(\sum_{i=1}^n q_i\, \lambda_i^\ell\, p_i) - C]^+\}$ satisfies the subgradient inequality

$$\sum_{\ell=1}^L \rho^\ell\, \mathbb{E}\{[\mathsf{Pois}(\textstyle\sum_{i=1}^n q_i\, \lambda_i^\ell\, p_i) - C]^+\} \geq \displaystyle\sum_{\ell=1}^L \rho^\ell\, \mathbb{E}\{[\mathsf{Pois}(\textstyle\sum_{i=1}^n q_i\, \lambda_i^\ell\, p_i^k) - C]^+\}$$
$$+ \sum_{\ell=1}^L \sum_{j=1}^n \rho^\ell\, q_j\, \lambda_j^\ell\, \mathbb{P}\{\mathsf{Pois}(\textstyle\sum_{i=1}^n q_i\, \lambda_i^\ell\, p_i^k) \geq C\}\, [p_j - p_j^k]$$

for all $p \in [0,1]^n$. In this case, we let $\Psi^k = \sum_{\ell=1}^L \rho^\ell\, \mathbb{E}\big\{[\mathsf{Pois}(\sum_{i=1}^n q_i\, \lambda_i^\ell\, p_i^k) - C]^+\big\}$ and $\Phi_j^k = \sum_{\ell=1}^L \rho^\ell\, q_j\, \lambda_j^\ell\, \mathbb{P}\big\{\mathsf{Pois}(\sum_{i=1}^n q_i\, \lambda_i^\ell\, p_i^k) \geq C\big\}$ for all $j = 1,\ldots,n$ so that we can write the right side of the subgradient inequality above as $\Psi^k + \sum_{j=1}^n \Phi_j^k\, [p_j - p_j^k]$ and the last expression gives a cutting plane approximation to the function $\sum_{\ell=1}^L \rho^\ell\, \mathbb{E}\big\{[\mathsf{Pois}(\sum_{i=1}^n q_i\, \lambda_i^\ell\, p_i) - C]^+\big\}$ at the point p^k.

Step 3 We solve the problem

$$
\begin{aligned}
\max \quad & \sum_{\ell=1}^{L}\sum_{j=1}^{n} \rho^\ell\, f_j\, \lambda_j^\ell\, p_j - \theta\, y \\
& y \geq \Psi^\kappa + \sum_{j=1}^{n} \Phi_j^\kappa\, [p_j - p_j^\kappa] && \text{for all } \kappa = 1,\ldots,k \\
& p_j \in [0,1],\ \ y \geq 0 && \text{for all } j = 1,\ldots,n,
\end{aligned}
$$

where the first set of constraints represent the function $\sum_{\ell=1}^L \rho^\ell\, \mathbb{E}\big\{[\mathsf{Pois}(\sum_{j=1}^n q_j\, \lambda_j^\ell\, p_j) - C]^+\big\}$ by the cutting plane approximations that we have constructed at iterations $1,\ldots,k$. Letting $p^{k+1} = (p_1^{k+1},\ldots,p_n^{k+1})$ and y^{k+1} be an optimal solution to the problem above, we increase the iteration counter k by one and go back to Step 2.

The cutting plane method above is adopted from Section 7.2 in Ruszczynski (2006). Theorem 7.7 in Ruszczynski (2006) shows that the sequence of solutions $\{p^k\}_k$ generated by the cutting plane method satisfy $\lim_{k\to\infty} \Pi(p^k) = \Pi(p^*)$, where p^* is a maximizer of the net expected revenue function in (17). Furthermore, in Step 2, if we stop whenever $\big|\sum_{\ell=1}^L \rho^\ell\, \mathbb{E}\big\{[\mathsf{Pois}(\sum_{j=1}^n q_j\, \lambda_j^\ell\, p_j^k) - C]^+\big\} - y^k\big| \leq \epsilon$, then the solution p^k satisfies $\Pi(p^*) - \Pi(p^k) \leq \epsilon$. Therefore, the net expected revenue provided by the solution p^k deviates from the optimal net expected revenue by no more than ϵ.

C Appendix: Binary Acceptance Probabilities

We consider a problem instance with three fare classes. The fares and show up probabilities are $(f_1, f_2, f_3) = (60, 100, 120)$ and $(q_1, q_2, q_3) = (0.1, 0.2, 0.3)$. These fares and show up probabilities indeed satisfy $f_1/q_1 > f_2/q_2 > f_3/q_3$. The penalty cost is $\theta = 1000$. The arrival rates and capacity on the flight leg are $(\lambda_1, \lambda_2, \lambda_3) = (50, 50, 50)$ and $C = 25$. For this problem instance, we proceed to showing that if we impose binary constraints on the decision variables, then the optimal solution to problem (4) is $(p_1^*, p_2^*, p_3^*) = (0, 1, 1)$. Since there are three fare classes, there are $2^3 = 8$ possible solutions to problem (4) when we impose binary constraints on the decision variables. Table 11 shows the net expected revenues provided by these eight possible solutions. From this table, we observe that the solution $(p_1^*, p_2^*, p_3^*) = (0, 1, 1)$ provides the largest net expected revenue so that if we impose binary constraints on the decision variables, then the optimal solution to problem (4) is $(p_1^*, p_2^*, p_3^*) = (0, 1, 1)$. As a side observation, we note that if we do not impose binary constraints on the decision variables, then the optimal solution to problem (4) is $(p_1^*, p_2^*, p_3^*) \approx (1, 1, 0.56)$. Therefore, rounding the fractional acceptance probability up or down does not necessarily provide the optimal solution to problem (4) when we impose binary constraints.

We build on the deterministic approximation in Section 5 to understand why rounding the fractional acceptance probability up or down in the optimal solution to problem (4) does not necessarily provide

p	$(0,0,0)$	$(1,0,0)$	$(0,1,0)$	$(0,0,1)$	$(1,1,0)$	$(1,0,1)$	$(0,1,1)$	$(1,1,1)$
$\Pi(p)$	0.00	3,000.00	4,999.97	5,987.15	7,987.15	8,669.17	9,011.93	8,508.33

Table 11: Net expected revenues obtained by the eight possible solutions.

the optimal solution when we impose binary constraints. We consider a sequence of problem instances $\{\mathcal{P}^k : k \in \mathbb{Z}_+\}$ indexed by the integer parameter k. In problem instance $\mathcal{P}^k$, the fares, show up probabilities and penalty cost are as given in the previous paragraph, denoted by $(f_1^k, f_2^k, f_3^k) = (60, 100, 120)$, $(q_1^k, q_2^k, q_3^k) = (0.1, 0.2, 0.3)$ and $\theta^k = 1000$. On the other hand, the arrival rates and capacity in problem instance $\mathcal{P}^k$ are scaled by a factor k so that the arrival rates and capacity are given by $(\lambda_1^k, \lambda_2^k, \lambda_3^k) = (k\,50, k\,50, k\,50)$ and $C^k = k\,25$. We let p^k and x^k respectively be the optimal solutions to problems (4) and (8) when we solve these problems for problem instance $\mathcal{P}^k$. Similarly, we let $V_0^k(\bar{0})$ be the optimal net expected revenue for this problem instance. For problem instance $\mathcal{P}^k$, the right side of the inequality in (9) evaluates to $1 - 0.6/\sqrt{2\,\pi\,k}$, in which case, we obtain

$$1 \geq \frac{\Pi(p^k)}{V_0^k(\bar{0})} \geq \frac{\Pi(x^k)}{V_0^k(\bar{0})} \geq 1 - \frac{0.6}{\sqrt{2\,\pi\,k}}, \tag{22}$$

where the first inequality above is by the fact that $V_0^k(\bar{0})$ is the optimal net expected revenue for problem instance $\mathcal{P}^k$, whereas $\Pi(p^k)$ is the net expected revenue from the best open loop policy and the second inequality is by the fact that the acceptance probabilities x^k characterize a feasible open loop policy, but not necessarily the best one. As $k \to \infty$, the right side of the inequality above converges to one, implying that the percent gaps among $V_0^k(\bar{0})$, $\Pi(p^k)$ and $\Pi(x^k)$ vanish. Therefore, the deterministic approximation in (8) approximates problem (4) well when k becomes large. In other words, we can get a feel for how the optimal solution to problem (4) behaves by studying the optimal solution to the deterministic approximation in (8), as long as k is reasonably large. We observe that even when k is as small as one, the right side of the inequality in (22) already evaluates to about 0.76.

For problem instance $\mathcal{P}^k$, the deterministic approximation in (8) becomes $\max_{x \in [0,1]^3}\ k\,3000\,x_1 + k\,5000\,x_2 + k\,6000\,x_3 - 1000\,[k\,5\,x_1 + k\,10\,x_2 + k\,15\,x_3 - k\,25]^+$. We observe that it would never be optimal to exceed the capacity $k\,25$ in the last optimization problem as long as the penalty cost is large enough. In particular, if we exceed the capacity, then for every unit of increase in the decision variable x_1, we incur a cost of $k\,5000$, but generate only a revenue of $k\,3000$. Similarly, for every unit of increase in the decision variable x_2, we incur a cost of $k\,10000$, but generate a revenue of $k\,5000$, whereas for every unit of increase in the decision variable x_3, we incur a cost of $k\,15000$, but generate a revenue of $k\,6000$. In this case, given that it is never optimal to exceed the capacity, we can find an optimal solution to the last optimization problem by solving

$$\max\left\{k\,3000\,x_1 + k\,5000\,x_2 + k\,6000\,x_3 :\ k\,5\,x_1 + k\,10\,x_2 + k\,15\,x_3 \leq k\,25,\ x \in [0,1]^3\right\}. \tag{23}$$

Problem (23) is a slight modification of a well known knapsack problem instance given in Chapter 16 in Cormen et al. (2009). The optimal solution to the problem above is $(x_1^*, x_2^*, x_3^*) \approx (1, 1, 0.67)$ with an objective value of $k\,12000$, but if we impose binary constraints on the decision variables, then the optimal solution is $(x_1^*, x_2^*, x_3^*) = (0, 1, 1)$ with an objective value of $k\,11000$. Therefore, rounding the fractional acceptance probability up or down in the optimal solution to problem (23) does not provide the optimal solution when we impose binary constraints. Furthermore, if we round the fractional acceptance probability in the optimal solution to problem (23) up or down, then the best solution we obtain has an optimality gap of more than 15% for the version of problem (23) with binary constraints. Noting that the deterministic approximation in (23) is a good approximation to problem (4) for large values of k, it is not surprising that similar observations would apply to the optimal solution to problem (4) as long as k is reasonably large.

The message from this discussion is the following. If we have binary constraints on the decision variables, then for many knapsack problem instances in practice, solving the continuous relaxation of the problem and rounding down the fractional decision variable in the optimal solution provides a good solution. This is the popular greedy heuristic for knapsack problems and Williamson and Shmoys (2011) show that a minor enhancement of the greedy heuristic provides an approximation guarantee of two, implying that the greedy heuristic always provides at least half of the optimal objective value. The practical performance of the greedy heuristic is generally much better than this theoretical worst case guarantee, but there are pathological problem instances where the performance of the greedy heuristic can indeed be as bad as half of the optimal. It turns out we can argue that a similar situation holds for problem (4). In particular, as the arrival rates and capacity are scaled linearly with the same rate, problem (4) can be approximated well by using the deterministic approximation in (8). The behavior of the deterministic approximation in (8) is, in turn, similar to that of a knapsack problem. As indicated by our findings in Table 6, solving problem (4) and rounding the fractional acceptance probability in the optimal solution up or down generally yields minimal performance loss when we impose binary constraints on the decision variables. However, it may still be possible to generate pathological problem instances where this approach results in more than minimal performance loss, though it appears to be a rare occurrence to get these pathological problem instances in practice.