

An Approximation Algorithm for Capacity Allocation over a Single Flight Leg with Fare-Locking

Mika Sumida

School of Operations Research and Information Engineering,
Cornell University, Ithaca, New York 14853
ms3268@cornell.edu

Huseyin Topaloglu

School of Operations Research and Information Engineering,
Cornell Tech, New York City, New York 10011
topaloglu@orie.cornell.edu

December 22, 2016

Abstract

In this paper, we study a revenue management model over a single flight leg, where the customers are allowed to lock an available fare. Each customer arrives into the system with an interest in purchasing a ticket for a particular fare class. If this fare class is available, the customer either immediately purchases the ticket by paying the fare or locks the fare by paying a fee. If the customer locks the fare, then the airline reserves the capacity for the customer for a certain duration of time. At the end of this duration of time, the customer makes her ultimate purchase decision at the locked fare. The goal of the airline is to find a policy to decide which set of fare classes to make available at each time period to maximize the total expected revenue. Such fare locking options are commonly offered by airlines today, but the dynamic programming formulation of the revenue management problem with the option to lock an available fare has a high dimensional state variable that keeps track of the locked fares. We develop an approximate policy that is guaranteed to obtain at least half of the optimal total expected revenue. Our approach is based on leveraging a linear programming approximation to decompose the problem by the seats on the flight and solving a dynamic program that controls the capacity on each seat separately. We also show that our results continue to hold when the airline makes pricing decisions, instead of fare class availability decisions. Our numerical study shows that the practical performance our approximate policy is remarkably good when compared with a tractable upper bound on the optimal total expected revenue.

A recent practice adopted by airlines is to offer the option of locking an available fare. With this option, the customers lock an available fare for a certain duration of time by paying a fee. In return for the fee, the airline reserves the capacity for the customer until she makes her ultimate purchase decision within a certain duration of time. If the customer ultimately purchases the ticket, then she pays the fare that she locked. If the customer ultimately does not purchase the ticket, then the fee for locking the fare is wasted. On the one hand, by using the option of locking an available fare, the customers can protect themselves against lack of capacity until their travel plans materialize. On the other hand, by offering the option of locking an available fare, the airlines can come up with another revenue stream. Nevertheless, the option of locking an available fare brings an additional source of uncertainty in the revenue management decisions of the airlines. In particular, since the ultimate purchase decisions of the customers with locked fares are uncertain, the capacity available for future customers also becomes uncertain.

In this paper, we study a revenue management model over a single flight leg when customers have the option of locking an available fare. Each customer arrives into the system with an interest in purchasing a ticket for a particular fare class. If this fare class is available, then the customer either purchases the ticket immediately or locks the fare by paying a fee. If the customer locks the fare, then the airline reserves the capacity for the customer for a certain duration of time, allowing the customer to delay her purchase decision. If the customer ultimately decides to purchase the ticket, then she pays the fare she locked. The goal is to find a policy to decide which fare classes to make available at each time period in the selling horizon to maximize the total expected revenue. The dynamic programming formulation of this problem requires a high dimensional state variable that keeps track of the customers with locked fares. Therefore, finding the optimal policy is computationally challenging. We construct an approximate policy that is guaranteed to obtain at least half of the optimal total expected revenue.

The construction of our approximate policy has three components. First, we give a linear programming approximation to the problem under the assumption that the demands for the fare classes and the locking decisions of the customers take on their expected values. Second, we use the optimal solution to the linear programming approximation to decide whether a customer arriving at a certain time period with an interest in a certain fare class is offered the fare class. If the customer is offered the fare class, then we “dispatch” the customer randomly to one of the seats on the flight. Therefore, each seat on the flight receives an exogenous arrival of customers, in which case, we can solve a dynamic program to manage the capacity on each seat individually. Third, we use the dynamic program that we solve for each seat to obtain approximations to the value functions in the dynamic programming formulation of the original problem. We use these approximations to the value functions to construct our approximate policy. We show that our approximate policy obtains at least half of the optimal total expected revenue.

To our knowledge, there are no policies that can be computed efficiently, while providing a constant factor approximation guarantee for the revenue management problem over a single flight

leg with the option of locking an available fare. Our approximate policy fills this gap. Our numerical study indicates that the practical performance of this policy is substantially better than the theoretical approximation guarantee. To obtain our approximation guarantee, we show that the optimal objective value of the linear programming approximation provides an upper bound on the optimal total expected revenue and the total expected revenue from the approximate policy is at least half of the optimal objective value of the linear programming approximation. So, letting OPT be the optimal total expected revenue, LP be the optimal objective value of the linear programming approximation and APP be the total expected revenue from the approximate policy, we show that $LP \geq OPT$ and $APP \geq LP/2$, which yield $APP \geq LP/2 \geq OPT/2$, implying that the approximate policy obtains at least half of the optimal total expected revenue. As a side note, since the total expected revenue from the approximate policy cannot exceed the optimal total expected revenue, we have $OPT \geq APP$, in which case, the last chain of inequalities yields $OPT \geq LP/2 \geq OPT/2$. Thus, the upper bound provided by the linear programming approximation also deviates from the optimal total expected revenue by at most a factor of two. Linear programming approximations are often used in the revenue management literature, but we are not aware of such a tightness result for the upper bound from the linear programming approximation.

The demand model that we use is the so called independent demand model, where each customer arrives with an interest in purchasing a ticket for a particular fare class and she leaves without a purchase when this fare class is not available. In this case, the decision that we make at each time period is the set of fare classes offered to the customers. Such a simple demand model allows us to explain the ideas behind our approximate policy clearly, but we show that all of our results continue to hold under the price dependent demand model, where each customer arrives with a particular willingness to pay and she leaves without a purchase when the price for the ticket exceeds the willingness to pay of the customer. In this case, the decision that we make at each time period is the price that we charge for the ticket. Under price dependent demand, we can still construct an approximate policy that obtains at least half of the optimal total expected revenue.

RELATED LITERATURE. There are two studies of the revenue management problem over a single flight leg with an option to lock an available fare. Aydin et al. (2016) consider the problem under independent demand, whereas Chen and Chen (2016) consider the problem under price dependent demand. The models in Aydin et al. (2016) and Chen and Chen (2016) precisely correspond to the models that we consider in this paper under independent demand and price dependent demand. Aydin et al. (2016) demonstrate that it is not always beneficial for the airline to offer the option of locking a fare, but there are cases when the revenue improvements from offering the option can be substantial. Chen and Chen (2016) study how the option of locking an available fare affects the pricing decisions of the airline. As far as we are aware, these papers are the first studies of the problem but they do not provide policies with performance guarantees.

The construction of our approximate policy borrows from Wang et al. (2015) and Gallego et al. (2015). These papers consider the problem of assigning resources to arriving customers,

when the reward from assigning a resource to a customer depends on the identities of the resource and the customer. In Wang et al. (2015), one chooses the resource to assign to each customer, whereas in Gallego et al. (2015), one chooses the set of resources to offer to each customer, among which the customer makes a choice. Both papers use a linear programming approximation to decompose the problem by the resources and to construct approximate policies with performance guarantees. The performance guarantees in these papers require a special choice of Lagrange multipliers in a suitable relaxation of the problem and use arguments specially tailored to the problem class. Our performance guarantee does not require a special choice of Lagrange multipliers and uses an induction argument that could be applicable in a variety of settings.

The option of locking an available fare is one approach to construct a revenue stream while providing some flexibility to the customers and the airlines, but there are other approaches studied in the literature. Karaesmen and van Ryzin (2004) study overbooking models with substitutable flights, where the customers that are denied boarding from one flight can be accommodated on the next one. Gallego and Phillips (2004) focus on flexible products, where the customers purchase ticket for a particular date without knowing which particular flight they eventually take. Gallego et al. (2008) give a model to study callable products, where the airline may buy back the tickets purchased by low fare customers to sell them to high fare customers. Gallego and Stefanescu (2009) study models for upgrading customers to higher fare classes when the airline runs out of capacity to serve the customers for their originally purchased fare classes. Balseiro et al. (2011) focus on call options for sports tournaments, which allow the customers to reserve tickets early by paying a small fee and ultimately purchase the ticket only if their teams are on the final game.

Our approximate policy makes use of a linear programming approximation that is formulated under the assumption that all random components of the problem take on their expected values. Similar deterministic approximations are often used in the literature. For example, Gallego and van Ryzin (1994) focus on dynamic pricing on a single flight leg, whereas Gallego and van Ryzin (1997) focus on dynamic pricing over a flight network. Talluri and van Ryzin (1998) focus on capacity allocation over a flight network. Gallego et al. (2004) focus on network revenue management with customer choice behavior. Kunnumkal et al. (2012) focus on overbooking over a flight network. The authors give deterministic approximations for each one of these settings, show that the deterministic approximation provides an upper bound on the optimal total expected revenue and develop policies that are asymptotically optimal as the capacity and the demand increases linearly with the same rate, but do not provide policies with constant factor approximation guarantees when the capacity and the demand do not necessarily get large.

When we give a dynamic programming formulation for the revenue management problem under the option of locking an available fare, the state variable ends up being high dimensional, since we need to keep track of the customers with locked fares over a suitable portion of the history of the system. High dimensional state variables appear in numerous setting in revenue management. For example, in revenue management problems over a network of flight legs, the state variable needs

to keep track of the remaining capacities on all of the flight legs. Adelman (2007) constructs linear approximations of the value functions in the dynamic programming formulation of the network revenue management problem. The author chooses the slopes of the value function approximations by solving a linear program that represents the dynamic programming formulation of the network revenue management problem. Liu and van Ryzin (2008) extract policies from a linear programming approximation of the network revenue management problem, but the authors consider the network revenue management problem under customer choice, where the customers are offered a set of itineraries, among which they make a choice. Zhang and Adelman (2009) construct linear approximations to the value functions in the network revenue management problem under customer choice, whereas Topaloglu (2009) constructs separable and piecewise linear approximations. Other strategies for approximating the value functions can be found in Kunnumkal and Topaloglu (2010), Meissner and Strauss (2012), Meissner et al. (2012) and Kunnumkal (2014). Vossen and Zhang (2015) and Kunnumkal and Talluri (2016) show the relationships between these approximation strategies and indicate that some of them are equivalent to each other.

Dynamic programming formulations of overbooking problems also require high dimensional state variables, since we need to keep track of the reservations for different fare classes and the reservations for different fare classes cancel and show up with different probabilities due to different restrictions. Subramanian et al. (1999) give a dynamic programming formulation of the overbooking problem over a single flight leg and show that the state variable collapses to a scalar when the cancellation and show up probabilities are the same for all fare classes. This model assumes that there can be at most one cancellation at each time period. Bertsimas and Popescu (2003) use a deterministic approximation to compute policies for overbooking over a flight network. Erdelyi and Topaloglu (2010) study overbooking problems over a network of flight legs and propose an approach to decompose the problem by the flight legs, but their approach does not have a performance guarantee. Lan et al. (2011) give a robust formulation of the overbooking problem, provide policies that minimize regret and characterize the structure of the optimal policy. Aydin et al. (2013) give an overbooking model over a single flight leg, where any one of the reservations can cancel at a particular time period, so that there can be multiple cancellations.

ORGANIZATION. In Section 1, we formulate the problem as a dynamic program. In Section 2, we give a linear programming approximation. In Section 3, we show how to decompose the problem by the seats on the flight so that we can solve a dynamic program to control the capacity on each seat separately. This approach provides approximations to the value functions in the original dynamic programming formulation. In Section 4, we use these approximations to construct an approximate policy and show that the approximate policy obtains at least half of the optimal total expected revenue. In Section 5, we extend our approach to deal with the case where the customers interested in different fare classes have different parameters governing their purchase behavior. In Section 6, we show that our results continue to hold under price dependent demand. In Section 7, we provide a numerical study. In Section 8, we conclude.

1 Problem Formulation

We have n fare classes indexed by $N = \{1, \dots, n\}$. The capacity available on the flight is C . The selling horizon has τ time periods indexed by $T = \{1, \dots, \tau\}$. There is at most one customer arrival at each time period. With probability $\lambda_{i,t}$, a customer with an interest in purchasing a ticket for fare class i arrives into the system at time period t . If fare class i is not available, then the customer leaves without a purchase. If fare class i is available, then the customer either purchases the ticket or locks the fare. In particular, with probability ρ , the customer purchases the ticket by paying the fare r_i for fare class i . With probability $1 - \rho$, the customer locks the fare by paying the fee h . Locking the fare ensures that the capacity on the flight is reserved for the customer and the customer makes her purchase decision for the ticket for fare class i after L time periods. With probability π , the customer ultimately purchases the ticket after L time periods, in which case, she pays the fare r_i . With probability $1 - \pi$, the customer does not purchase the ticket and the capacity reserved for this customer becomes available for other customers. We want to find a policy to decide which fare classes to make available at each time period so that we maximize the total expected revenue over the selling horizon.

We proceed to describing the sequence of events that take place at time period t . Considering the capacity available on the flight and the customers who locked fares at time periods $t - L, \dots, t - 1$, we decide which fare classes to make available at time period t . Next, we observe the customer arrival at time period t , along with whether this customer purchases the ticket for the fare class she is interested in or locks the fare. Finally, if the customer arriving at time period $t - L$ locked the fare at this time period, then we observe the ultimate purchase decision of this customer. If the customer ultimately decides not to make a purchase, then the capacity reserved for this customer becomes available for other customers. At the beginning of a time period, we use x to denote the remaining capacity on the flight and $y_\ell \in \{0, 1\}$ to capture whether the customer arriving ℓ time periods ago locked the fare that is of interest to her. In particular, we have $y_\ell = 1$ if and only if the customer arriving ℓ time periods ago locked the fare that is of interest to her. At the beginning of a time period, we observe the customers who locked the fares at the previous L time periods. Therefore, the vector (y_1, \dots, y_L) captures the customers with locked fares at the previous time periods who have not yet made their purchase decisions.

To formulate the problem as a dynamic program, we use (x, y_1, \dots, y_L) as the state variable at a time period. We capture the decisions that we make at a time period by using the vector $u = (u_1, \dots, u_n) \in \{0, 1\}^n$, where $u_i = 1$ if and only if we make fare class i available. Since we can make a fare class available only when there is remaining capacity on the flight, the set of feasible decisions at a time period is given by $\mathcal{U}(x) = \{u \in \{0, 1\}^n : u_i \leq x \ \forall i \in N\}$. If we make fare class i available and a customer with an interest in fare class i arrives into the system, then the customer purchases the ticket with probability ρ , whereas the customer locks the fare with probability $1 - \rho$ and she ultimately purchases the ticket after L time periods with probability π . Therefore, if we make fare class i available, then we obtain an expected revenue of $\rho r_i + (1 - \rho)(h + \pi r_i)$ from a

customer with an interest in fare class i . If the fee to lock the fare depends on the fare class, then we can replace h with h_i . For notational brevity, we let $f_i = \rho r_i + (1 - \rho)(h + \pi r_i)$, capturing the expected revenue from a customer with an interest in fare class i . Given that there is a customer arrival at time period t and the fare class that is of interest to this customer is available for purchase, we use the random variable B_t to capture whether this customer locks the fare that is of interest to her. In particular, B_t is a Bernoulli random variable with parameter $1 - \rho$, taking value 1 if and only if the customer locks the fare. As a function of whether the customer arriving at time period $t - L$ locked the fare that is of interest to her, we use the random variable $D_t(y_L)$ to capture whether the customer ultimately decides not to purchase the ticket at time period t . In particular, $D_t(y_L)$ is a Bernoulli random variable with parameter $(1 - \pi)y_L$, taking value 1 if and only if the customer who locked a fare at time period $t - L$ ultimately decides not to purchase the ticket at time period t . If $y_L = 0$ so that we do not have a customer who locked a fare L time periods ago, then $D_t(y_L) = 0$ with probability one. Letting $V_t(x, y_1, \dots, y_L)$ be the optimal total expected revenue over time periods t, \dots, τ given that the state of the system at time period t is (x, y_1, \dots, y_L) , we can compute the value functions $\{V_t(\cdot) : t \in T\}$ by solving the dynamic program

$$V_t(x, y_1, y_2, \dots, y_L) = \max_{u \in \mathcal{U}(x)} \left\{ \sum_{i \in N} \lambda_{i,t} u_i \left[f_i + \mathbb{E} \left\{ V_{t+1}(x - 1 + D_t(y_L), B_t, y_1, \dots, y_{L-1}) \right\} \right] + \left[1 - \sum_{i \in N} \lambda_{i,t} u_i \right] \mathbb{E} \left\{ V_{t+1}(x + D_t(y_L), 0, y_1, \dots, y_{L-1}) \right\} \right\}, \quad (1)$$

with the boundary condition that $V_{\tau+1}(\cdot) = 0$. Noting that the capacity available on the flight is C , the optimal total expected revenue is given by $V_1(C, 0, \dots, 0)$.

In the dynamic program in (1), if the customer who locked the fare at time period $t - L$ ultimately decides not to purchase the ticket at time period t , then we have $D_t(y_L) = 1$, in which case, the capacity reserved for this customer becomes available for other customers. Also, if a customer is interested in fare class i at time period t and this fare class is available, then we account for the expected revenue f_i from this customer at time period t , although the customer may lock the fare and ultimately make her purchase decision after L time periods, in which case, the fare is collected after L time periods. Section 4.4.2 in Talluri and van Ryzin (2005) focuses on overbooking problems over a single flight leg and shows that we can account for the expected refund to a customer who does not show up for the flight when the customer purchases the ticket rather than when the customer makes her decision to show up for the flight. By using the same argument, it is indeed possible to show that we can account for the expected revenue from a customer at the time period she arrives into the system, although a portion of the expected revenue from the customer is obtained after L time periods. Due to the high dimensional state variable in the dynamic program in (1), it is difficult to compute the optimal policy. In the rest of the paper, we focus on developing an approximate policy with a performance guarantee.

In our model, the probability $1 - \rho$ of locking the fare, the probability $1 - \pi$ of ultimately not purchasing the ticket and the number of time periods L that a customer with a locked fare

waits until she makes a purchase decision do not depend on the fare class that is of interest to the customer. In Section 5, we explain that our results continue to hold when we allow $1 - \rho$, $1 - \pi$ and L to depend on the fare class, but the state variable becomes more involved. By allowing L to depend on the fare class, we can also capture the situation where the customers can lock the fare for different durations, where we define different fare classes corresponding to different durations to lock the fare. Lastly, we note that the decisions of the customers to lock a fare and to ultimately purchase the ticket after locking the fare are probabilistic. So, the customers do not anticipate the behavior of the policy that is used to decide which fare classes to make available. This approach is precisely the same one adopted by Aydin et al. (2016) and Chen and Chen (2016). In our model, we can interpret $1 - \rho$ as the fraction of customers with uncertain travel plans. These customers lock the fare. The travel plans of these customers become certain after L time periods. We interpret $1 - \pi$ as the fraction of customers whose uncertain travel plans do not materialize. These customers do not purchase the ticket L time periods after locking the fare. By using a probabilistic model of the customer behavior, we capture the customers as boundedly rational and we obtain a tractable model to make operational decisions.

2 Upper Bound on the Optimal Total Expected Revenue

We construct a linear program to obtain an upper bound on the optimal total expected revenue. In the linear program, we use the decision variable $z_{i,t}$ to capture the expected number of customers arriving at time period t with an interest in fare class i and finding this fare class available. Note that these customers either purchase the ticket for fare class i or lock the fare. Consider the customers arriving with an interest in fare class i and finding this fare class available. These customers consume the capacity on the flight. Using $\mathbf{1}(\cdot)$ to denote the indicator function, the total expected capacity consumed due to these customers over time periods $1, \dots, t$ is $\sum_{\kappa \in T} \mathbf{1}(\kappa \leq t) z_{i,\kappa}$. Again, consider the customers arriving with an interest in fare class i and finding this fare class available. Each one of these customers locks the fare with probability $1 - \rho$ and ultimately decides not to purchase the ticket after L time periods with probability $1 - \pi$. The total expected capacity released due to these customers over time periods $1, \dots, t$ is $\sum_{\kappa \in T} \mathbf{1}(\kappa \leq t - L) (1 - \rho) (1 - \pi) z_{i,\kappa}$, where we use the fact that customers locking the fare after time period $t - L$ release the capacity after time period t . Therefore, the net total expected capacity consumed over time periods $1, \dots, t$ is $\sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t) z_{i,\kappa} - \sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t - L) (1 - \rho) (1 - \pi) z_{i,\kappa}$. Using the decision variables $z = \{z_{i,t} : i \in N, t \in T\}$, we consider the linear program

$$\max_{z \in \mathbb{R}_+^{n \times \tau}} \left\{ \sum_{t \in T} \sum_{i \in N} f_i z_{i,t} : z_{i,t} \leq \lambda_{i,t} \quad \forall i \in N, t \in T, \right. \\ \left. C - \sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t) z_{i,\kappa} + \sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t - L) (1 - \rho) (1 - \pi) z_{i,\kappa} \geq 0 \quad \forall t \in T \right\}. \quad (2)$$

The objective function accounts for the total expected revenue over the selling horizon. The first constraint ensures that the expected number of customers arriving at time period t with

an interest in fare class i and finding this fare class available cannot exceed the expected number of customers arriving with an interest in fare class i . Since $\sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t) z_{i,\kappa} - \sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t - L) (1 - \rho) (1 - \pi) z_{i,\kappa}$ is the net total expected capacity consumed over time periods $1, \dots, t$, the second constraint ensures that the net total expected capacity consumed cannot exceed the capacity available on the flight. We use the linear program above to obtain an upper bound on the optimal total expected revenue. Later on, we construct an approximate policy that allows us to obtain a total expected revenue that is at least half of this upper bound on the optimal total expected revenue. Therefore, the approximate policy that we construct obtains at least half of the optimal total expected revenue.

In the next lemma, we show that the optimal objective value of problem (2) provides an upper bound on the optimal total expected revenue.

Lemma 1 *Letting ζ^* be the optimal objective value of problem (2) and noting that $V_1(C, 0, \dots, 0)$ is the optimal total expected revenue, we have $\zeta^* \geq V_1(C, 0, \dots, 0)$.*

Proof. Under the optimal policy, we let $Z_{i,t}^* = 1$ if a customer arrives at time period t with an interest in fare class i and finds this fare class available. Otherwise, $Z_{i,t}^* = 0$. Note that $Z_{i,t}^*$ is a random variable. As a function of whether a customer arrives at time period t with an interest in fare class i and finds this fare class available, we use the random variable $Q_t(Z_{i,t}^*)$ to capture whether the customer locks the fare at time period t for fare class i and ultimately decides not to purchase the ticket after L time periods. In particular, $Q_t(Z_{i,t}^*)$ is a Bernoulli random variable with parameters $(1 - \rho) (1 - \pi) Z_{i,t}^*$, taking value 1 if and only if a customer arrives at time period t with an interest in fare class i , finds this fare class available, locks the fare and ultimately does not purchase the ticket after L time periods. Under the optimal policy, for all $t \in T$, we have

$$C - \sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t) Z_{i,\kappa}^* + \sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t - L) Q_{\kappa}(Z_{i,\kappa}^*) \geq 0. \quad (3)$$

In particular, the customers arriving with an interest in a fare class and finding the fare class available either purchase the ticket or lock the fare. In either case, they consume the capacity. In (3), $\sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t) Z_{i,t}^*$ is the total capacity consumed by all customers at or before time period t . The customers arriving with an interest in a fare class, finding this fare class available, locking the fare and ultimately not purchasing the ticket release the capacity after L time periods. In (3), $\sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t - L) Q_{\kappa}(Z_{i,\kappa}^*)$ is the total capacity released by all customers at or before time period t . In (3), we state that the remaining capacity at the end of time period t under the optimal policy is nonnegative. Taking expectations in (3), since $\mathbb{E}\{Q_{\kappa}(Z_{i,\kappa}^*)\} = (1 - \rho) (1 - \pi) \mathbb{E}\{Z_{i,\kappa}^*\}$, the solution $\{\mathbb{E}\{Z_{i,t}^*\} : i \in N, t \in T\}$ satisfies the second constraint in problem (2).

We let $D_{i,t} = 1$ if a customer arrives at time period t with an interest in fare class i . Otherwise, $D_{i,t} = 0$. Note that $D_{i,t}$ is a random variable satisfying $\mathbb{E}\{D_{i,t}\} = \lambda_{i,t}$. By the definition of $Z_{i,t}^*$, we have $Z_{i,t}^* \leq D_{i,t}$. Taking expectations in the last inequality, it follows that the

solution $\{\mathbb{E}\{Z_{i,t}^*\} : i \in N, t \in T\}$ satisfies the first constraint in problem (2) as well. In this case, the solution $\{\mathbb{E}\{Z_{i,t}^*\} : i \in N, t \in T\}$ is feasible to problem (2). From each customer that arrives with an interest in fare class i and finds this fare class available, we obtain an expected revenue of f_i . Thus, noting the definition of $Z_{i,t}^*$ once more, we can write the total expected revenue obtained by the optimal policy as $V_1(C, 0, \dots, 0) = \sum_{t \in T} \sum_{i \in N} f_i \mathbb{E}\{Z_{i,t}^*\}$, which implies that the objective value provided by the solution $\{\mathbb{E}\{Z_{i,t}^*\} : i \in N, t \in T\}$ for problem (2) is equal to $V_1(C, 0, \dots, 0)$. Since the solution $\{\mathbb{E}\{Z_{i,t}^*\} : i \in N, t \in T\}$ is feasible to problem (2) and the objective value provided by this solution for problem (2) is equal to $V_1(C, 0, \dots, 0)$, the optimal objective value ζ^* for problem (2) is at least $V_1(C, 0, \dots, 0)$. \square

In the revenue management literature, it is common to construct deterministic optimization problems to obtain upper bounds on the optimal total expected revenue. Gallego and van Ryzin (1997) give an example in dynamic pricing, Talluri and van Ryzin (1998) give an example in capacity allocation, Gallego et al. (2004) give an example in network revenue management with customer choice and Kunnumkal et al. (2012) give an example in overbooking. In this paper, we construct an approximate policy that allows us to obtain a total expected revenue that is at least half of the upper bound provided by the optimal objective value of problem (2). This result immediately implies that the upper bound provided by the optimal objective value of problem (2) never exceeds the optimal total expected revenue by more than a factor of two.

In the next section, we approximate the value functions in the dynamic program in (1), which allows us to construct an approximate policy.

3 Approximations to the Value Functions

Our approach for approximating the value functions in the dynamic program in (1) is based on decomposing the dynamic program by the seats on the flight and managing the availability of the fare classes for each seat separately. We use $\{z_{i,t}^* : i \in N, t \in T\}$ to denote an optimal solution to problem (2). When we manage the availability of the fare class on a particular seat, with probability $z_{i,t}^*/C$, a customer arrives at time period t with an interest in fare class i . The interpretation of the probability $z_{i,t}^*/C$ will shortly become clear. If fare class i is not available, then the customer immediately leaves without a purchase. If fare class i is available, then the customer purchases the ticket or locks the fare. In particular, as in Section 1, with probability ρ , the customer purchases the ticket by paying the fare r_i , whereas with probability $1 - \rho$, the customer locks the fare by paying the fee h . If the customer locks the fare, then she makes her ultimate purchase decision for the ticket for fare class i after L time periods. With probability π , the customer ultimately purchases the ticket for fare class i after L time periods by paying the fare r_i . With probability $1 - \pi$, the customer ultimately does not purchase the ticket, in which case, the capacity reserved for this customer becomes available for other customers. Therefore, when we manage the availability of the fare classes for each seat separately, the mechanics of the problem is identical to the mechanics in the original problem. Since $z_{i,t}^* \leq \lambda_{i,t}$, the probability $z_{i,t}^*/C$ implies that we “reject” a customer

interested in fare class i with probability $(\lambda_{i,t} - z_{i,t}^*)/\lambda_{i,t}$. If we “accept” the customer, then we randomly dispatch her to one of the seats. Since we manage one unit of capacity, we can find the optimal policy to manage the capacity in a tractable fashion. In particular, when we manage the capacity on one seat, we use v_t to denote the optimal total expected revenue over time periods t, \dots, τ given that the capacity on the seat is available at time period t . Using $[\cdot]^+ = \max\{\cdot, 0\}$, we can compute the value functions $\{v_t : t \in T\}$ by solving the dynamic program

$$\begin{aligned} v_t &= \sum_{i \in N} \frac{z_{i,t}^*}{C} \max \left\{ f_i + (1 - \rho)(1 - \pi) v_{t+L+1}, v_{t+1} \right\} + \left\{ 1 - \sum_{i \in N} \frac{z_{i,t}^*}{C} \right\} v_{t+1} \\ &= \sum_{i \in N} \frac{z_{i,t}^*}{C} \left[f_i + (1 - \rho)(1 - \pi) v_{t+L+1} - v_{t+1} \right]^+ + v_{t+1}, \quad (4) \end{aligned}$$

with the boundary condition that $v_{\tau+1} = \dots = v_{\tau+L+1} = 0$. Note that we continue using the fact that f_i is the expected revenue from a customer interested in fare class i .

In the dynamic program in (4), the two terms in the maximum operator correspond to making and not making fare class i available. Consider a customer arriving at time period t with an interest in fare class i . If we make fare class i available, then with probability $(1 - \rho)(1 - \pi)$, the customer locks the fare and ultimately decides at time period $t + L$ not to purchase the ticket for fare class i , in which case, the capacity on the seat becomes available for use at time period $t + L + 1$. If we do not make fare class i available, then the customer leaves without a purchase, in which case, the capacity on the seat is available at time period $t + 1$. Intuitively speaking, since the capacity on the flight is C , the total expected revenue obtained by managing the availability of the fare classes on each seat separately is Cv_1 . In the next lemma, we show that Cv_1 is at least half of the upper bound on the optimal total expected revenue provided by problem (2).

Lemma 2 *Letting ζ^* be the optimal objective value of problem (2), if $\{v_t : t \in T\}$ are obtained by solving the dynamic program in (4), then we have $Cv_1 \geq \zeta^*/2$.*

Proof. By (4), we observe that $C(v_t - v_{t+1}) = \sum_{i \in N} z_{i,t}^* [f_i + (1 - \rho)(1 - \pi) v_{t+L+1} - v_{t+1}]^+ \geq \sum_{i \in N} z_{i,t}^* [f_i + (1 - \rho)(1 - \pi) v_{t+L+1} - v_{t+1}]$. Adding this chain of inequalities over all $t \in T$, we obtain $Cv_1 \geq \sum_{t \in T} \sum_{i \in N} z_{i,t}^* [f_i + (1 - \rho)(1 - \pi) v_{t+L+1} - v_{t+1}]$. Since $\sum_{t \in T} \sum_{i \in N} f_i z_{i,t}^*$ corresponds to the optimal objective value of problem (2), this inequality is equivalent to $Cv_1 \geq \zeta^* + \sum_{t \in T} \sum_{i \in N} z_{i,t}^* [(1 - \rho)(1 - \pi) v_{t+L+1} - v_{t+1}]$. In the last inequality, we can write the expression $\sum_{t \in T} \sum_{i \in N} z_{i,t}^* v_{t+L+1}$ equivalently as

$$\begin{aligned} \sum_{\kappa \in T} \sum_{i \in N} z_{i,\kappa}^* v_{\kappa+L+1} &= \sum_{\kappa \in T} \sum_{i \in N} z_{i,\kappa}^* \sum_{t \in T} \mathbf{1}(t \geq \kappa + L + 1) (v_t - v_{t+1}) \\ &= \sum_{t \in T} (v_t - v_{t+1}) \sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t - 1 - L) z_{i,\kappa}^*. \end{aligned}$$

Repeating the argument in the chain of equalities above after replacing L with zero, we can also write the expression $\sum_{t \in T} \sum_{i \in N} z_{i,t}^* v_{t+1}$ equivalently as $\sum_{t \in T} (v_t - v_{t+1}) \sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t - 1) z_{i,\kappa}^*$. In

this case, replacing $\sum_{t \in T} \sum_{i \in N} z_{i,t}^* v_{t+L+1}$ and $\sum_{t \in T} \sum_{i \in N} z_{i,t}^* v_{t+1}$ in the inequality $C v_1 \geq \zeta^* + \sum_{i \in N} \sum_{t \in T} z_{i,t}^* [(1-\rho)(1-\pi)v_{t+L+1} - v_{t+1}]$ with their equivalent expressions, we get

$$C v_1 \geq \zeta^* + \sum_{t \in T} (v_t - v_{t+1}) \left\{ \sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t-1-L) (1-\rho)(1-\pi) z_{i,\kappa}^* - \sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t-1) z_{i,\kappa}^* \right\}.$$

Since $\{z_{i,t}^* : i \in N, t \in T\}$ is a feasible solution to problem (2), noting the second constraint in this problem for time period $t-1$, we have $\sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t-1-L) (1-\rho)(1-\pi) z_{i,\kappa}^* - \sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t-1) z_{i,\kappa}^* \geq -C$. Furthermore, by (4), since $[a]^+ \geq 0$ for all $a \in \mathfrak{R}$, we have $v_t - v_{t+1} \geq 0$ for all $t \in T$. In this case, the inequality above implies that $C v_1 \geq \zeta^* - \sum_{t \in T} (v_t - v_{t+1}) C = \zeta^* - C v_1$. Focusing on the first and last expressions in the last chain of inequalities, it follows that $2 C v_1 \geq \zeta^*$. \square

In the next section, we develop an approximate policy that uses $\{v_t : t \in T\}$ computed through the dynamic program in (4) to obtain at least half of the optimal total expected revenue.

4 Performance Guarantee for the Approximate Policy

Our approximate policy is, intuitively speaking, based on following the decision rule in the dynamic program in (4). We let $\{v_t : t \in T\}$ be computed through the dynamic program in (4). In the approximate policy, if $f_i + (1-\rho)(1-\pi)v_{t+L+1} - v_{t+1} \geq 0$ and there is capacity on the flight, then we make fare class i available at time period t . Otherwise, we do not make fare class i available. In this section, we show that if we use the approximate policy, then we obtain at least half of the optimal total expected revenue. Our computational experiments indicate that the practical performance of the approximate policy can be substantially better than the theoretical guarantee of obtaining half of the optimal total expected revenue. We can compute the total expected revenue from the approximate policy by using a dynamic program similar to the one in (1), but the decision in the dynamic program is fixed by the approximate policy. We use $\hat{u}_{i,t}^x \in \{0, 1\}$ to denote the decision of the approximate policy for fare class i at time period t given that the remaining capacity on the flight is x . In particular, we have $\hat{u}_{i,t}^x = 1$ if and only if $f_i + (1-\rho)(1-\pi)v_{t+L+1} - v_{t+1} \geq 0$ and $x > 0$. Letting $\Theta_t(x, y_1, \dots, y_L)$ be the total expected revenue from the approximate policy over the time periods t, \dots, τ given that the state of the system at time period t is (x, y_1, \dots, y_L) , we can compute the value functions $\{\Theta_t(\cdot) : t \in T\}$ by solving the dynamic program

$$\begin{aligned} \Theta_t(x, y_1, y_2, \dots, y_L) = & \sum_{i \in N} \lambda_{i,t} \hat{u}_{i,t}^x \left[f_i + \mathbb{E} \left\{ \Theta_{t+1}(x-1 + D_t(y_L), B_t, y_1, \dots, y_{L-1}) \right\} \right] \\ & + \left\{ 1 - \sum_{i \in N} \lambda_{i,t} \hat{u}_{i,t}^x \right\} \mathbb{E} \left\{ \Theta_{t+1}(x + D_t(y_L), 0, y_1, \dots, y_{L-1}) \right\}, \quad (5) \end{aligned}$$

with the boundary condition that $\Theta_{\tau+1}(\cdot) = 0$. The total expected revenue obtained by the approximate policy is given by $\Theta_1(C, 0, \dots, 0)$.

In the next theorem, we show that the total expected revenue obtained by the approximate policy is at least half of the optimal total expected revenue.

Theorem 3 Noting that $V_1(C, 0, \dots, 0)$ is the optimal total expected revenue, we have $\Theta_1(C, 0, \dots, 0) \geq V_1(C, 0, \dots, 0)/2$.

Proof. We show that $\Theta_t(x, y_1, \dots, y_L) \geq v_t x + (1 - \pi)[v_{t+L} y_1 + \dots + v_{t+1} y_L]$ for all $t \in T$ by using induction over the time periods. In this case, using this inequality for the first time period and the state $(x, y_1, \dots, y_L) = (C, 0, \dots, 0)$, the desired result follows by noting that we have $\Theta_1(C, 0, \dots, 0) \geq C v_1 \geq \zeta^*/2 \geq V_1(C, 0, \dots, 0)/2$, where the last two inequalities use Lemmas 1 and 2. Letting $\tilde{V}_t(x, y_1, \dots, y_L) = v_t x + (1 - \pi)[v_{t+L} y_1 + \dots + v_{t+1} y_L]$ for notational brevity, we proceed to using induction over the time periods to show that $\Theta_t(x, y_1, \dots, y_L) \geq \tilde{V}_t(x, y_1, \dots, y_L)$ for all $t \in T$. By the boundary conditions of the dynamic programs in (4) and (5), we $v_{\tau+1} = \dots = v_{\tau+L+1} = 0$ and $\Theta_{\tau+1}(\cdot) = 0$, in which case, it immediately follows that the result holds at time period $\tau + 1$. Assuming that the result holds at time period $t + 1$, we show that the result holds at time period t . Noting that $\tilde{V}_t(x, y_1, \dots, y_L)$ is a linear function of (x, y_1, \dots, y_L) of the form $v_t x + (1 - \pi)[v_{t+L} y_1 + \dots + v_{t+1} y_L]$, we have

$$\begin{aligned} & \mathbb{E}\left\{\tilde{V}_{t+1}(x - 1 + D_t(y_L), B_t, y_1, \dots, y_{L-1}) - \tilde{V}_{t+1}(x + D_t(y_L), 0, y_1, \dots, y_{L-1})\right\} \\ &= \mathbb{E}\left\{-v_{t+1} + (1 - \pi)v_{t+L+1} B_t\right\} = -v_{t+1} + (1 - \rho)(1 - \pi)v_{t+L+1}, \end{aligned} \quad (6)$$

where the last equality uses the fact that B_t is a Bernoulli random variable with parameter $1 - \rho$. In this case, noting (5) and using the induction assumption, we obtain

$$\begin{aligned} \Theta_t(x, y_1, y_2, \dots, y_L) &\geq \sum_{i \in N} \lambda_{i,t} \hat{u}_{i,t}^x \left[f_i + \mathbb{E}\left\{\tilde{V}_{t+1}(x - 1 + D_t(y_L), B_t, y_1, \dots, y_{L-1})\right\} \right] \\ &\quad + \left\{1 - \sum_{i \in N} \lambda_{i,t} \hat{u}_{i,t}^x\right\} \mathbb{E}\left\{\tilde{V}_{t+1}(x + D_t(y_L), 0, y_1, \dots, y_{L-1})\right\} \\ &= \sum_{i \in N} \lambda_{i,t} \hat{u}_{i,t}^x \left[f_i - v_{t+1} + (1 - \rho)(1 - \pi)v_{t+L+1} \right] \\ &\quad + \mathbb{E}\left\{\tilde{V}_{t+1}(x + D_t(y_L), 0, y_1, \dots, y_{L-1})\right\}, \end{aligned} \quad (7)$$

where the equality uses (6). For the first expression on the right side of (7), note that $\hat{u}_{i,t}^x = 1$ if and only if $f_i - v_{t+1} + (1 - \rho)(1 - \pi)v_{t+L+1} \geq 0$ and $x > 0$. Thus, for any $x \in [0, C]$, we obtain

$$\begin{aligned} & \sum_{i \in N} \lambda_{i,t} \hat{u}_{i,t}^x \left[f_i - v_{t+1} + (1 - \rho)(1 - \pi)v_{t+L+1} \right] \\ &= \sum_{i \in N} \lambda_{i,t} \mathbf{1}(x > 0) \left[f_i - v_{t+1} + (1 - \rho)(1 - \pi)v_{t+L+1} \right]^+ \\ &\geq \sum_{i \in N} z_{i,t}^* \frac{x}{C} \left[f_i - v_{t+1} + (1 - \rho)(1 - \pi)v_{t+L+1} \right]^+ = (v_t - v_{t+1})x. \end{aligned} \quad (8)$$

In the chain of inequalities above, the first equality uses the fact that $\hat{u}_{i,t}^x = 1$ if and only if $f_i - v_{t+1} + (1 - \rho)(1 - \pi)v_{t+L+1} \geq 0$ and $x > 0$. The first inequality uses the fact that $\{z_{i,t}^* : i \in N, t \in T\}$ is

a feasible solution to problem (2) so that we have $\lambda_{i,t} \geq z_{i,t}^*$ and $\mathbf{1}(x > 0) \geq x/C$ for any $x \in [0, C]$. The second equality is by (4). For the second expression on the right side of (7), we have

$$\begin{aligned} & \mathbb{E}\left\{\tilde{V}_{t+1}(x + D_t(y_L), 0, y_1, \dots, y_{L-1})\right\} \\ &= \mathbb{E}\left\{v_{t+1}[x + D_t(y_L)] + (1 - \pi)[v_{t+L+1} 0 + v_{t+L} y_1 + \dots + v_{t+2} y_{L-1}]\right\}, \end{aligned}$$

where we use the definition of $\tilde{V}_t(x, y_1, \dots, y_L)$. Noting that $D_t(y_L)$ is a Bernoulli random variable with parameter $(1 - \pi)y_L$ and rearranging the expression, the equality above yields $\mathbb{E}\{\tilde{V}_{t+1}(x + D_t(y_L), 0, y_1, \dots, y_{L-1})\} = v_{t+1}x + (1 - \pi)[v_{t+L}y_1 + \dots + v_{t+2}y_{L-1} + v_{t+1}y_L]$. Using this equality and the inequality on (8) in the right side of (7), we obtain $\Theta_t(x, y_1, \dots, y_L) \geq (v_t - v_{t+1})x + v_{t+1}x + (1 - \pi)[v_{t+L}y_1 + \dots + v_{t+2}y_{L-1} + v_{t+1}y_L] = \tilde{V}_t(x, y_1, \dots, y_L)$, where the equality follows from the definition of $\tilde{V}_t(x, y_1, \dots, y_L)$. \square

Theorem 3 shows that the total expected revenue obtained by the approximate policy is at least half of the optimal total expected revenue. The proof of Theorem 3 is based on showing that $v_t x + (1 - \pi)[v_{t+L}y_1 + \dots + v_{t+1}y_L]$ is a lower bound on $\Theta_1(x, y_1, y_2, \dots, y_L)$. Note that we can interpret $v_t x + (1 - \pi)[v_{t+L}y_1 + \dots + v_{t+1}y_L]$ as an approximation to the value function $V_t(x, y_1, \dots, y_L)$. In particular, if we manage the availability of the fare classes on each seat separately, then v_t captures the optimal total expected revenue from a seat available at time period t . If the state of the system is (x, y_1, \dots, y_L) at time period t , then we have x seats available at this time period and the total expected revenue from each one of these seats is v_t , which yields the term $v_t x$ in the approximation $v_t x + (1 - \pi)[v_{t+L}y_1 + \dots + v_{t+1}y_L]$. Furthermore, if $y_1 = 1$, then we have one seat reserved for the customer who locked the fare at time period $t - 1$. This customer makes her ultimate purchase decision at time period $t + L - 1$. With probability $1 - \pi$, the customer ultimately does not purchase the ticket, in which case, the seat becomes available for other customers at the end of time period $t + L - 1$ and it can be used for a customer arriving at time period $t + L$. This reasoning yields the term $(1 - \pi)v_{t+L}y_1$ in the approximation $v_t x + (1 - \pi)[v_{t+L}y_1 + \dots + v_{t+1}y_L]$. A similar reasoning yields the other terms in the approximation $v_t x + (1 - \pi)[v_{t+L}y_1 + \dots + v_{t+1}y_L]$. It also turns out that the approximate policy is equivalent to approximating the value function $V_t(x, y_1, \dots, y_L)$ in (1) by using $\tilde{V}_t(x, y_1, \dots, y_L) = v_t x + (1 - \pi)[v_{t+L}y_1 + \dots + v_{t+1}y_L]$. To see this equivalence, by the same argument in (6), we have $\mathbb{E}\{\tilde{V}_{t+1}(x - 1 + D_t(y_L), B_t, y_1, \dots, y_{L-1})\} - \mathbb{E}\{\tilde{V}_{t+1}(x + D_t(y_L), 0, y_1, \dots, y_{L-1})\} = -v_{t+1} + (1 - \rho)(1 - \pi)v_{t+L+1}$. In this case, approximating $V_t(x, y_1, \dots, y_L)$ on the right side of (1) with $\tilde{V}_t(x, y_1, \dots, y_L)$, we obtain the problem

$$\max_{u \in \mathcal{U}(x)} \left\{ \sum_{i \in N} \lambda_{i,t} u_i \left[f_i + (1 - \rho)(1 - \pi)v_{t+L+1} - v_{t+1} \right] \right\} + \mathbb{E}\{\tilde{V}_{t+1}(x + D_t(y_L), 0, y_1, \dots, y_{L-1})\}.$$

Setting $u_i = 1$ if and only if $f_i + (1 - \rho)(1 - \pi)v_{t+L+1} - v_{t+1} \geq 0$ and $x > 0$ yields an optimal solution to the problem above, which is the decision rule used by the approximate policy.

5 Purchase Decisions Depending on Fare Classes

We can extend our results to deal with the case where $1 - \rho$, $1 - \pi$ and L depend on the fare class that is of interest to a customer. Consider the case where a customer with an interest in fare class i locks the fare with probability $1 - \rho_i$, makes her purchase decision after L_i time periods and ultimately decides not to purchase the ticket with probability $1 - \pi_i$. In this case, the state variable in our dynamic programming formulation needs to keep track of the fare classes that are of interest to the customers who locked the fare. We use $y_{i,\ell} \in \{0, 1\}$ to capture whether the customer arriving ℓ time periods ago is interested in fare class i and locks the fare. Using the vector $y_\ell = (y_{1,\ell}, \dots, y_{n,\ell}) \in \{0, 1\}^n$ and letting $L = \max\{L_i : i \in N\}$, we capture the state of the system at a time period by using (x, y_1, \dots, y_L) . We emphasize that y_ℓ in the state variable in this section is a vector. We continue using the vector $u = (u_1, \dots, u_n) \in \{0, 1\}^n$ to capture the decisions that we make at a time period, where $u_i = 1$ if and only if we make fare class i available.

We proceed to describing the changes that we need to do in the dynamic program in (1). Given that there is a customer arrival with an interest in fare class i at time period t and this fare class is available, we use the random variable $B_{i,t}$ to capture whether this customer locks the fare. In particular, $B_{i,t}$ is a Bernoulli random variable with parameter $1 - \rho_i$, taking value 1 if and only if the customer locks the fare. As a function of whether the customer arriving at time period $t - L_i$ is interested in fare class i and locks the fare, we use the random variable $D_{i,t}(y_{i,L_i})$ to capture whether the customer with an interest in fare class i ultimately decides not to purchase the ticket at time period t . In particular, $D_{i,t}(y_{i,L_i})$ is a Bernoulli random variable with parameter $(1 - \pi_i) y_{i,L_i}$, taking value 1 if and only if the customer ultimately decides not to purchase. In this case, using $e_i \in \mathbb{R}_+^n$ to denote a unit vector with a one in the i -th component, all we need to do is to replace B_t and $D_t(y_L)$ in the dynamic program in (1) with $B_{i,t} e_i$ and $\sum_{i \in N} D_{i,t}(y_{i,L_i})$.

The only change that we need to do in the linear program in (2) is to replace $1 - \rho$, $1 - \pi$ and L in the second constraint with $1 - \rho_i$, $1 - \pi_i$ and L_i . In this case, Lemma 1 continues to hold. Similarly, the only change that we need to do in the dynamic program in (4) is to replace $1 - \rho$, $1 - \pi$ and L with $1 - \rho_i$, $1 - \pi_i$ and L_i . In this case, Lemma 2 continues to hold. In the approximate policy, we make fare class i available at time period t if and only if $f_i + (1 - \rho_i)(1 - \pi_i)v_{t+L_i+1} - v_{t+1} \geq 0$ and there is capacity on the flight. To compute the total expected revenue obtained by the approximate policy, we can use the dynamic program in (5) after replacing B_t and $D_t(y_L)$ with $B_{i,t} e_i$ and $\sum_{i \in N} D_{i,t}(y_{i,L_i})$. In this case, Theorem 3 continues to hold. Thus, the total expected revenue from the optimal policy is at least half of the optimal total expected revenue when $1 - \rho$, $1 - \pi$ and L depend on the fare class that is of interest to a customer, but the proof of Theorem 3 gets substantially more tedious since we have to deal with vectors y_1, \dots, y_L .

In our model, each customer arrives with an interest in a fare class and we pick the fare classes that are available. Our results continue to hold when we pick the price for the ticket and the customers decide to purchase based on the price. We give this extension in the next section.

6 Capacity Control by Using Pricing Decisions

We describe the changes that need to be made in our results when we pick the price for the ticket and the purchase decisions of the customers depend on the price.

6.1 Problem Formulation

We have n possible price levels indexed by $N = \{1, \dots, n\}$. The capacity on the flight is C . The selling horizon has τ time periods indexed by $T = \{1, \dots, \tau\}$. For notational brevity, we assume that there is one customer arrival at each time period. If we charge price level i for the ticket, then the customer arriving at time period t purchases the ticket or locks the fare with probability $\lambda_{i,t}$. We can interpret $\lambda_{i,t}$ as the probability that the willingness to pay of the customer arriving at time period t exceeds price level i . We assume that there is one price level $\phi \in N$ such that $\lambda_{\phi,t} = 0$ for all $t \in T$. In this case, if there is no capacity available on the flight, then we can charge price level ϕ to ensure that the customers arriving into the system do not purchase the ticket or lock the fare. Given that we charge price level i at a time period, if the willingness to pay of the arriving customer exceeds price level i , then the customer purchases the ticket with probability ρ by paying the price r_i corresponding to price level i , whereas the customer locks the fare with probability $1 - \rho$ by paying the fee h . A customer with a locked fare makes her purchase decision after L time periods. With probability π , the customer purchases the ticket, in which case, she pays the price corresponding to the fare she locked. With probability $1 - \pi$, the customer decides not to purchase the ticket. So, if we charge price level i and the willingness to pay of a customer exceeds price level i , then we obtain an expected revenue of $f_i = \rho r_i + (1 - \rho)(h + \pi r_i)$. We continue using (x, y_1, \dots, y_L) to capture the state of the system, where x and y_ℓ are as defined in Section 1. To capture the decisions that we make at a time period, we use the vector $u = (u_1, \dots, u_n) \in \{0, 1\}^n$, where $u_i = 1$ if and only if we charge price level i . At any time period, we can charge only one price level. Furthermore, if there is no capacity available on the flight, then we need to charge a price level that ensures that no customer purchases the ticket or locks the fare. Therefore, the set of feasible decisions at time period t is given by $\mathcal{U}_t(x) = \{u \in \{0, 1\}^n : \sum_{i \in N} u_i = 1, \lambda_{i,t} u_i \leq x \forall i \in N\}$. We continue using the random variables B_t and $D_t(y_L)$ as defined in Section 1. In this case, letting $V_t(x, y_1, \dots, y_L)$ be the optimal total expected revenue over time periods t, \dots, T given that the state of the system at time period t is (x, y_1, \dots, y_L) , we can compute the value functions $\{V_t(\cdot) : t \in T\}$ by solving the dynamic program

$$V_t(x, y_1, y_2, \dots, y_L) = \max_{u \in \mathcal{U}_t(x)} \left\{ \sum_{i \in N} \lambda_{i,t} u_i \left[f_i + \mathbb{E} \left\{ V_{t+1}(x - 1 + D_t(y_L), B_t, y_1, \dots, y_{L-1}) \right\} \right] \right. \\ \left. + \left[1 - \sum_{i \in N} \lambda_{i,t} u_i \right] \mathbb{E} \left\{ V_{t+1}(x + D_t(y_L), 0, y_1, \dots, y_{L-1}) \right\} \right\}, \quad (9)$$

with the boundary condition that $V_{\tau+1}(\cdot) = 0$. This dynamic program is identical to the one in (1), but the set of feasible decisions $\mathcal{U}_t(x)$ and the interpretations of $\lambda_{i,t}$ and u_i are different.

6.2 Upper Bound on the Optimal Total Expected Revenue

We formulate a linear program to obtain an upper bound on the optimal total expected revenue. We use the decision variable $z_{i,t}$ to capture the probability of charging price level i at time period t . Using the vector $z = \{z_{i,t} : i \in N, t \in T\}$, we consider the linear program

$$\max_{z \in \mathfrak{R}_+^{n \times \tau}} \left\{ \sum_{t \in T} \sum_{i \in N} f_i \lambda_{i,t} z_{i,t} : \sum_{i \in N} z_{i,t} = 1 \quad \forall t \in T, \right. \\ \left. C - \sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t) \lambda_{i,\kappa} z_{i,\kappa} + \sum_{\kappa \in T} \sum_{i \in N} \mathbf{1}(\kappa \leq t - L) (1 - \rho) (1 - \pi) \lambda_{i,\kappa} z_{i,\kappa} \geq 0 \quad \forall t \in T \right\}. \quad (10)$$

The willingness to pay of a customer arriving at time period t exceeds price level i with probability $\lambda_{i,t}$, in which case, the expected revenue obtained from the customer is f_i . Thus, the objective function accounts for the total expected revenue over the selling horizon. The first constraint ensures that we charge some price level at time period t . The second constraint is similar to the second constraint in problem (2) and ensures that the net total expected capacity consumed over time periods $1, \dots, t$ cannot exceed the capacity available on the flight. Using the same argument in the proof of Lemma 1, we can show that the optimal objective value of problem (10) is an upper bound on the optimal total expected revenue. The only difference is that we let $Z_{i,t}^* = 1$ if we charge price level i at time period t under the optimal policy.

6.3 Approximations to the Value Functions

We can approximate the value functions in the dynamic program in (9) by making the pricing decisions for each seat separately. Using $\{z_{i,t}^* : i \in N, t \in T\}$ to denote an optimal solution to problem (10), we compute the value functions $\{v_t : t \in T\}$ through the dynamic program

$$v_t = \sum_{i \in N} \frac{\lambda_{i,t}}{C} z_{i,t}^* \max \left\{ f_i + (1 - \rho) (1 - \pi) v_{t+L+1}, v_{t+1} \right\} + \left\{ 1 - \sum_{i \in N} \frac{\lambda_{i,t}}{C} z_{i,t}^* \right\} v_{t+1} \\ = \sum_{i \in N} \frac{\lambda_{i,t}}{C} z_{i,t}^* \left[f_i + (1 - \rho) (1 - \pi) v_{t+L+1} - v_{t+1} \right]^+ + v_{t+1}. \quad (11)$$

In the dynamic program above, we charge price level i at time period t with probability $z_{i,t}^*$. The willingness to pay for a customer arriving at time period t exceeds price level i with probability $\lambda_{i,t}/C$. Noting the maximum operator in (11), even if the willingness to pay for a customer exceeds the price level we charge, we can deny the customer from purchasing the ticket or locking the fare. In reality, when we control the capacity on the flight by making pricing decisions, we do not have the option of denying a customer from purchasing the ticket or locking the fare, but we use the value functions that are obtained through the dynamic program in (11) only to construct an approximate policy. Using the same argument in the proof of Lemma 2, we can show that Cv_1 is least half of the optimal objective value of problem (10).

6.4 Performance Guarantee for the Approximate Policy

We let $\{v_t : t \in T\}$ be computed through the dynamic program in (11). In the approximate policy, we let the vector $\hat{u}_t^x = (\hat{u}_{1,t}^x, \dots, \hat{u}_{n,t}^x) \in \{0,1\}^n$ be the optimal solution to the problem $\max_{u \in \mathcal{U}_t(x)} \left\{ \sum_{i \in N} \lambda_{i,t} u_i [f_i + (1-\rho)(1-\pi)v_{t+L+1} - v_{t+1}] \right\}$. Since $\hat{u}_t^x \in \mathcal{U}_t(x)$, noting the definition of $\mathcal{U}_t(x)$, there is a unique price level i such that $\hat{u}_{i,t}(x) = 1$, which corresponds to the price level that we charge in the approximate policy. Note that it is simple to obtain the optimal solution to the problem $\max_{u \in \mathcal{U}_t(x)} \left\{ \sum_{i \in N} \lambda_{i,t} u_i [f_i + (1-\rho)(1-\pi)v_{t+L+1} - v_{t+1}] \right\}$. If $x = 0$, then we set $\hat{u}_{\phi,t}^x = 1$ and $\hat{u}_{i,t}^x = 0$ for all $i \in N \setminus \{\phi\}$. If $x > 0$, then we find $k \in N$ such that $k = \arg \max_{i \in N} \left\{ \lambda_{i,t} [f_i + (1-\rho)(1-\pi)v_{t+L+1} - v_{t+1}] \right\}$ and set $\hat{u}_{k,t}^x = 1$ and $\hat{u}_{i,t}^x = 0$ for all $i \in N \setminus \{k\}$. We let $\Theta_t(x, y_1, \dots, y_L)$ be the total expected revenue from the approximate policy over time periods t, \dots, τ given that the state of the system at time period t is (x, y_1, \dots, y_L) . We can compute the value functions $\{\Theta_t(\cdot) : t \in T\}$ by solving the dynamic program

$$\begin{aligned} \Theta_t(x, y_1, y_2, \dots, y_L) = & \sum_{i \in N} \lambda_{i,t} \hat{u}_{i,t}^x \left[f_i + \mathbb{E} \left\{ \Theta_{t+1}(x-1 + D_t(y_L), B_t, y_1, \dots, y_{L-1}) \right\} \right] \\ & + \left\{ 1 - \sum_{i \in N} \lambda_{i,t} \hat{u}_{i,t}^x \right\} \mathbb{E} \left\{ \Theta_{t+1}(x + D_t(y_L), 0, y_1, \dots, y_{L-1}) \right\}, \end{aligned} \quad (12)$$

with the boundary condition that $\Theta_{\tau+1}(\cdot) = 0$. Note that the dynamic programs in (12) and (5) are identical, but $\hat{u}_{i,t}^x$ has a different interpretation in each dynamic program.

Using the same argument in the proof of Theorem 3, we can show that the total expected revenue obtained by the approximate policy is at least half of the optimal total expected revenue. The only slight deviation occurs when we obtain a chain of inequalities analogous to the one in (8). In particular, let $k \in N$ be such that $\hat{u}_{k,t}^x = 1$. By the discussion at the beginning of this section, we have $\lambda_{k,t} [f_k + (1-\rho)(1-\pi)v_{t+L+1} - v_{t+1}] \geq \mathbf{1}(x > 0) \lambda_{i,t} [f_i + (1-\rho)(1-\pi)v_{t+L+1} - v_{t+1}]$ for all $i \in N$. Since $\lambda_{\phi,t} = 0$, the last inequality also implies $\lambda_{k,t} [f_k + (1-\rho)(1-\pi)v_{t+L+1} - v_{t+1}] \geq 0$ so that $\lambda_{k,t} [f_k + (1-\rho)(1-\pi)v_{t+L+1} - v_{t+1}] \geq \mathbf{1}(x > 0) \lambda_{i,t} [f_i + (1-\rho)(1-\pi)v_{t+L+1} - v_{t+1}]^+$ for all $i \in N$. Since $\{z_{i,t}^* : i \in N, t \in T\}$ is a feasible solution to problem (10), we have $\sum_{i \in N} z_{i,t}^* = 1$, in which case, multiplying the last inequality by $z_{i,t}^*$ and adding over all $i \in N$, we get $\lambda_{k,t} [f_k + (1-\rho)(1-\pi)v_{t+L+1} - v_{t+1}] \geq \sum_{i \in N} \lambda_{i,t} z_{i,t}^* \mathbf{1}(x > 0) [f_i + (1-\rho)(1-\pi)v_{t+L+1} - v_{t+1}]^+$. Thus, since $\hat{u}_{k,t}^x = 1$ and $\hat{u}_{i,t}^x = 0$ for all $i \in N \setminus \{k\}$, for any $x \in [0, C]$, we obtain

$$\begin{aligned} \sum_{i \in N} \lambda_{i,t} \hat{u}_{i,t}^x [f_i - v_{t+1} + (1-\rho)(1-\pi)v_{t+L+1}] &= \lambda_{k,t} \hat{u}_{k,t}^x [f_k - v_{t+1} + (1-\rho)(1-\pi)v_{t+L+1}] \\ &= \sum_{i \in N} \lambda_{i,t} z_{i,t}^* \mathbf{1}(x > 0) [f_i - v_{t+1} + (1-\rho)(1-\pi)v_{t+L+1}]^+ \\ &\geq \sum_{i \in N} \lambda_{i,t}^* z_{i,t}^* \frac{x}{C} [f_i - v_{t+1} + (1-\rho)(1-\pi)v_{t+L+1}]^+ = (v_t - v_{t+1})x, \end{aligned}$$

where the first equality is by the fact that $\hat{u}_{i,t}^x = 1$ if and only if $i = k$, the inequality is by the fact that $\mathbf{1}(x > 0) \geq x/C$ for any $x \in [0, C]$ and the last equality is by (11).

7 Numerical Study

In this section, we give a numerical study to demonstrate that the practical performance of the approximate policy can be significantly better than the theoretical guarantee.

7.1 Numerical Setup

In all of our test problems, the number of fare classes is $n = 4$, the number of time periods in the selling horizon is $\tau = 300$, the capacity on the flight is $C = 100$ and the fares associated with the four fare classes are $(r_1, r_2, r_3, r_4) = (1000, 750, 500, 250)$. For the fee h of locking the fare, we use $h \in \{40, 80\}$. Noting the discussion in the previous section, we allow the probability $1 - \rho$ of locking the fare to depend on the fare class that is of interest to the customer. In particular, we use $1 - \rho_i$ to denote the probability that a customer with an interest in fare class i locks the fare. We use $(1 - \rho_1, 1 - \rho_2, 1 - \rho_3, 1 - \rho_4) \in \{(0.15, 0.2, 0.25, 0.3), (0.4, 0.45, 0.5, 0.55)\}$. The number of time periods L that a customer with a locked fare waits until she makes a purchase decision and the probability π that a customer with a locked fare ultimately purchases the ticket do not depend on the fare class that is of interest to the customer. We use $L \in \{25, 50\}$ and $\pi \in \{0.4, 0.7\}$. To come up with the probability $\lambda_{i,t}$ of a customer arrival with an interest in fare class i at time period t , we let $\gamma_{1,t} = 0.1 + \mathbf{1}(t \geq \frac{\tau}{3}) \times 0.1 + \mathbf{1}(t \geq \frac{2\tau}{3}) \times 0.2$ for fare class 1 so that we have $\gamma_{1,t} = 0.1$ for all $t \in [1, \frac{\tau}{3})$, $\gamma_{1,t} = 0.2$ for all $t \in [\frac{\tau}{3}, \frac{2\tau}{3})$ and $\gamma_{1,t} = 0.4$ for all $t \in [\frac{2\tau}{3}, \tau]$. Similarly, we let $\gamma_{2,t} = 0.2 + \mathbf{1}(t \geq \frac{2\tau}{3}) \times 0.1$ for fare class 2, $\gamma_{3,t} = 0.3 - \mathbf{1}(t \geq \frac{2\tau}{3}) \times 0.1$ for fare class 3 and $\gamma_{4,t} = 0.4 - \mathbf{1}(t \geq \frac{\tau}{3}) \times 0.1 - \mathbf{1}(t \geq \frac{2\tau}{3}) \times 0.2$ for fare class 4. Note that $\gamma_{1,t}$ and $\gamma_{2,t}$ are increasing in t , whereas $\gamma_{3,t}$ and $\gamma_{4,t}$ are decreasing in t . A customer with an interest in fare class i immediately purchases the ticket with probability ρ_i , whereas she locks the fare with probability $1 - \rho_i$ and ultimately purchases after L time periods with probability π . Therefore, if we make all of the fare classes available at all of the time periods in the selling horizon, then the total expected capacity consumption on the flight is $\sum_{t \in T} \sum_{i \in N} \lambda_{i,t} (\rho_i + (1 - \rho_i) \pi)$, which implies that the load factor of the flight is $\sum_{t \in T} \sum_{i \in N} \lambda_{i,t} (\rho_i + (1 - \rho_i) \pi) / C$. In our test problems, we set $\lambda_{i,t} = \gamma_{i,t} / \Delta$ for some value of Δ to achieve a load factor of 1.2. Noting that $\gamma_{1,t}$ and $\gamma_{2,t}$ are increasing in t , whereas $\gamma_{3,t}$ and $\gamma_{4,t}$ are decreasing in t and fare classes 1 and 2 are the more expensive fare classes, whereas fare classes 3 and 4 are the less expensive fare classes, our numerical setup corresponds to the situation where the demand for the more expensive fare classes tend to arrive more frequently later in the selling horizon, whereas the demand for the less expensive fare classes tend to arrive more frequently earlier in the selling horizon

For notational brevity, we let $1 - \rho = (1 - \rho_1, 1 - \rho_2, 1 - \rho_3, 1 - \rho_4)$, $(1 - \rho)^L = (0.15, 0.2, 0.25, 0.3)$ and $(1 - \rho)^H = (0.4, 0.45, 0.5, 0.55)$, where the superscripts L and H stand for low and high. Varying $(h, 1 - \rho, L, \pi)$ so that $h \in \{40, 80\}$, $1 - \rho \in \{(1 - \rho)^L, (1 - \rho)^H\}$, $L \in \{25, 50\}$ and $\pi \in \{0.4, 0.7\}$, we obtain 16 test problems in our numerical setup.

7.2 Numerical Results

We summarize our numerical results in Table 1. The first column in this table shows the parameters for each test problem by using the tuple $(h, 1 - \rho, L, \pi)$. The second column shows the total expected revenue obtained by the approximate policy. We estimate the total expected revenue obtained by the approximate policy by simulating the performance of this policy over 1,000 sample paths. The third column shows the upper bound on the optimal total expected revenue provided by the optimal objective value of problem (2). The fourth column shows the percent gap between the upper bound on the optimal total expected revenue and the total expected revenue obtained by the approximate policy, which gives an estimate of the optimality gap of the approximate policy. This estimate of the optimality gap is somewhat pessimistic in the sense that we compare the total expected revenue obtained by the approximate policy with an upper bound on the optimal total expected revenue, rather than the optimal total expected revenue itself. The fifth column shows the CPU seconds required to solve the linear program in (2), whereas the sixth column shows the CPU seconds required to solve the dynamic program in (4).

The results in Table 1 indicate that the approximate policy performs quite well. Over all of our test problems, the average gap between the upper bound on the optimal total expected revenue and the total expected revenue obtained by the approximate policy is about 3.23%. The state variable in the dynamic program in (1) has $L + 1$ dimensions, where L is the number of time periods that a customer with a locked fare waits to make her ultimate purchase decision. Therefore, the number of dimensions of the state variable gets larger as L gets larger. The optimality gaps of the approximate policy remain quite stable as L gets larger. The CPU seconds to implement the approximate policy are quite reasonable. The main bulk of the computational work requires solving the linear program in (2), which can be carried out in a fraction of a second.

8 Conclusions

We considered a revenue management problem on a single flight leg, where the customers have the option of locking an available fare. The dynamic programming formulation of the problem requires a high dimensional state variable. We gave an approximate policy that is guaranteed to obtain at least half of the optimal total expected revenue. Our approach is based on using a linear programming approximation to decompose the problem by the seats on the flight and managing the availability of the fare classes for each seat separately. For future research, one can try to use this general approach to obtain approximate policies for other dynamic programs with high dimensional state variables. Linear programming approximations are often used for constructing approximate policies for managing resources in various settings, but we are not aware of any constant factor performance guarantees for the approximate policies constructed by using solely these linear programming approximations. It is exciting to see what problems will be amenable to leveraging a linear programming approximation to decompose the problem by the resources and to obtain approximate policies through simpler dynamic programs.

Prob. Params. ($h, 1 - \rho, L, \pi$)	App. Rev.	Rev. Bnd.	Perc. Gap	Lin. Secs.	Dyn. Secs.
$(40, (1 - \rho)^L, 25, 0.4)$	67,487	69,759	3.26	0.023	0.004
$(40, (1 - \rho)^L, 25, 0.7)$	66,866	69,177	3.34	0.024	0.004
$(40, (1 - \rho)^L, 50, 0.4)$	67,272	69,539	3.26	0.023	0.006
$(40, (1 - \rho)^L, 50, 0.7)$	66,744	69,074	3.37	0.024	0.006
$(40, (1 - \rho)^H, 25, 0.4)$	69,062	71,196	3.00	0.028	0.005
$(40, (1 - \rho)^H, 25, 0.7)$	68,145	70,426	3.24	0.023	0.004
$(40, (1 - \rho)^H, 50, 0.4)$	68,019	70,255	3.18	0.023	0.007
$(40, (1 - \rho)^H, 50, 0.7)$	67,760	70,043	3.26	0.027	0.006
$(80, (1 - \rho)^L, 25, 0.4)$	68,213	70,472	3.21	0.024	0.005
$(80, (1 - \rho)^L, 25, 0.7)$	67,557	69,861	3.30	0.029	0.005
$(80, (1 - \rho)^L, 50, 0.4)$	67,942	70,242	3.28	0.024	0.006
$(80, (1 - \rho)^L, 50, 0.7)$	67,437	69,753	3.32	0.023	0.007
$(80, (1 - \rho)^H, 25, 0.4)$	71,388	73,650	3.07	0.024	0.005
$(80, (1 - \rho)^H, 25, 0.7)$	70,229	72,543	3.19	0.023	0.005
$(80, (1 - \rho)^H, 50, 0.4)$	70,252	72,601	3.24	0.025	0.006
$(80, (1 - \rho)^H, 50, 0.7)$	69,794	72,124	3.23	0.025	0.007

Table 1: Performance of the approximate policy.

References

- Adelman, D. (2007), ‘Dynamic bid-prices in revenue management’, *Operations Research* **55**(4), 647–661.
- Aydin, N., Birbil, S. I., Frenk, J. B. G. and Noyan, N. (2013), ‘Single-leg airline revenue management with overbooking’, *Transportation Science* **47**(4), 560–583.
- Aydin, N., Birbil, S. I. and Topaloglu, H. (2016), ‘Delayed purchase options in single-leg revenue management’, *Transportation Science* (to appear).
- Balseiro, S. R., Gallego, G., Gocmen, C. and Phillips, R. (2011), Revenue management of consumer options for sporting events, Technical report, Columbia University, New York, NY.
- Bertsimas, D. and Popescu, I. (2003), ‘Revenue management in a dynamic network environment’, *Transportation Science* **37**(3), 257–277.
- Chen, Z.-L. and Chen, M. (2016), Dynamic pricing with a fare-lock option. *INFORMS Annual Conference*, Nashville, TN.
- 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., Iyengar, G., Phillips, R. and Dubey, A. (2004), Managing flexible products on a network, Computational Optimization Research Center Technical Report TR-2004-01, Columbia University.
- Gallego, G., Kou, S. G., and Phillips, R. (2008), ‘Revenue management of callable products’, *Management Science* **54**(3), 550–564.
- Gallego, G., Li, A., Truong, V.-A. and Wang, X. (2015), Online personalized resource allocation with customer choice, Technical report, Columbia University, New York, NY.
- Gallego, G. and Phillips, R. (2004), ‘Revenue management with flexible products’, *Manufacturing & Service Operations Management* **6**(4), 321–337.
- Gallego, G. and Stefanescu, C. (2009), Upgrades, upsells and pricing in revenue management, Technical report, Columbia University.
- 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), ‘Overbooking with substitutable inventory classes’, *Operations Research* **52**(1), 83–104.
- Kunnumkal, S. (2014), ‘Randomization approaches for network revenue management with customer choice behavior’, *Production and Operations Management* **23**(9), 1617–1633.
- Kunnumkal, S. and Talluri, K. (2016), ‘On a piecewise-linear approximation for network revenue management’, *Mathematics of Operations Research* **41**(1), 72–91.
- Kunnumkal, S., Talluri, K. and Topaloglu, H. (2012), ‘A randomized linear programming method for network revenue management with product-specific no-shows’, *Transportation Science* **46**(1), 90–108.
- Kunnumkal, S. and Topaloglu, H. (2010), ‘A new dynamic programming decomposition method for the network revenue management problem with customer choice behavior’, *Production and Operations Management* **19**, 575–590.
- Lan, Y., Ball, M. O. and Karaesmen, I. Z. (2011), ‘Regret in overbooking and fare-class allocation for single leg’, *Manufacturing & Service Operations Management* **13**(2), 194–208.
- Liu, Q. and van Ryzin, G. (2008), ‘On the choice-based linear programming model for network revenue management’, *Manufacturing & Service Operations Management* **10**(2), 288–310.
- Meissner, J. and Strauss, A. K. (2012), ‘Network revenue management with inventory-sensitive bid prices and customer choice’, *European Journal of Operational Research* **216**, 459–468.
- Meissner, J., Strauss, A. and Talluri, K. (2012), ‘An enhanced concave program relaxation for choice network revenue management’, *Production and Operations Management* **22**(1), 71–87.
- 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.
- Topaloglu, H. (2009), ‘Using Lagrangian relaxation to compute capacity-dependent bid-prices in network revenue management’, *Operations Research* **57**(3), 637–649.
- Vossen, T. M. V. and Zhang, D. (2015), ‘Reductions of approximate linear programs for network revenue management’, *Operations Research* **63**(6), 1352–1371.
- Wang, X., Truong, V.-A. and Bank, D. (2015), Online advance admission scheduling for services with customer preferences, Technical report, Columbia University, New York, NY.
- Zhang, D. and Adelman, D. (2009), ‘An approximate dynamic programming approach to network revenue management with customer choice’, *Transportation Science* **42**(3), 381–394.