Proceedings of the 2019 Winter Simulation Conference
N. Mustafee, K.-H.G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, eds.

MITIGATING CUSTOMERS' DOWNSELL RISK FOR SINGLE-LEG REVENUE MANAGEMENT WITH DEMAND DEPENDENCIES

Lin Li

Department of Systems and Industrial Engineering
Kennesaw State University
650 Clair Harris Road
Marietta, GA 30060, USA

Yuan Zhou

Department of Industrial, Manufacturing and Systems Engineering
University of Texas at Arlington
701 South Nedderman Drive
Arlington, TX 76019, USA

ABSTRACT

With increased complexity of customers choice behaviors, practical optimization approaches often involve decomposing a network revenue management problem into multiple single-leg problems. While dynamic programming approaches can be used to solve single-leg problems exactly, they are not scalable and require precise information about the customers' arrival rates. On the other hand, the traditional heuristics are often static which do not explicitly consider the remaining time horizon in the optimization. This motivates us to find scalable and efficient dynamic heuristics that work well with the complex customers choice models. We develop two expected marginal seat revenue type heuristics for the single-leg dynamic revenue management problems in airline industry and evaluate its performances using Monte Carlo simulation. The initial simulation results indicate that our proposed heuristics are computationally efficient and fairly robust. This study provides a foundation for potential future extensions to solve larger network problems.

1 INTRODUCTION

Revenue management (RM) is a discipline originated in the airline industry, following the deregulation of the US air travel in 1970s. It applies analytical techniques, such as stochastic optimization and dynamic programming, to manage demand either through price or product availability so that the revenue performance is optimized. Today, the theory and practice of RM have been widespread in many other areas such as cruise lines, care rentals, hospitality and so on.

In airline industry, the revenues of airline companies can be maximized through dynamically determining the best set of fares to offer to customers. Here, the RM involves bucketing airfares into multiple classes (i.e., prices from low to high) and managing seat inventory in each fare bucket dynamically. The customer demand, and seat availability as well as time horizon remaining prior to the flight departure are important factors in RM decision making. Most of research prior to 2004 is based on independent demand assumption, that is, the customer demand of any product is independent of the product availability (see Lautenbacher and Stidham (1999) and Cooper (2002)). In this context, the products refer to the air tickets sold at different fares. In other words, customers will not substitute one product for another if their preferred one is not available. However, with the rapid growth of low-cost carriers, e.g., Spirit airlines, Frontier airlines, and the proliferation of online travel search engines, e.g., Expedia, Priceline, customers' choice of product is likely to be influenced by the price and product availability. Therefore, the independent demand assumption is no longer valid nor reasonable as indicated by Joyner (2007).

Given the updated demand forecasts, seat availability and time prior to flight departure, airlines constantly change their product availability by opening, closing, and reopening of fare classes. All of these changes are likely to trigger fare shopping on the consumer side based on the probability of prices going up or down prior to flight departure. This new area of RM with dependent demand that incorporates customer choice behavior has seen significant development in recent years. Four related effects, upsell, downsell, recapture and competition, need to be considered in airline RM in order to determine the optimal solution with customer choice behavior. To illustrate the choice behavior, let's consider a simple example of a direct flight from Salt Lake City (SLC) to Atlanta (ATL). Figure 1 shows the different customer purchase behaviors that affect revenue maximization. Tickets on this flight can be sold either at a low fare of $250 or a high fare at $350. The customers who purchased the low fare are called low-fare customers and can be further divided into two groups, group A consists of customers who purchase the $250 fare and group B are the ones who prefer $250 but will purchase $350 fare if the low fare is not available. The group B's behavior is called upsell, which occurs when a lower fare customer purchase a higher fare on the same flight if the lower class is unavailable for sale. In general, the only barrier for customer to upsell is the product price. Similarly, the customers who normally purchase the $350 fare are called high-fare customers and a fraction them will downsell to low-fare if the low fare $250 is available. Downsell happens when a higher fare customer purchases a lower fare that is also available for sale. In a mild restriction environment where very limited restrictions are imposed on the discount fare, downsell effects can be quite pronounced resulting in significant revenue dilution for airlines. Recapture occurs when customers purchase an alternative flight from the same carrier if their preferred flight is not available. Demand dependencies are further complicated by the presence of market competitors. In many cases where an airline closes discount fares to promote upsell to higher fares, the competing carriers may continue to offer discount alternatives, and these alternatives impact the airline's ability to achieve upsell. A two-carrier upsell simulation study reported by KLM Royal Dutch Airlines, Hartmans (2006) demonstrated this negative effect.

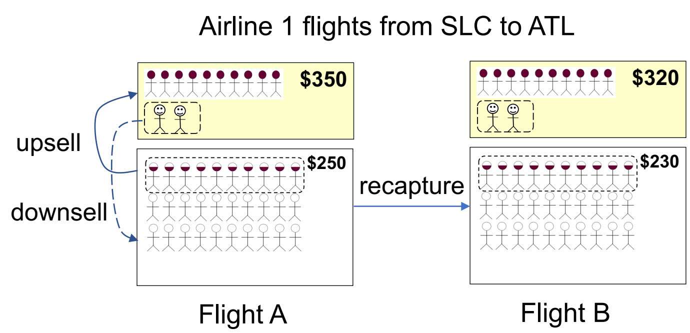

Figure 1: Example of usell, downsell and recapure.

While managing airlines' revenues over a network of flights is at the heart of RM, solving single-leg(a single flight resource) problems is still crucial because (1) the network RM with customer choice model is intractable for realistically-sized number of flights in the network, and hence, the methods require solving a set of single-leg problems are frequently applied; (2) some small airline companies, like charter flight companies, accept booking requests only for single-leg itineraries.

By incorporating time element into optimization control, dynamic RM models allow an arbitrary order of demand arrivals. With the assumption of Markovian demand and sufficiently fine discretization of time, a dynamic model determines whether to accept or reject a particular demand by examining the remaining capacity and remaining sales horizon. In this case, the decision variables depend on the state of current system. In the light of customer choice behavior, the demand for each fare product is a function of the set of fare products available to the customers.

In this paper, we focus on the single-leg dynamic RM with dependent demand as presented in the seminal publication by Talluri and van Ryzin (2004). Our work expands the current understanding of this problem by proposing a new heuristic that are computationally efficient and applicable in practice. Customers choice behavior is modeled using a multinomial logit (MNL) model which allows estimation of market share for each of the available airline flight fare classes as well as the subsequent demand shifts arising from fare class closures as described by van Ryzin and Vulcano (2008) and Gallego et al. (2015). Instead of solving the dynamic program exactly as in Talluri and van Ryzin (2004), we develop heuristics to approximate the marginal seat revenue at each state of the system to mitigate the downsell risk and avoid revenue dilution. The initial numerical study indicates that our proposed heuristics dominate the existing EMSR-b type heuristic that has been frequently used in the dynamic case by aggregating the total demands and averaging fares.

The rest of the paper is organized as follows: a brief literature review is presented in section 2, followed by problem formulation in section 3. In section 4, we describe two heuristics and discuss some of their properties. Then, we present the results of numerical simulation in section 5 and finally conclude with a summary in section 6.

2 LITERATURE REVIEW

Our paper is mainly related to three main streams of literature: 1) dynamic revenue management, 2) multinomial logit (MNL) customer choice model, and 3) expected marginal seat revenue (EMSR) type heuristic. Dynamic single-leg model was first analyzed by Lee and Hersh (1993) who study a discrete time Markov model with arbitrary order of demand arrivals. Lautenbacher and Stidham (1999) later provide further analysis of both static and dynamic models. Liang (1999) study the dynamic model in a continuous time version. All these work share a common assumption that the demand for each fare product is independent of the fare classes that are available to the customers.

To overcome the deficiency of independent demand assumption, many researchers have started to address this issue by incorporating customer choice behavior into their models. Andersson (1998) and Algers and Besser (2001) apply logit choice model to estimate upsell and recapture effects at Scandinavian Airline Systems. A milestone work is done by Talluri and van Ryzin (2004) who formulate a dynamic programming (DP) model of single-leg RM with a general discrete customer choice model. They provide structural property of the optimal solution and identify that the optimal policy can be characterized by a sequence of 'efficient' sets. They also develop an estimation procedure based on expectation-maximization (EM) method to estimate the choice-model parameters. Feldman and Topaloglu (2017) also investigate the single-leg RM problem with Markov Chain as the underlying customer choice model. They reach the same conclusion that the optimal policy can be presented by protection levels.

There are many studies on the second theme of modeling customer choice behavior when facing multiple substitutable products. The MNL model as described by Ben-Akiva and Lerman (1994) is derived from a random utility theory and has been used to model discrete choices. The customers utility for each product and subsequent choice probabilities are measured by certain probability distribution. The MNL models allow researchers to study various operational decisions with varying customer behaviors. Readers can refer to Dong et al. (2009) and Wang and Wang (2017) and references therein for more details. In airline RM, the MNL demand model allows estimation of market share for each of the available airline flight fare class alternatives as well as the resulting demand shifts arising from class closures. As reported by Mishra et al. (2005), MNL model can be calibrated using air shopping and purchase activity from a leading travel reservations system. A recent review by Strauss et al. (2018) offers an excellent overview of theory and recent development for choice-based revenue management.

The final stream of literature is on EMSR type algorithms. The dynamic heuristics we proposed in this paper share some similarities with the approximating marginal seat revenue. An EMSR algorithm appropriate for dependent demands is provided by Brumelle et al. (1990) where they study a simple two fare class model with dependent demand and focuses on upsells. They show that the state-dependent demands need to be incorporated to generate the correct solution. Fiig et al. (2005) analyze the downsell effect in single-leg RM and propose a novel fare adjustment approach known as Displacement Adjusted Virtual Nesting Marginal Revenue (DAVN-MR) to mitigate the revenue dilution. Gallego et al. (2009) develop new extensions of the traditional static EMSR-based approach for the single-leg RM with customer choice model. They also study the variations of the algorithm under both low-to-high and mixed arrival demand patterns. For a complete review on static EMSR-type heuristics for airline RM with dependent demand, we refer readers to Weatherford and Ratliff (2010).

Although DP generates optimal revenue performance for single-leg RM with customer choice behavior, it is computationally intensive, difficult to scale to large problems, and requires estimations of time dependent arrival rates. To overcome these shortcomings, the heuristics developed in this paper are designed to mitigate customer downsell risk and achieve good revenue performance with much less data and computational effort. Our heuristics are more robust than DP in a sense that they allow us to work with aggregate demands and forecast updates. Thus, these heuristics are promising candidates to be used as subroutines in a network optimization where processing speed are much more important.

3 PROBLEM FORMULATION

Similar to Talluri and van Ryzin (2004), we consider a multiple-fare-class, single-leg, revenue management problem where consumers make selections according to a choice model. The n fares are $p_1 > p_2 > \ldots > p_n$. We assume that customers arrive to the system at a rate of λ per unit time and that selling occurs over T periods $\{0, 1, \ldots, T-1\}$ with the plane departing at the beginning of period T. By selecting appropriate units of time, e.g., by measuring time in days or hours, we can make λ small while keeping λT constant. In particular, by selecting T such that $\lambda << 1$ we can think of λ as the probability that a single customer arrives in a period. Under this construction the probability of two or more arrivals over a period is negligible. While this construction is not strictly necessary, it is helpful to write a discrete time dynamic programming formulation of the problem (the alternative being a continuous time Hamilton Jacobi Bellman formulation).

3.1 Choice Model

We start with a brief description of a customer choice model. The complete set of products is denoted by $N = \{1, 2, ..., n\}$ with prices $p_1 > p_2 > \ldots > p_n$. For any subset $S \subset N$, let $\pi_j(S)$ be the probability that alternative $j \in S$ is selected where $\pi_o(S)$ denotes the probability that an outside alternative is selected.

As an example, consider the multinomial logit (MNL) choice model where the utility of product j is derived from the fare p_j and other attributes, such as schedule quality, summarized by a vector s_j. We assume homogenous customers choice behavior. Let β_s and β_p be the schedule and price sensitivities of fare j respectively, this results in utility $u_j = \beta_s s_j + \beta_p p_j$ and attractiveness $a_j = e^{\phi u_j}$ for product $j \in N$ where ϕ is a scale parameter. Let a_0 be the attractiveness of the no purchase alternative, then the probability of selecting $j \in S$ under the MNL is given by $\pi_j(S) = \frac{a_j}{a_0 + a(S)}$ with $\pi_0(S) = \frac{a_0}{a_0 + a(S)}$ where $a(S) = \sum_{j \in S} a_j$. This probability arises from a random utility model with $U_j = u_j + \varepsilon_j$ where the u_j is a deterministic representative component and the ε_j's are IID mean zero Gumbel random variables. The scale parameter ϕ is inversely proportional to the variance of the Gumbel distribution. Note that in this MNL model, the no-purchase alternative is intended to represent an aggregate attractiveness of all the competitor offerings in the market. This aggregate no-purchase alternative approach provides a simple means of representing competitive effects, and by increasing (or decreasing) the attractiveness of the no purchase alternative, the carrier's upsell rates can be decreased (or increased). This approach provides a practical means of including competitive effects with only minimal additional date collection effort.

3.2 Dynamic Programming Formulation

Let $V(t, x)$ be the optimal expected revenue over periods $\{t, t+1, \ldots, T\}$ starting with x units of inventory we can write the dynamic programming formulation as

$$V(t,x) = \max_{S \subseteq N} \sum_{j \in S} \lambda \pi_j(S)[p_j + V(t+1, x-1)] + (1 - \lambda + \lambda \pi_o(S))V(t+1, x), \tag{1}$$

where $\lambda\pi_j(S)$ is the probability that a customer arrives during the period and selects alternative $j \in S$ and $1-\lambda+\lambda\pi_o(S)$ is the probability of no purchase during a period. It is easy to see that

$$V(t,x) = V(t+1,x) + \max_{S\subseteq N} \lambda(S)[q(S) - \Delta V(t+1,x)]$$

where $\Delta V(t,x) = V(t,x) - V(t,x-1)$, $\lambda(S) = \lambda \sum_{j\in S} \pi_j(S)$, $r(S) = \sum_{j\in S} p_j \lambda \pi_j(S)$ is the revenue rate under set S and

$$q(S) = \frac{r(S)}{\lambda(S)} = \sum_{j\in S} p_j \frac{a_j}{a(S)}$$

is the average fare conditional on a sale under set S. The boundary conditions of the problem are $V(T,x) = V(t,0) = 0$.

Notice that formulation (1) allows the decision maker to open and close fares at any time. In particular, it implicitly relaxes the traditional assumption that low fare demands arrive before high fare demands. Alternatively, it is possible to formulate a more restricted (and more difficult to solve) dynamic program if the fares cannot be opened once they are closed. This is more difficult to solve because there is an option value that is exercised every time a fare is closed and this option value has to be implicitly computed in the calculations of the resulting dynamic program. In practice, it leads to smaller booking limits or equivalently larger protection levels. It is also possible to modify the formulation to put restrictions on the subsets of fares that are available at time t. This helps, for example, if some fares are not available after a certain period because of their time of purchase restrictions.

3.3 Efficient Fares

The optimization problem (1) is over all 2^n subsets of N. For the MNL, Talluri and van Ryzin (2004) have shown that only sets of the form $S_j = \{j,\ldots,1\}$ (This is called 'complete set' in Talluri and van Ryzin (2004).), $j = 1,\ldots,n$ can be part of the optimal solution. In other words, other subsets of N are inefficient in the sense that they generate a revenue rate $r(S)$ that can be replicated by a combination of other sets in such a way that the combination has a consumption rate lower than or equal to $\lambda(S)$. For simplicity we will write $r_j = r(S_j)$, $\lambda_j = \lambda(S_j)$ and $q_j = q(S_j)$, and for convenience we define $r_0 = \lambda_0 = 0$.

Even though the efficient set can only be of form $S_j = \{j,\ldots,1\}$ under MNL choice model, there is no guarantee that all these n sets $\{S_n,\ldots,S_1\}$ are efficient. It is possible that some of these sets are inefficient and thus will not be included in the optimal solution. It helps to identify these sets/fares and narrow the search region before solving the DP.

Let

$$u_j = \frac{r_j - r_{j-1}}{\lambda_j - \lambda_{j-1}}.$$

Gallego et al. (2009) have proved that that a sufficient condition for all sets $S_j, j = 1,\ldots,n$ to be efficient is that $u_1 > u_2 > \ldots > u_n > u_{n+1} \equiv 0$. If this condition does not hold, then there is at least one inefficient fare that can be removed without loss of optimality. We can continue removing fares until this condition is satisfied after relabeling the fares.

There are choice models for which the set of efficient fares is nested (i.e., the sets $\hat{S}_1,\ldots,\hat{S}_m$ are called nested if there is an increasing order of $\hat{S}_1 \subseteq \hat{S}_2 \subseteq \cdots \hat{S}_m$.), although not necessarily nested by fare. For example, Talluri and van Ryzin (2004) consider a formulation in which the efficient fares are $\{1\},\{1,3\}$ and $\{1,2,3\}$. The results of this paper are applicable to this case by reordering the labeling of the products. Note that performing these checks to eliminate inefficient fare classes is crucial for proper application of the algorithms presented in this document, and this step should not be bypassed.

3.4 Linear Programming Formulation and Dual Program

In this section, we study the deterministic linear programming formulation under the assumption that customers arrive at a constant rate λ and that a fixed fraction, namely $\pi_k(S_j)$ select choice $k \in S_j$. The study of this deterministic model will shed light on developing the heuristics in the next section.

Let t_j be the length of time we use set S_j, $j = 1, \ldots, n$. The corresponding linear program with decision variable t_j is given by

$$\begin{array}{lrcl} \max & \sum_{j=1}^{n} r_j t_j & & \\ \text{s. t.} & \sum_{j=1}^{n} \lambda_j t_j & \leq & C, \\ & \sum_{j=1}^{n} t_j & \leq & T, \\ & t_j & \geq & 0, \quad j = 1, \ldots, n \end{array}$$

Notice that this same linear program is valid regardless of the assumptions made in terms of the flexibility to open and close fares in the stochastic problems.

The dual program is

$$\begin{array}{lrcl} \min & Cu + Tv & & \\ \text{s. t.} & \lambda_j u + v & \geq & r_j \quad j = 1, \ldots, n, \\ & u, v & \geq & 0. \end{array}$$

So $v = \max_j (r_j - \lambda_j u)^+$ and the program reduces to

$$\begin{array}{ll} \min & Cu + T \max_j (r_j - \lambda_j u)^+, \\ \text{s. t.} & u \geq 0. \end{array} \tag{2}$$

Notice that $u = u_{n+1} = 0$ is optimal if $\lambda_n T \leq C$, resulting in objective function Tr_n. This corresponds to $t_n = T$ and $t_j = 0$, $j < n$. Otherwise, let $m = m(T,C) = \min\{i : \lambda_i T > C\}$. Then $u = u_m = \frac{r_m - r_{m-1}}{\lambda_m - \lambda_{m-1}}$ is optimal as the function (2) is decreasing up to u_m and is increasing thereafter. As capacity becomes larger, the index $m(T,C)$ goes up and u_m goes down, so the marginal value of capacity goes down. Likewise, as the time horizon becomes larger, the index $m(T,C)$ decreases and u_m increases, so v decreases.

The primal solution involves only sets S_{m-1} and S_m with $t_m = \frac{C - \lambda_{m-1} T}{\lambda_m - \lambda_{m-1}}$ and $t_{m-1} = T - t_m$. Thus $V_d(T,C) = Cu_m + T(r_m - \lambda_m u_m) = r_{m-1}(T - t_m) + r_m t_m$. Moreover, it can be shown that $V(t,x) \leq V_d(t,x)$ for all (t,x). It is well known that $V(T,C) \leq V_d(T,C)$ proved by Gallego and van Ryzin (1994), so the deterministic solution gives an upper bound on the expected revenue under any policy.

4 DYNAMIC PROTECTION LEVELS

We analyze two different dynamic models and provide heuristics for them respectively. The first heuristic, Central Ray Heuristic (CRH), is designed for the dynamic program formulation (1) where fares can open and close at any given time depending on the state of the system. The other heuristic, Unidirectional Closure Heuristic (UCH), assumes that the fares can not be reopened once they are closed. The corresponding dynamic program formulation is different from equation (1). Because of this restriction, the revenue generated by UCH is generally lower than that of CRH when fares can be opened and closed. The UCH applies when airlines can not freely reopen lower fares due to certain business constraints.

4.1 Central Ray Heuristic (CRH): Fares Can Be Opened and Closed

Suppose that at state (t,x) we receive a sales request under set S_k. If we accept the sales request and then follow an optimal policy the resulting expected revenue is given by

$$V_{t+1}(x) + \lambda_k [q_k - \Delta V_{t+1}(x)].$$

If instead we switch to set S_{k-1} and then follow an optimal policy, the revenue is given by

$$V_{t+1}(x) + \lambda_{k-1} [q_{k-1} - \Delta V_{t+1}(x)].$$

Taking the difference and manipulating the terms indicates that it is better to switch to set S_{k-1} at time t if and only if

$$\Delta V_{t+1}(x) = V_{t+1}(x) - V_{t+1}(x-1) > u_k.$$

It helps to visualize what happens in the deterministic problem before develop a heuristic. If the state (t,x) is such that $\lambda_{k-1}t < x < \lambda_k t$, then the solution that insists on using set k before set $k-1$ would sell at rate λ_k until the remaining inventory sits on the the ray $\{\lambda_{k-1}s, s \geq 0\}$ defined by λ_{k-1}. When we have the flexibility of opening and closing set S_k, we find the following alternative solution to the linear program useful. Indeed, the solution we have in mind is one that switches between efficient set S_{k-1} and S_k, using set S_k when the inventory is above the central ray $\{\tilde{\lambda}_{k-1}s : s \geq 0\}$ and uses set S_{k-1} when the inventory is below the ray, where $\tilde{\lambda}_{k-1} = (\lambda_{k-1} + \lambda_k)/2$.

Once this ray is reached for the first time, the inventory can be drawn down to zero along this ray by continuously switching between sets $k-1$ and k, spending half of the time on each. The modified heuristic that we propose will use the ray $\tilde{\lambda}_{k-1}$ and the corresponding average fare $\tilde{q}_{k-1} = (r_{k-1} + r_k)/(\lambda_{k-1} + \lambda_k)$ along this ray instead of λ_{k-1} and q_{k-1} to approximate the marginal value of capacity. This leads to following approximation of $\Delta V_{t+1}(x)$.

$$\Delta V_t(x) \simeq \tilde{q}_{k-1}\mathbb{P}(\tilde{D}_{k-1}(t) \geq x),$$

resulting in protection level

$$\tilde{y}_{k-1}(t) = \max\left\{y \in \mathcal{N} : \mathbb{P}(\tilde{D}_{k-1}(t) \geq y) > \frac{u_k}{\tilde{q}_{k-1}}\right\}$$

where $\tilde{D}_{k-1}(t)$ is defined as the remaining demand distribution and is Poisson with parameter $\tilde{\lambda}_{k-1}(T-t)$.

Proposition 1 For any fixed k, protection level $\tilde{y}_k(t)$ is decreasing in t.

Proof. The result is straight-forward since $\tilde{D}_{k-1}(t)$ is Poisson with parameter $\tilde{\lambda}_{k-1}(T-t)$, so $\mathbb{P}(\tilde{D}_{k-1}(t) \geq y)$ is decreasing in t. Thus, $\tilde{y}_k(t)$ is decreasing in t. □

Proposition 2 If either $\frac{u_k}{\tilde{q}_{k-1}} \geq \frac{u_{k+1}}{\tilde{q}_k}$ or $\tilde{q}_{k-1}\mathbb{P}(\tilde{D}_{k-1}(t) \geq y) \leq \tilde{q}_k\mathbb{P}(\tilde{D}_k(t) \geq y)$, then $\tilde{y}_k(t) \geq \tilde{y}_{k-1}(t)$.

Proof. Let's look at the first sufficient condition. Note that $\tilde{\lambda}_{k-1} \leq \tilde{\lambda}_k$, we have

$$\mathbb{P}(\tilde{D}_{k-1}(t) \geq y) \leq \mathbb{P}(\tilde{D}_k(t) \geq y). \tag{3}$$

For any fixed t, let

$$\tilde{y}^a(t) = \max\left\{y \in \mathcal{N} : \mathbb{P}(\tilde{D}_{k-1}(t) \geq y) > \frac{u_{k+1}}{\tilde{q}_k}\right\}.$$

If $\frac{u_k}{\tilde{q}_{k-1}} \geq \frac{u_{k+1}}{\tilde{q}_k}$ is true, then $\tilde{y}_{k-1}(t) \leq \tilde{y}^a(t)$ since $\mathbb{P}(\tilde{D}_{k-1}(t) \geq y)$ is decreasing in y. From equation (3), we also have $\tilde{y}_k(t) \geq \tilde{y}^a(t)$. Thus, $\tilde{y}_k(t) \geq \tilde{y}_{k-1}(t)$.

The proof for the second sufficient condition is similar to the first one. Let

$$\tilde{y}^b(t) = \max\left\{y \in \mathcal{N} : \tilde{q}_{k-1}\mathbb{P}(\tilde{D}_{k-1}(t) \geq y) > u_{k+1}\right\},$$

then $\tilde{y}_{k-1}(t) \leq \tilde{y}^b(t)$ since $u_k \geq u_{k+1}$ and $\tilde{q}_{k-1}\mathbb{P}(\tilde{D}_{k-1}(t) \geq y)$ is decreasing in y. If for any fixed time t, we have $\tilde{q}_{k-1}\mathbb{P}(\tilde{D}_{k-1}(t) \geq y) \leq \tilde{q}_k\mathbb{P}(\tilde{D}_k(t) \geq y)$ for all y, then $\tilde{y}_k(t) \geq \tilde{y}^b(t)$. Thus, $\tilde{y}_k(t) \geq \tilde{y}_{k-1}(t)$. □

Using this heuristic, we can compute a table of protection levels $\tilde{y}_k(t)$, $k = 1, \ldots, n-1$ and $t = 0, 1, \ldots, T-1$. At state (t,x), we will offer set $S_{k(t,x)}$ where the index $k(t,x)$ is set to n if $\tilde{y}_n(t) < x$ and otherwise it is set to the smallest k such that $\tilde{y}_k(t) \geq x$.

Up to now we have not used the specific form of the choice model except for the fact that we have assumed the efficient sets are nested. Thus, the above applies to any situation where the efficient sets are nested even if they are not nested by fare values (but this may require relabeling the fares). As we shall see in the numerical examples in section (5), this heuristic performs extremely well with performance that is almost as good as solving the corresponding dynamic program.

4.2 Unidirectional Closure Heuristic(UCH): Fares Cannot Be Reopened Once They Are Closed

The heuristic approximates the marginal value of capacity when only set S_{k-1} is used over periods $\{t+1,\dots,T-1\}$. Let $D_k(t)$ be Poisson with parameter $\lambda_k(T-t)$. Then $V_t(x) \simeq q_{k-1} E\min(D_{k-1}(t),x)$ resulting in $\Delta V_t(x) \simeq q_{k-1}\mathbb{P}(D_{k-1}(t) \geq x)$. This leads a protection of

$$y_{k-1}(t) = \max\left\{ y \in \mathscr{N} : \mathbb{P}(D_{k-1}(t) \geq y) > \frac{u_k}{q_{k-1}} \right\} \tag{4}$$

units of resource at the beginning of period $t \in \{0,\dots,T-1\}$ for set S_{k-1}.

Similar to CRH heuristic, we can compute a table of protection levels $y_k(t)$, $k = 1,\dots,n-1$ and $t = 0,1,\dots,T-1$. At state (t,x), we will offer set $S_{k(t,x)}$ where the index $k(t,x)$ is set to n if $y_n(t) < x$ and otherwise it is set to the smallest k such that $y_k(t) \geq x$. Under the multinomial logit model, the above heuristic has the flavor of the EMSR-b with buy up developed by Belobaba and Weatherford (1996).

The protection levels $y_k(t)$ resulted from the EMSR-b type heuristic, see equation (4), tend to be more conservative compared to the optimal protection levels computed by DP. This is because the EMSR-b heuristic and optimal solutions for the case where fares cannot be reopened take into account the value of exercising the option to switch out of a fare.

5 SIMULATION AND NUMERICAL RESULTS

Monte Carlo Simulation is implemented to evaluate the performance of our proposed heuristics against the existing EMSRb heuristic that are adjusted to customer choice models Talluri and van Ryzin (2004). While EMSRb heuristic is developed under static-model assumptions where demand for each fare classes arrives in a low-to-high revenue order, it is frequently used as a heuristic in the dynamic settings by aggregating the total demand-to-come for each fare classes. Also, this heuristic can be adjusted to incorporate customer choice behavior, so it serves as a useful benchmark for comparison. Customer arrivals are simulated over a booking horizon according to a Poisson process. Customers upsell and downsell behaviors are captured in our simulation with the specified probabilities given by the MNL model. We compare the total revenue generated using availability policies specified by our proposed dynamic heuristics against the ones found by optimal DP as well as the existing EMSRb heuristic. Extensive study of a two-fare class, a three-fare-class and a ten-fare-class simulations over a wide range of parameters is conducted. We tested over 300 scenarios with different parameters varying across a wide range of market conditions. For example, the demand load factor (defined as aggregate customer demand divided by flight capacity) varies from 50 to 200 per cent, and price elasticity changes from -1.0 (low) to -4 (high). In all examples, the number of simulation runs is fixed at $40,000$. With this sample size, the revenue estimates have relative errors in the range of 0.05% to 0.13% with 99% confidence. In each case, customers choice behavior follows a MNL choice model.

5.1 Two-fare-class Simulations

The results here are based on a single flight with two fare classes. The high fare $p_1 = \$1,000$ has schedule quality $s_1 = \$1,200$, while the low fare $p_2 = \$600$ has schedule quality s_2. Price and schedule sensitivities are $\beta_p = -.005, \beta_s = .005$ respectively. The attractiveness (exponential utility) of these two products are $a_1 = a_2 = e^1$, and the outside alternative has attraction $a_o = e^1$. The booking time horizon is $T = 100$ and per period demand rate is $\lambda = .4$. We compare the revenue performance of CRH and UCH against EMSRb in different scenarios by varying schedule quality s_2 from \$650 to $\$1,200$ and flight capacity C from 10 to

40. Based on the simulation results, the CRH consistently outperforms ESMRb with an average revenue improvement of 1.36%, while UCH performs similarly to EMSRb. The average revenue gain of CRH and UCH over EMSRb at different load factors is shown in Table 1 below. Figure 2 presents the average suboptimality gaps of these heuristics over optimal DP.

Table 1: Average revenue gain of CRH and UCH over EMSRb in two-fare simulation.

Load Factor Range	CRH over EMSRb	UCH over EMSRb
$[0, 1.0]$	-0.01%	-0.02%
$(1.0, 1.5]$	0.44%	0.01%
$(1.5, 2.0]$	2.12%	0.00%
$(2.0, 2.5]$	3.42%	-0.01%
$(2.5, 3.0]$	3.77%	0.01%
$(3.0, 3.5]$	3.36%	0.02%
$(3.5, 4.0]$	2.59%	-0.02%
$(4.0, 5.0]$	1.59%	0.02%

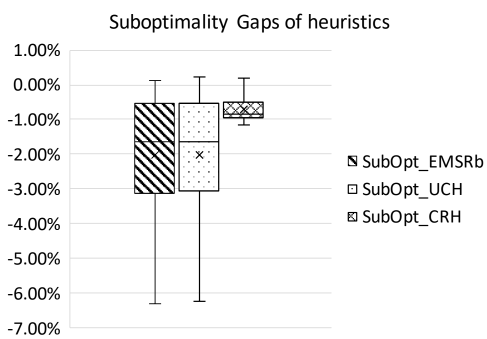

Figure 2: Suboptimality gap of CRH, UCH and EMSRb for two-fare simulations.

5.2 Three-fare-class Simulations

This example is taken fromTalluri and van Ryzin (2004) and involves three products Y, K and M with fares \$800, \$450 and \$500. In this example, only the sets $\{Y\}, \{Y,K\}$ and $\{Y,K,M\}$ are efficient. We relabel the fares as 1, 2 and 3 with 1 corresponding to fare Y, 2 corresponding to fare K and 3 corresponding to fare M. Then the three efficient sets are $S_k = \{1,\dots,k\}$ for $k = 1,2,3$. Our simulation uses per period demand rate $\lambda \in \{.15, .20, .25\}$. For a capacity level $C = 20$ and $\lambda = .25$, table 2 illustrates the optimal dynamic protection levels together with protection levels obtained by heuristics respectively at the beginning of period $t \in \{0,10,20,40,60,80\}$. For example, at $t = 0$, we see that $y_1 < C \le y_2$ so we start by offering set S_2.

Table 3 compares the results of heuristics at different capacity levels at when $\lambda T = 20$. We find that the performance of UCH is similar to EMSRb adjusted to choice models. Both of these perform well when the capacity becomes large. When capacity is tight, the largest revenue gap of UCH compared to optimal DP solution is around -2.82% when $\lambda T = 20$ and capacity level of 10. The undesirable performance of both UCH and EMSRb is probably due to the fact that it cannot reopen fares once they are closed, and this limitation precludes UCH and EMSRb from capturing more high fare demands, thus it leads to lower total revenue. The optimal policy, on the other hand, is more aggressive in protecting high fare demand, because it has the option of reopening low fares when the actual demand turns out to be low. The central ray heuristic(CRH), on the other hand, is more robust. The largest suboptimality gap of -1.04% occurs

Table 2: Dynamic protection levels when $\lambda T = 25$.

t	y_1^*	y_2^*	y_1^{EMSRb}	y_2^{EMSRb}	y_1^{UCH}	y_2^{UCH}	y_1^{CRH}	y_2^{CRH}
0	12	20	7	21	7	21	12	20
10	11	20	6	19	6	19	10	20
20	10	18	5	17	5	17	9	19
40	7	14	4	13	4	13	6	14
60	5	9	3	9	3	9	4	10
80	2	5	1	4	1	4	2	5

at capacity $C = 18$. The average revenue gap of CRH is around -0.6%. The revenue improvement of CRH over existing EMSRb is shown in the last column of table 3. The largest improvement occurs when the capacity is tight when the load factor is high. The similar conclusion can be drawn by examining the suboptimality gaps of all heuristics compared to the optimal value at different capacity level and arrival rates. The results are presented in Figure 3. In summary, CRH performs much better than UCH and ESMRb when the capacity is limited, and when capacity is abundance and load factor is low, all heuristics perform equally well.

Table 3: Suboptimality gaps of CRH, UCH, EMSRb at different capacity when $\lambda T = 20$.

Capacity	Load Factor	SubOpt EMSRb	SubOpt UCH	SubOpt CRH	$\frac{UCH-EMSRb}{EMSRb}$	$\frac{CRH-EMSRb}{EMSRb}$
10	2.0	-3.26%	-3.22%	-0.94%	0.04%	2.39%
12	1.67	-2.19%	-2.23%	-0.99%	-0.04%	1.23%
14	1.43	-1.57%	-1.51%	-0.85%	0.05%	0.72%
16	1.25	-1.47%	-1.41%	-1.04%	0.06%	0.43%
18	1.11	-1.29%	-1.36%	-0.98%	-0.07%	0.31%
20	1.00	-1.17%	-1.16%	-0.87%	0.00%	0.30%
22	0.91	-0.96%	-0.96%	-0.58%	0.00%	0.39%
24	0.83	-0.69%	-0.60%	-0.45%	0.10%	0.25%
26	0.77	-0.34%	-0.24%	-0.40%	0.09%	-0.06%

5.3 Ten-fare-class Simulation

A more complicated ten-fare example is also extracted from Talluri and van Ryzin (2004), but is modified to allow the no-purchase alternative in customer choice model. Under the assumption that customers behave according to a MNL model, the only efficient sets are nested by fare order, $S_k = 1,\ldots,k$ where $1 \leq k \leq 10$. The arrivals follow a homogeneous Poisson process with mean 205 and total booking periods are $T = 410$, so the arrival rate is $\lambda = 0.5$ per period. We again test both heuristics under multiple scenarios by varying both capacity levels and choice model parameters, and compare them against existing EMSRb heuristic. These simulation results again suggest that while UCH performs similar to EMSRb (with an average revenue gaps less than 0.1%), CRH performs much better than EMSRb and is fairly robust (all revenue performances are within 0.5% of the optimal solution) across all scenarios that we tested . The revenue gains of CRH over EMSRb range from 0.1% to 2.3% with the biggest improvement occurring when the capacity is tight which confirms with our previous simulation findings.

6 CONCLUSION

We develop two computationally efficient heuristics for dynamic single-leg revenue management problems with dependent demand. The product availability policies generated by these heuristics can help mitigate

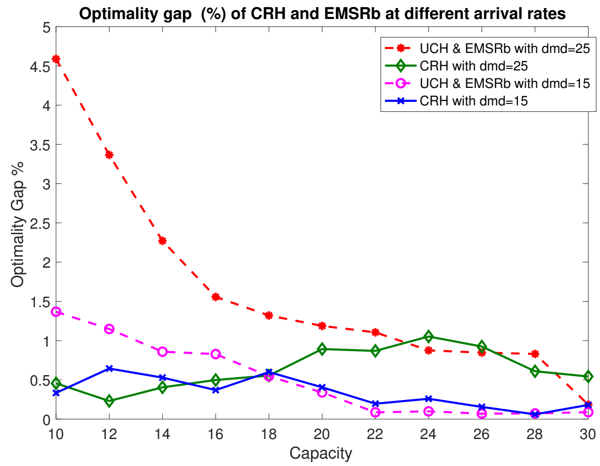

Figure 3: Suboptimality gap of CRH and EMSRb at $\lambda T = 15$ and $\lambda T = 25$.

revenue dilution caused by customer downsells. The key idea underlying the heuristics is to find good approximations of marginal revenue of capacity. Central Ray Heuristic (CRH) is developed for the cases where fares can be opened and closed freely during the booking horizon. For other cases where fares cannot be reopened once they are closed, we propose a different heuristic, namely Unidirectional Closure Heuristic (UCH), to compute the protection levels. Computational experiments indicate that while UCH performs similar to EMSRb (adjusted to customer choice models), CRH outperforms existing EMSRb heuristic in almost all market conditions. When capacity is tight (i.e., load factor is high), significant revenue improvement (1.5%-4.5%) can be gained using CRH over EMSRb. When capacity is abundant (i.e., load factor is low), both heuristics generate similar protection levels so their performance are comparable. Our simulation study also indicates that CRH performs close to the optimal solution generated by dynamic programming, which makes it a promising candidate that can be used as subroutines in large size network revenue management problems to obtain good solutions within a reasonable amount of time.

The current study has two limitations. First, each choice model will only be valid for a certain amount of time and the corresponding efficient sets will be different from the other choice models. Second, the capacity allocation is more complicated. Not only do managers need to decide when to switch between different efficient sets, they also have to decide how much capacity to allocated to each choice model. Future work may include 1) validating our proposed CRH heuristic in a more complicated network setting, 2) further extending our models to consider time-varying arrival rate λ_t, and incorporating time heterogeneous choice models into this framework and develop efficient heuristics for the corresponding problems.

REFERENCES

Algers, S., and M. Besser. 2001. "Modelling Choice of Flight and Booking Class: A Study Using Stated Preference and Revealed Preference Date". *International Journal of Services Technology and Management* 2(1/2):28–45.

Andersson, S. E. 1998. "Passenger Choice Analysis for Seat Capacity Control: A Pilot Project in Scandinavian Airlines". *International Transactions in Operational Research* 5:471–486.

Belobaba, P. P., and L. R. Weatherford. 1996. "Comparing Decision Rules that Incorporate Customer Diversion in Perishable Asset Revenue Management Situations". *Decision Sciences* 27:343–363.

Ben-Akiva, M., and S. Lerman. 1994. *Discrete Choice Analysis: Theory and Applications to Travel Demand*. Sixth ed. Cambridge, MA: The MIT Press.

Brumelle, S. L., J. I. McGill, T. H. Oum, K. Sawaki, and M. W. Tretheway. 1990, August. "Allocation of Airline Seats between Stochastically Dependent Demands". *Transportation Science* 24(3):183–192.

Cooper, W. L. 2002. "Asymptotic Behavior of an Allocation Policy for Revenue Management". *Operations Research* 50:720–727.

Dong, L., P. Kouvelis, and Z. Tian. 2009. "Dynamic Pricing and Inventory Control of Substitute Products". *Manufacturing and Service Operations Management* 11(2):317–339.

Feldman, J., and H. Topaloglu. 2017. "Revenue Management under the Markov Chain Choice Model". *Operations Research* 65(5):1322–1342.

Fiig, T., K. Isler, C. Hopperstad, and R. Cleaz-Savoyen. 2005. "DAVN-MR: A Unified Theory of O&D Optimization in a Mixed Network with Restricted and Unrestricted Fare Products". In *AGIFORS Revenue Management and Distribution Study Group Meeting*. May 15th-19th, Cape Town, South Africa.

Gallego, G., L. Li, and R. Ratliff. 2009. "Choice-based EMSR Methods for Single-leg Revenue Management with Demand Dependencies". *Journal of Revenue and Pricing Management* 8(1):207–240.

Gallego, G., R. Ratliff, and S. Shebalov. 2015. "A General Attraction Model and Sales-based Linear Program for Network Revenue Management under Customer Choice". *Operations Research* 63:212–232.

Gallego, G., and G. J. van Ryzin. 1994. "Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons". *Management Science* 40:999–1020.

Hartmans, G. 2006. "Sell-up Simulation". In *AGIFORS Revenue Management Study Group Meeting*. May 19th-22nd, Cancun, Mexico.

Joyner, H. 2007. "Airline Industry Outlook". In *32nd FAA Aviation Forecast Conference*. March 16th Washington DC, USA.

Lautenbacher, C. J., and S. J. Stidham. 1999. "The Underlying Markov Decision Process in the Single-leg Airline Yield Management Problem". *Transportation Science* 34:136–146.

Lee, T. G., and M. Hersh. 1993. "A Model for Dynamic Airline Seat Inventory Control with Multiple Seat Bookings". *Transportation Science* 27:252–265.

Liang, Y. 1999. "Solution to the Continuous Time Dynamic Yield Management Model". *Transportation Science* 33:117–123.

Mishra, S., R. Ratliff, and B. Vinod. 2005. "New Generation in Demand Forecasting". In *AGIFORS 45th Annual Symposium*. Oct 1st-4th, Sao Paulo, Brazil.

Strauss, A. K., R. Klein, and C. Steinhardt. 2018. "A Review of Choice-based Revenue Management: Theory and Methods". *European Journal of Operational Research* 271(2):375–387.

Talluri, K., and G. van Ryzin. 2004. "Revenue Management under a General Discrete Choice Model of Consumer Behavior". *Management Science* 50:15–33.

van Ryzin, G., and G. Vulcano. 2008. "Computing Virtual Nesting Controls for Network Revenue Management under Customer Choice Behavior". *Manufacturing and Service Operations Management* 10(3):448–467.

Wang, R., and Z. Wang. 2017. "Consumer Choice Models with Endogenous Network Effects". *Management Science* 11(2):3944–3960.

Weatherford, L. R., and R. Ratliff. 2010. "Review of Revenue Management Methods with Dependent Demands". *Journal of Revenue and Pricing Management* 9(4):326–340.

AUTHOR BIOGRAPHIES

Lin Li is Assistant Professor of Systems and Industrial Engineering at Kennesaw State University. She received her Ph.D. in Operations Research from Columbia University and later worked at IBM research center in New York as well as Sabre Inc. in Dallas. Her research primarily focuses on modelling, optimization, and simulation in areas such as supply chain management and revenue management. Her professional work experience is closely related to applying statistical methods and simulations to study demand modelling in travel industry. Dr. Li is a member of Institute of Industrial and Systems Engineer (IISE) and Institute for Operations Research and the Management Sciences (INFORMS). Her email address is Lin.Li@kennesaw.edu.

Yuan Zhou is an Assistant Professor of the Department of Industrial, Manufacturing, and Systems Engineering at the University of Texas at Arlington. Her research areas focus on complex adaptive system modeling, discrete-event and agent-based simulation, operations and process improvement, performance measurement, and data analytics. Her previous work includes agent-based modeling and simulation for infectious disease transmission and management, statistical modeling for nursing workload assessment, discrete-event simulation for evaluating electronic health record interoperability, and data analytics for home care service improvement. Also, Dr. Zhou is actively working with service organizations, where she applies methods and tools of systems engineering, operations management, and predictive analysis to help organizations identify operational bottlenecks and improve processes efficiency and quality. Dr. Zhou is a member of Institute of Industrial and Systems Engineer (IISE) and Institute for Operations Research and the Management Sciences (INFORMS). Her email address is yuan.zhou@uta.edu.