

Choice Based Revenue Management for Parallel Flights

Jim Dai

Georgia Institute of Technology and Cornell University, dai@gatech.edu,

Weijun Ding

Georgia Institute of Technology, wding34@gatech.edu,

Anton J. Kleywegt

Georgia Institute of Technology, anton@isye.gatech.edu,

Xinchang Wang

Georgia Institute of Technology, xwang336@gatech.edu,

Yi Zhang

Georgia Institute of Technology, yzhang33@gatech.edu,

This paper describes a revenue management project with a major airline that operates in a fiercely competitive market involving two major hubs and having more than 30 parallel daily flights. The market has a number of unusual characteristics including (1) almost half of customers choose not to purchase the tickets after booking; (2) about half of customers purchase their tickets within 3 days of departure; and (3) a significant number of customers no-show or go-show.

We formulate choice based stochastic optimization problems to maximize the expected revenue for the airline. The inputs of the stochastic models include booking arrival rates, competitor assortment selection, booking choice probabilities for the airline's own flights as well as competitors' flights, booking-to-ticketing conversion probabilities, and go-show and no-show probabilities. We build a number of booking choice models, including multinomial logit models, nested logit models, and mixed logit models. The latter two types of models are aimed at incorporating unobserved heterogeneous customer preferences for different departure times. We formulate corresponding deterministic (fluid) optimization problems under each of the three booking choice models. We designed a column generation algorithm to compute optimal or near-optimal solutions for the deterministic problems, and the solutions are used to make assortment selections for the stochastic problem. The models used as input for the optimization problems are calibrated using 2011 data or 2012 data. Simulation studies using 2012 data show that the fluid based booking policies generate significantly more revenue than the airline's existing policy, and that policies based on the simpler multinomial logit models perform better than policies based on nested logit and mixed logit models, even when the simulation is based on the latter models, and even though the latter models seem to be more realistic.

Key words: Pricing and Revenue Management, Transportation, Consumer Behavior

1. Introduction

In this paper we describe an airline revenue management project. With “revenue management” we mean the modeling of a seller’s decisions and the seller’s resulting sales and revenue and/or profit, the calibration of the models with available data, the optimization of the seller’s objective, and the testing of the decisions obtained from the optimization, possibly allowing for influences not captured in the model. Although revenue management is aimed at optimizing the objective of a single seller, the influence of other sellers and the resulting competition is taken into account in our work.

Revenue management as described above started in the airline industry; many authors refer to Littlewood (1972) as the origin of airline revenue management. Littlewood (1972) has indeed been very influential in revenue management — the intuitive Littlewood rule to allow bookings in a low fare class if the low fare is greater than the product of the high fare and the probability that the remaining seats can be filled charging the high fare, has formed the core of the EMSRb method and airline revenue management software for many years. Littlewood (1972) also pointed out that it had become “increasingly important for the airline objective to be to maximise revenue instead” of traditional performance measures such as revenue passenger miles per seat mile or load factor (revenue passengers per seat). This may seem obvious to us today, but at the time the use of traditional performance measures was so deeply embedded that Littlewood’s suggestion was truly revolutionary. Thus, the origin of airline revenue management can be traced back more than 40 years ago. It may be surprising to an outsider that there remains anything of significance to be done in revenue management, and especially airline revenue management, but we think that researchers have barely begun to address many of the most important aspects of revenue management, such as modeling consumer behavior and competitor behavior. In this paper we address some important issues, some of which we only became aware of after analyzing the data. We give an overview of all major parts of the project, with emphasis on aspects that are novel.

Project overview. This project considers one of the busiest origin-destination markets in the world, involving about 30 flights per day by three major airlines that we call XX, YY, and ZZ. The airline that we optimize for is XX, with YY and ZZ as its competitors. Airline XX has implemented versions of our models and booking policies in its booking control system.

The project consisted of the following major activities:

1. We acquired and analyzed data of the origin-destination market, including booking data and fare class availability data of all three competitors, as well as data on ticketing, cancellations, no-shows, and go-shows of airline XX. Section 4 provides more detail about the data.

2. We developed models of customer arrivals, competitor assortment selection, customer booking choice behavior, customer purchasing behavior, and cancellation, no-show, and go-show behavior. Separate models were calibrated with 2011 data and 2012 data. The models and estimation results are presented in Section 5.

3. We formulated a customer choice based stochastic optimization problem to optimize the expected revenue for airline XX. The stochastic optimization problem is given in Section 3.1. We consider three versions of the problem based on three classes of booking choice models.

4. Corresponding to the stochastic optimization problem, we formulated a deterministic (or fluid) optimization problem to optimize the expected revenue for airline XX. The deterministic optimization problem is given in Section 3.2. We also consider three versions of the deterministic problem based on three classes of booking choice models.

5. The deterministic optimization problem is challenging because of the large number of possible assortments. For the deterministic problem under a multinomial logit choice model, we present an efficient column generation algorithm which converges to an optimal solution. For the deterministic problem under nested logit and mixed logit choice models, we use heuristics to generate promising columns. The algorithms are discussed in more detail in the online companion to this paper. The resulting solutions are called *fluid booking policies*. We solved each of the three versions of the deterministic optimization problem for the models calibrated with 2011 data and for the models calibrated with 2012 data, obtaining six fluid booking policies.

6. We built a computer simulation for each version of the stochastic optimization problem. The simulation uses models calibrated with 2012 data. We evaluated each of the six fluid booking policies as well as the actual 2012 booking policy of airline XX using the simulation for each of the three versions of the stochastic optimization problem. Thereby we obtained numerical results regarding the performance of the fluid booking policies when the booking choice model used to obtain a fluid booking policy is the same as the booking choice model used in the simulation, as well as when the fluid booking policy and the simulation are based on different booking choice models, and when the models used to obtain a fluid booking policy are calibrated with older data than the simulation, as well as when the models used to obtain a fluid booking policy are calibrated with the same data as the simulation. This gives some evidence of the robustness of the performance results with respect to modeling errors as well as with respect to errors in parameter estimates.

Modeling innovations. Our models include the following features that we think are innovative: (1) competition, (2) booking but not ticketing, (3) go-shows, and (4) modeling of unobservable customer choice sets. Relatively few papers in the revenue management literature have considered

competition, and as far as we know, the other three features are new in the revenue management literature. Next we discuss these features in a little more detail.

1. There is fierce competition in the market between the three airlines. Therefore, the customer choice sets in the booking choice models include alternatives offered by airline XX, as well as YY and ZZ. Also, we forecast the alternatives that airlines YY and ZZ will make available during the remainder of the booking horizon.

2. A large number of potential passengers make bookings but do not buy the corresponding ticket. As far as we know, this issue has not been addressed in the revenue management literature. As a result there is a need to make a distinction between booking and purchasing (also called “ticketing”) in the models, including a model of the conversion of booking to purchasing.

3. In this market there are a significant number of no-shows and go-shows. “Go-shows” are customers who show up at the airport and then request a change of flight. As far as we know go-shows have not been addressed in the revenue management literature. Another issue that we addressed that is often neglected in the revenue management literature, is no-show behavior. In this case, four types of no-show customer are distinguished: (1) no-shows who go-show, (2) no-shows who change to another flight and pay the required fee, (3) no-shows who cancel after the flight and request a (partial) refund, and (4) no-shows who never request a refund or ticket change. Therefore we developed no-show and go-show models, and incorporated the impact of no-show and go-show behavior on capacity utilization into the optimization model.

4. Another topic that is of general importance in airline revenue management, but that has not received attention in the literature, is the following: Typically, individual transaction data are used to calibrate discrete choice models. Such data identify, for each customer, the alternative chosen by the customer, as well as various attribute values of the customer and the chosen alternative. However, the traditional theory of discrete choice model estimation assumes that for each customer, the set of alternatives (called the choice set) among which the customer chooses is observed. This assumption hardly ever holds in practice, but the violation of this assumption is ignored in most discrete choice work. In the airline revenue management setting described, different customers are interested in flying during different time intervals, but the time interval of interest to a customer is usually not observed. This issue is especially important in a market such as the one described, where there are many flights per day, and most passengers travel for business purposes and consider fairly narrow time intervals for travel. If the discrete choice model incorporates parameters for time interval preference, without explicitly taking into account that the time interval preference is different for different customers, it could cause severe bias in the estimates of the other parameters.

To illustrate the intuition, consider a simple setting with two time intervals, in which half the customers choose the first time interval and half the customers choose the second time interval, and then each customer chooses the cheapest flight in the customer's chosen time interval, with the prices in one time interval being less than the prices in the other time interval. If one would estimate a discrete choice model with a price parameter (as well as parameters for time interval preference), with the set of all flights as the choice set for each customer, then the price parameter would indicate that the customers are not very price sensitive, because half the customers choose a relatively expensive flight. On the other hand, if one would estimate a model with the correct (but typically unobserved) choice set for each customer, the resulting model would indicate that customers are much more price sensitive than the previous model. Since the true choice sets are not observed, multinomial logit models are of the first (biased) type. We develop models to capture unobserved variation in travel time preferences among the population of customers. As the intuition suggests, we found that the price parameters are significantly more negative in the models that incorporate preference variation than in the multinomial logit models.

Important phenomena. Next we discuss a number of phenomena in the considered market that play an important role in the models.

1. Both origin-destination airports involved are hubs, and therefore the itineraries of almost all passengers on the flights between the two hubs are between these two airports only. (If one can travel from an origin O via hub A to hub B , then one can also travel directly from origin O to hub B , and in general the latter itinerary is preferred over the former. Thus, almost all the passengers on flights from A to B have origin A and destination B , that is, they are not connecting passengers.) In the literature this is called the setting with parallel flights and it is the setting considered in this paper.

2. A large fraction of bookings (and purchasing) takes place close to departure time. (This leaves little time for recourse actions, and makes improvements in revenue management more challenging.) For example, in 2011, 69% of bookings took place during the last week before departure, 49% of bookings took place during the last 3 days before departure, and 23% of bookings took place on the day of departure. That makes good forecasting especially important, because as mentioned there is not much time for recourse.

3. Ticket prices tend to be cheaper earlier in the booking horizon and more expensive closer to departure time. For example, Figure 1 shows that 20 days and 14 days before departure, more than 50% of the time the cheapest fare class (class 13) is the cheapest available fare class on a flight, but 1 day before departure, more than 70% of the time the most expensive fare class (class 1) is the cheapest available fare class on a flight.

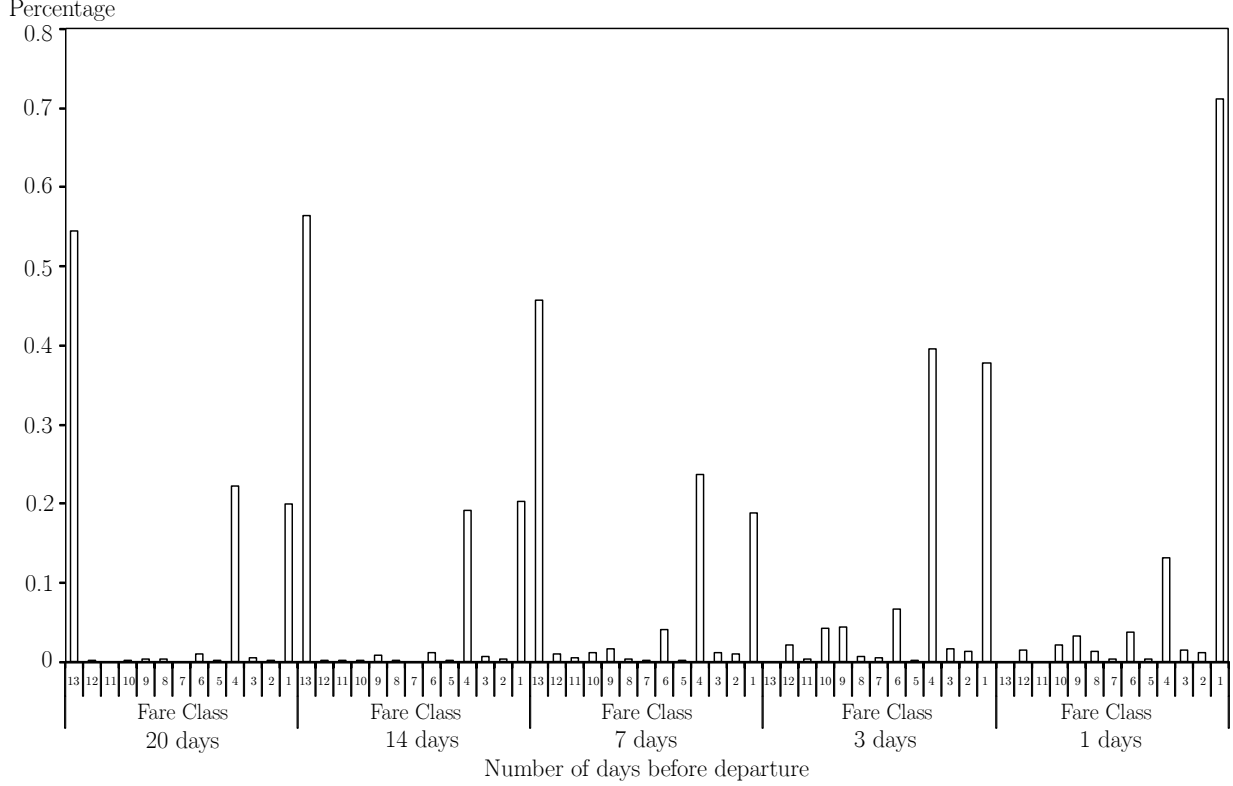


Figure 1 Fraction of days that each fare class is the cheapest available fare class for different numbers of days before departure. Class 1 is the most expensive, and class 13 is the cheapest.

4. The previous two observations together mean that a large fraction of bookings occur close to departure time, in spite of the steep discounts that are often offered for early bookings. These two observations together suggest that many passengers travel for business reasons and are not very price sensitive. The fare class selection behavior discussed in Section 5.2.6 provides further evidence of a quite clear distinction between price sensitive and price insensitive customers, and that the price insensitive customers constitute a significant fraction of customers.

5. The first two observations together also suggest that the booking arrival rate does not change much if prices change. Thus, it seems reasonable to assume that the arrival rate is exogenous, independent of the booking control policy. This assumption is standard in the literature and is also assumed in this paper.

6. In this market, the price of a round-trip ticket is the sum of the prices of the two one-way tickets. Data show that the majority of customers buy one one-way ticket at a time (the ratio of one-way tickets sold to round-trip tickets sold is 13:1). Therefore, the products we consider are tickets for a single flight from a specific origin to a specific destination.

7. Most flights are not full at departure time.

8. The first two observations together with the previous observation indicate that if a customer wants to travel, the customer will travel. Also, note that customers may book and then not ticket. This behavior is addressed by our booking-to-ticketing conversion model. Thus, most customers' decisions to book or not to book are not affected by the offered assortments. Therefore, it seems reasonable for our booking choice models to ignore the no-booking option.

9. Very few ticketed customers cancel or change their tickets before departure time, except for customers who go-show before departure time. For example, the cancellation rate in 2011 was only 2.1% for airline XX. This behavior is possibly related to the fact that most customers book close to departure time, and the availability of the option to book but not ticket. Therefore, in this paper, we include only go-shows before departure time, and ignore ticket cancellations or changes before the departure day.

2. Literature Review

There is a large literature on revenue management with airline applications. Here we give a brief overview of some of this literature that intersects with our paper in various ways.

A classic paper on the implementation of a revenue management system at an airline is Smith et al. (1992), based on work that won the Franz Edelman prize in 1991. A few other papers that report revenue management research implemented at airlines are Alstrup et al. (1989), Smith et al. (2001), Slager and Kapteijns (2004), and Vulcano et al. (2010). Our paper is similar to that of Vulcano et al. (2010) in the sense that we also use discrete choice models for parallel flights calibrated with airline data and we also use simulation to evaluate the performance of our policies. However, we model some aspects not discussed in Vulcano et al. (2010): we explicitly model competitor assortments, customers who book but do not purchase, no-shows, and go-shows. In addition, our choice models address some issues not addressed in Vulcano et al. (2010): we develop models that incorporate the idea that different customers have different preferences (taste heterogeneity) for different departure times, we allow differences in price sensitivity depending on when the customer books and what channel the customer uses, and we identified and modeled discontinuous demand responses (spikes) for the cheapest available fare classes as well as for fully refundable fare classes. Zhang and Cooper (2005) also considered discrete choice models for parallel flights. They studied structural properties of a Markov decision process formulation, and they compare a number of heuristics for their model.

We use discrete choice models to model booking choice, booking-to-ticketing conversion, and no-show and go-show behavior. Discrete choice models have been used to predict travelers' choices

in transportation (Ben-Akiva and Lerman 1985) and customers’ choices from a set of products in revenue management (Talluri and van Ryzin 2004). The classic multinomial logit (MNL) model has been widely used due to its tractability, but it has a number of shortcomings, including the independence from irrelevant alternatives (IIA) property, the assumption that each customer’s choice set is known, and the assumption that all customers have the same preferences or taste coefficients. To address these shortcomings a variety of other discrete choice models have been developed, such as the nested logit (NL) model and mixed logit (ML) models (Train 2003, Greene and Hensher 2003).

A number of recent papers, such as Rusmevichientong and Topaloglu (2012), Davis et al. (2013a,b), Gallego and Topaloglu (2013), Li et al. (2013), and Rusmevichientong et al. (2013), have addressed static assortment optimization problems under a variety of discrete choice models. In this paper we consider dynamic assortment optimization problems under MNL, NL, and ML choice models, where “dynamic” refers to the chosen assortments being a function of time.

3. The Optimization Problem

In this section we describe two optimization problems, a stochastic optimization problem and a deterministic optimization problem. The stochastic optimization problem is much too intractable to solve exactly, but is described here as the problem we would like to solve, and the objective values of which are computed by a simulation described in Section 6 and used to compare the performance of various policies. The deterministic optimization problem also involves probabilities, but it is a more tractable approximation of the stochastic optimization problem.

3.1. The Stochastic Optimization Problem

There are a number of airlines, indexed by $i \in I$, selling tickets for travel on parallel flights in a single origin-destination market. In this section, for ease of notation, airline XX for whom the optimization problem is solved is indexed by $i = 1$, and -1 denotes the competitors $\{YY, ZZ\}$. Airline i sells tickets for a set F_i of flights. Each flight $f \in F_i$ can accommodate up to b_f passengers. Airline i has a set J_i of flight-fare class combinations, also called products, that the airline can offer. Also, let $f(j) \in F_i$ denote the flight associated with flight-fare class combination j . Each unit of product $j \in J_i$ requires B_j units of capacity on flight $f(j)$ (or any other flight if the customer ends up traveling on another flight). The net revenue earned by selling a unit of product $j \in J_i$ is denoted by r_j . The selling horizon is denoted with $[0, T]$, where T denotes the scheduled departure time of the last flight in the time horizon. Let $T_f \in [0, T]$ denote the scheduled departure time of $f \in F_i$. For each airline i there is a set K_i of sales channels that can be used to sell the airline’s

products, for example, the airline's own web site, the airline's call center, various third party web sites, and independent travel agents. Some channels, such as an airline's own web site or call center, are used by only one airline, and some, such as third party web sites, are used by multiple airlines. Let $I_k \subset I$ denote the set of airlines that use channel $k \in K := \cup_{i \in I} K_i$.

Customer booking requests arrive in each channel k according to independent nonhomogeneous Poisson processes with rates $\lambda_k(t)$. For each time $t \in [0, T]$ and each sales channel $k \in K$, each airline $i \in I_k$ chooses a set $A_{i,k}(t) \subset J_i$ of flight-fare class combinations to offer; $A_{i,k}(t)$ is called the assortment offered in channel k at time t . Of course, products in $A_{i,k}(t)$ cannot belong to flights that depart before time t . (Specifically, the airline in the case study closes all flights 30 minutes before the scheduled departure time of the flight.) An airline may restrict attention to only certain assortments, for example, many airlines restrict attention to nested-by-fare assortments, as follows: Consider any two products j^1 and j^2 with $f(j^1) = f(j^2)$ and $r_{j^1} > r_{j^2}$. A nested-by-fare assortment A has the property that if $j^2 \in A$, then $j^1 \in A$. Let $\mathcal{A}_{i,k}(t) \subset 2^{J_i}$ denote the collection of assortments that airline i considers at time t for channel k . Let $A_k(t) := (A_{i,k}(t), i \in I_k)$ denote the list of assortments offered by the airlines in channel k at time t . For any list of assortments $A = (A_i, i \in I_k)$ in channel k , we sometimes write $\cup A$ for $\cup_{i \in I_k} A_i$. Given the offered assortments $A_k(t) = A$, a customer who arrives at time t using channel k books alternative j with probability $\mathbb{P}_{j:A}(k, t)$ ($\mathbb{P}_{j:A}(k, t) = 0$ if $j \notin \cup A$). Then, the customer purchases alternative j with probability $\mathbb{P}_{j,k}^{\text{b2t}}(t)$, called the booking-to-ticketing conversion probability.

The airline formulating and solving the revenue optimization problem does not know in advance what assortments the other airlines are going to offer. Here we assume that the airline estimates a probability distribution on the collection of assortments that the other airlines can offer at each time in the booking horizon. For airline i and channel $k \in K_i$, let $\mathcal{A}_{-i,k}(t) := \prod_{i' \in I_k \setminus \{i\}} \mathcal{A}_{i',k}(t)$ denote the collection of assortments that airlines other than i consider at time t for channel k , let $A_{-i,k}(t) := (A_{i',k}(t), i' \in I_k \setminus \{i\}) \in \mathcal{A}_{-i,k}(t)$ denote the list of assortments offered by the airlines other than i in channel k at time t , and let $\hat{\mathbb{P}}_{i,k}(t; \cdot)$ denote the probability distribution used by airline i for the competitors' assortments $A_{-i,k}(t)$.

Booking requests are indexed by $n \in \mathcal{N}$. The corresponding customer is referred to as customer n . In the context of the stochastic process, \mathcal{N} denotes the set of booking requests in $[0, T]$. Later we also refer to a data set of bookings used for model estimation, and then \mathcal{N} denotes the set of transactions in the data set, and n is the transaction index. Let t_n denote the arrival time of booking request n , let k_n denote the channel used, and let j_n denote the alternative chosen (booked)

by customer n , and let indicator ζ_n denote whether customer n pays for the ticket ($\zeta_n = 1$) or not ($\zeta_n = 0$).

Most customers with purchased tickets show up for their flights. The rest of the customers do not show up for their flights due to various actions by these customers. Four actions that have significant frequencies in our data set are modeled in this paper. They are (1) go-show before departure, (2) go-show after departure or no-show and change (not at the airport) to a later flight (these two actions are treated in the same way and are henceforth referred to as “go-show after departure”), (3) no-show with refund, and (4) no-show without refund. We now describe each of these four actions in detail.

Show-up for departure. The customer departs with the originally ticketed flight unless the flight is oversold. If the flight is oversold and the customer is bumped to a later flight, then the airline pays a penalty e_{j_n, k_n}^1 . Otherwise, if the flight is oversold and the customer does not get space on a later flight, then the airline pays a penalty e_{j_n, k_n}^0 .

Go-show before departure. A ticketed customer who go-shows before departure time shows up at an airport counter on the departure date and then requests a seat on an earlier flight. The time $T_n^{\text{gb}} < T_{f(j_n)}$ (in the case study, $T_n^{\text{gb}} < T_{f(j_n)} - 30$ minutes) at which customer n go-shows before departure time has cumulative distribution function $F_{j_n, k_n}^{\text{gb}}(t_n; \cdot)$. Thus $F_{j, k}^{\text{gb}}(t; T_{f(j)}) \leq 1$ can be thought of as the fraction of customers who use channel k at time t to book and pay for product j who go-show before the departure time of their flight. If there is a flight with an available seat scheduled to depart in time interval $(T_n^{\text{gb}}, T_{f(j_n)})$ (in the case study, $(T_n^{\text{gb}} + 30 \text{ minutes}, T_{f(j_n)})$), then the airline gives the customer a seat on such a flight, and the customer pays the change fee of c_{j_n, k_n}^{gb} . Otherwise the customer departs with the originally ticketed flight in which case the customer pays no change fee, or the customer is bumped to an even later flight and the airline pays a penalty e_{j_n, k_n}^1 , or the customer does not get space on a later flight, and the airline pays a penalty e_{j_n, k_n}^0 .

Go-show after departure. A ticketed customer who go-shows after departure time shows up at an airport counter after the departure time $T_{f(j_n)}$ of the customer’s originally ticketed flight (and thus such a customer no-shows for the customer’s originally ticketed flight), but before T on the departure date, and then requests a seat on a flight before T on the same day. The time $T_n^{\text{ga}} > T_{f(j_n)}$ (in the case study, $T_n^{\text{ga}} > T_{f(j_n)} - 30$ minutes) at which customer n go-shows after departure time has cumulative distribution function $F_{j_n, k_n}^{\text{ga}}(t_n; \cdot)$ on the interval $(T_{f(j_n)}, T)$. Thus $F_{j, k}^{\text{ga}}(t; T) \leq 1$ can be thought of as the fraction of customers who use channel k at time t to book and pay for product j who go-show after the departure time of their flight. If there is a flight with an available seat scheduled to depart after time T_n^{ga} (in the case study, $T_n^{\text{ga}} + 30$ minutes), then the

airline gives the customer a seat on such a flight, and the customer pays the change fee of c_{j_n, k_n}^{ga} . If there is no seat on any flight before T , then customer receives a refund of e_{j_n, k_n}^c .

No-show with refund. A ticketed customer no-shows for the customer's originally ticketed flight, and requests a refund after the flight departs. The customer receives a refund of e_{j_n, k_n}^c .

No-show without refund. A ticketed customer no-shows for the customer's originally ticketed flight, and never requests a refund.

For each customer n with $\zeta_n = 1$, let

$$\xi_n := \begin{cases} \text{s} & \text{if the customer shows up for the departure of the customer's ticketed flight,} \\ \text{gb} & \text{if the customer go-shows before departure of the customer's ticketed flight,} \\ \text{ga} & \text{if the customer go-shows after departure of the customer's ticketed flight,} \\ \text{c} & \text{if the customer no-shows with refund,} \\ \text{n} & \text{if the customer no-shows without refund.} \end{cases}$$

We assume that each customer makes independent decisions, and therefore ξ_1, ξ_2, \dots are independent. The probabilities \mathbb{P}_{j_n, k_n}^c and \mathbb{P}_{j_n, k_n}^n that $\xi_n = c$ or $\xi_n = n$ respectively are allowed to depend on the product j_n and the channel k_n . Given $j_n = j$ and $k_n = k$, probabilities $\mathbb{P}_{j, k}^c$ and $\mathbb{P}_{j, k}^n$ are assumed to be independent of n . Note that $1 - (F_{j, k}^{\text{gb}}(t; T_{f(j)}) + F_{j, k}^{\text{ga}}(t; T) + \mathbb{P}_{j, k}^c + \mathbb{P}_{j, k}^n)$ represents the fraction of customers who use channel k at time t to book and pay for product j who show up on time for their flight (but do not go-show before departure).

For each customer n with $\zeta_n = 1$, let f_n denote the flight on which customer n travels, where $f_n = 0$ denotes that customer n does not travel on any flight, and by convention $T_0 > T$. If $\xi_n = s$, then $T_{f_n} \geq T_{f(j_n)}$, and if $T_{f_n} > T_{f(j_n)}$ then it denotes that customer n was bumped off the customer's ticketed flight. If $\xi_n = \text{gb}$, then $T_{f_n} \geq T_n^{\text{gb}}$, if $T_{f_n} = T_{f(j_n)}$ then it denotes that customer n travels on the customer's originally ticketed flight in spite of the customer's go-show before departure time, and if $T_{f_n} > T_{f(j_n)}$ then it denotes that customer n was bumped off the customer's ticketed flight. If $\xi_n = \text{ga}$, then $T_{f_n} > T_n^{\text{ga}} > T_{f(j_n)}$. If $\xi_n = c$ or $\xi_n = n$, then $f_n = 0$.

For each $k \in K_1$, $t \in [0, T]$, and $A \in \mathcal{A}_{1, k}(t)$, let $u_{A, k}(t) = 1$ if airline 1 decides to offer assortment A to customers in channel k at time t , and $u_{A, k}(t) = 0$ otherwise. Decision variable $u_{A, k}(t)$ is allowed to be random, but must be a function of the information available to airline 1 at time t .

Then, the stochastic optimization problem can be formulated as follows:

$$\begin{aligned} \max_{u, f_n} \quad & \mathbb{E} \left[\sum_{n \in \mathcal{N}} \zeta_n \left\{ r_{j_n} + c_{j_n}^{\text{gb}}(k_n) \mathbf{1}_{\{\xi_n = \text{gb}\}} \mathbf{1}_{\{T_{f_n} < T_{f(j_n)}\}} + c_{j_n}^{\text{ga}}(k_n) \mathbf{1}_{\{\xi_n = \text{ga}\}} \mathbf{1}_{\{f_n \neq 0\}} \right. \right. \\ & \quad - e_{j_n, k_n}^1 (\mathbf{1}_{\{\xi_n = s\}} + \mathbf{1}_{\{\xi_n = \text{gb}\}}) \mathbf{1}_{\{T_{f(j_n)} < T_{f_n} \leq T\}} - e_{j_n, k_n}^0 (\mathbf{1}_{\{\xi_n = s\}} + \mathbf{1}_{\{\xi_n = \text{gb}\}}) \mathbf{1}_{\{f_n = 0\}} \\ & \quad \left. \left. - e_{j_n, k_n}^c (\mathbf{1}_{\{\xi_n = \text{ga}\}} \mathbf{1}_{\{f_n = 0\}} + \mathbf{1}_{\{\xi_n = c\}}) \right\} \right] \end{aligned}$$

$$\begin{aligned}
\text{s.t. } \quad & \sum_{A \in \mathcal{A}_{1,k}(t)} u_{A,k}(t) \leq 1 \quad \forall k \in K_1, t \in [0, T] \\
& T_n^{\text{gb}} < T_{f_n} \quad \forall n \in \mathcal{N} : \xi_n = \text{gb} \\
& T_n^{\text{ga}} < T_{f_n} \quad \forall n \in \mathcal{N} : \xi_n = \text{ga} \\
& \sum_{n \in \mathcal{N}} \zeta_n B_{j_n} \mathbf{1}_{\{f_n=f\}} \leq b_f \quad \forall f \in F_1 \\
& u_{A,k}(t) \geq 0 \quad \forall k \in K_1, t \in [0, T], A \in \mathcal{A}_{1,k}(t)
\end{aligned}$$

Note that j_n depends on $(u_{A,k_n}(t_n), A \in \mathcal{A}_{1,k_n}(t_n))$ and on $A_{-1,k_n}(t_n)$, and that the expectation $\mathbb{E}[\cdot]$ above is determined by all the probabilities, including the probabilities $\hat{\mathbb{P}}_{1,k}(t; \cdot)$ used by airline 1 to forecast the assortments offered by the other airlines.

3.2. The Deterministic Optimization Problem

In this section we present a deterministic or fluid optimization problem analagous with the stochastic optimization problem described in Section 3.1. Customers are modeled as a fluid that arrives in each channel k with rate $\lambda_k(t)$. Given the offered assortments $A_k(t) = A$, a fraction $\mathbb{P}_{j:A}(k, t)$ of these customers book alternative j , and a fraction $\mathbb{P}_j^{\text{b2t}}(k, t)$ of these customers purchase alternative j . As before, the airline estimates a probability distribution on the collection of assortments that the other airlines can offer at each time in the booking horizon. Specifically, let $\hat{\mathbb{P}}_{i,k}(t; A)$ denote the probability used by airline i that the competitors will choose assortment $A_{-i,k}(t) = A$ for $k \in K_i, t \in [0, T]$. (Thus this deterministic or fluid optimization problem involves various probabilities.)

For each $k \in K_1, t \in [0, T]$, and $A \in \mathcal{A}_{1,k}(t)$, let $u_{A,k}(t) = 1$ if airline 1 decides to offer assortment A to customers in channel k at time t , and $u_{A,k}(t) = 0$ otherwise. For each $j \in J_1, k \in K_1$, and $f \in F_1 \cup \{0\}$, let $y_{j,k,f}^s$ denote the number of customers who use channel k to book and pay for product j , who show up on time for their flight, and who end up traveling on flight f (recall that $f = 0$ denotes that the customers are bumped and do not get to travel); let $y_{j,k,f}^{\text{gb}}$ denote the number of customers who use channel k to book and pay for product j , who go-show before the departure time of their flight, and who end up traveling on flight f ; let $y_{j,k,f}^{\text{ga}}$ denote the number of customers who use channel k to book and pay for product j , who go-show after the departure time of their flight, and who end up traveling on flight f ; let $y_{j,k}^c$ denote the number of customers who use channel k to book and pay for product j , who no-show for their flight, and who get a cancellation refund; and let $y_{j,k}^n$ denote the number of customers who use channel k to book and pay for product j , who no-show for their flight, and who do not get a cancellation refund.

Then, the deterministic or fluid optimization problem can be formulated as follows:

$$\max_{u,y} \int_0^T \sum_{k \in K_1} \lambda_k(t) \sum_{A \in \mathcal{A}_{1,k}(t)} u_{A,k}(t) \sum_{A' \in \mathcal{A}_{-1,k}(t)} \hat{\mathbb{P}}_{1,k}(t; A') \sum_{j \in A} \mathbb{P}_{j:A \cup A'}(k, t) \mathbb{P}_j^{\text{b2t}}(k, t) r_j dt$$

$$\begin{aligned}
& + \sum_{j \in J_1} \sum_{k \in K_1} \left\{ \sum_{\{f \in F_1 : T_f < T_{f(j)}\}} c_{j,k}^{\text{gb}} y_{j,k,f}^{\text{gb}} + \sum_{\{f \in F_1 : T_f > T_{f(j)}\}} [c_{j,k}^{\text{ga}} y_{j,k,f}^{\text{ga}} - e_{j,k}^1 (y_{j,k,f}^{\text{s}} + y_{j,k,f}^{\text{gb}})] \right\} \\
& - \sum_{j \in J_1} \sum_{k \in K_1} \{ e_{j,k}^0 (y_{j,k,0}^{\text{s}} + y_{j,k,0}^{\text{gb}}) + e_{j,k}^c (y_{j,k,0}^{\text{ga}} + y_{j,k}^c) \} \\
\text{s.t. } & \sum_{A \in \mathcal{A}_{1,k}(t)} u_{A,k}(t) \leq 1 \quad \forall k \in K_1, t \in [0, T] \\
& \sum_{\{f' \in F_1 : T_{f'} \leq T_f\}} y_{j,k,f'}^{\text{gb}} \leq \int_0^{T_f} \lambda_k(t) \sum_{A \in \mathcal{A}_{1,k}(t)} u_{A,k}(t) \sum_{A' \in \mathcal{A}_{-1,k}(t)} \hat{\mathbb{P}}_{1,k}(t; A') \mathbb{P}_{j:A \cup A'}(k, t) \mathbb{P}_j^{\text{b}2t}(k, t) F_{j,k}^{\text{gb}}(t; T_f) dt \\
& \quad \forall j \in J_1, k \in K_1, f \in F_1 : T_f < T_{f(j)} \\
& \sum_{f \in F_1 \cup \{0\}} y_{j,k,f}^{\text{gb}} = \int_0^{T_{f(j)}} \lambda_k(t) \sum_{A \in \mathcal{A}_{1,k}(t)} u_{A,k}(t) \sum_{A' \in \mathcal{A}_{-1,k}(t)} \hat{\mathbb{P}}_{1,k}(t; A') \mathbb{P}_{j:A \cup A'}(k, t) \mathbb{P}_j^{\text{b}2t}(k, t) F_{j,k}^{\text{gb}}(t; T_{f(j)}) dt \\
& \quad \forall j \in J_1, k \in K_1 \\
& \sum_{\{f' \in F_1 : T_{f(j)} < T_{f'} \leq T_f\}} y_{j,k,f'}^{\text{ga}} \leq \int_0^{T_{f(j)}} \lambda_k(t) \sum_{A \in \mathcal{A}_{1,k}(t)} u_{A,k}(t) \sum_{A' \in \mathcal{A}_{-1,k}(t)} \hat{\mathbb{P}}_{1,k}(t; A') \mathbb{P}_{j:A \cup A'}(k, t) \mathbb{P}_j^{\text{b}2t}(k, t) F_{j,k}^{\text{ga}}(t; T_f) dt \\
& \quad \forall j \in J_1, k \in K_1, f \in F_1 : T_{f(j)} < T_f \\
& \sum_{\{f \in F_1 \cup \{0\} : T_{f(j)} < T_f\}} y_{j,k,f}^{\text{ga}} = \int_0^{T_{f(j)}} \lambda_k(t) \sum_{A \in \mathcal{A}_{1,k}(t)} u_{A,k}(t) \sum_{A' \in \mathcal{A}_{-1,k}(t)} \hat{\mathbb{P}}_{1,k}(t; A') \mathbb{P}_{j:A \cup A'}(k, t) \mathbb{P}_j^{\text{b}2t}(k, t) F_{j,k}^{\text{ga}}(t; T) dt \\
& \quad \forall j \in J_1, k \in K_1 \\
& y_{j,k}^c = \int_0^{T_{f(j)}} \lambda_k(t) \sum_{A \in \mathcal{A}_{1,k}(t)} u_{A,k}(t) \sum_{A' \in \mathcal{A}_{-1,k}(t)} \hat{\mathbb{P}}_{1,k}(t; A') \mathbb{P}_{j:A \cup A'}(k, t) \mathbb{P}_j^{\text{b}2t}(k, t) \mathbb{P}_{j,k}^c dt \quad \forall j \in J_1, k \in K_1 \\
& y_{j,k}^{\text{n}} = \int_0^{T_{f(j)}} \lambda_k(t) \sum_{A \in \mathcal{A}_{1,k}(t)} u_{A,k}(t) \sum_{A' \in \mathcal{A}_{-1,k}(t)} \hat{\mathbb{P}}_{1,k}(t; A') \mathbb{P}_{j:A \cup A'}(k, t) \mathbb{P}_j^{\text{b}2t}(k, t) \mathbb{P}_{j,k}^{\text{n}} dt \quad \forall j \in J_1, k \in K_1 \\
& \sum_{\{f \in F_1 \cup \{0\} : T_{f(j)} \leq T_f\}} y_{j,k,f}^{\text{s}} + \sum_{f \in F_1 \cup \{0\}} y_{j,k,f}^{\text{gb}} + \sum_{\{f \in F_1 \cup \{0\} : T_{f(j)} < T_f\}} y_{j,k,f}^{\text{ga}} + y_{j,k}^c + y_{j,k}^{\text{n}} \\
& = \int_0^{T_{f(j)}} \lambda_k(t) \sum_{A \in \mathcal{A}_{1,k}(t)} u_{A,k}(t) \sum_{A' \in \mathcal{A}_{-1,k}(t)} \hat{\mathbb{P}}_{1,k}(t; A') \mathbb{P}_{j:A \cup A'}(k, t) \mathbb{P}_j^{\text{b}2t}(k, t) dt \quad \forall j \in J_1, k \in K_1 \\
& \sum_{k \in K_1} \left(\sum_{\{j \in J_1 : T_{f(j)} \leq T_f\}} B_j y_{j,k,f}^{\text{s}} + \sum_{j \in J_1} B_j y_{j,k,f}^{\text{gb}} + \sum_{\{j \in J_1 : T_{f(j)} < T_f\}} B_j y_{j,k,f}^{\text{ga}} \right) \leq b_f \quad \forall f \in F_1 \\
& u_{A,k}(t) \geq 0 \quad \forall k \in K_1, t \in [0, T], A \in \mathcal{A}_{1,k}(t) \\
& y_{j,k,f}^{\text{s}} \geq 0 \quad \forall j \in J_1, k \in K_1, f \in F_1 \cup \{0\} \\
& y_{j,k,f}^{\text{gb}} \geq 0 \quad \forall j \in J_1, k \in K_1, f \in F_1 \cup \{0\} \\
& y_{j,k,f}^{\text{ga}} \geq 0 \quad \forall j \in J_1, k \in K_1, f \in F_1 \cup \{0\} \\
& y_{j,k}^c \geq 0 \quad \forall j \in J_1, k \in K_1 \\
& y_{j,k}^{\text{n}} \geq 0 \quad \forall j \in J_1, k \in K_1
\end{aligned}$$

4. Data

Recall that there are three major airlines that we call XX, YY, and ZZ, in the market considered in our project. The revenue management problem for airline XX is considered. For airline XX, we have both booking and ticketing data. It shows for each customer, how many tickets the customer booked

and how many the customer actually paid for. These data are used to estimate the booking-to-ticketing conversion models. For airline YY and ZZ, we have only booking data, which is sufficient because once a customer made a booking with a competitor airline, whether the customer pays for the ticket or not will not have an impact on airline XX's revenue. These booking data contain the values of various factors that are important for the estimation of the models discussed in Section 5.2.4.

Another type of data that are not usually available to researchers are availability data. Availability data shows a snapshot, typically once per day, of the assortment being offered by each airline at that time. The offer set sometimes changes during a day, and we also use customers' booking data to identify when such changes took place, and to construct the historical assortments $A_{i,k}(t)$ for each airline i in channel k as a function of time t . As discussed in Sections 5.3 and 5.2, availability data are used to forecast competitor behavior, and to construct a choice set for each customer used in the estimation of booking choice models.

We also have no-show and go-show data. For each no-show customer, the data indicate which of the following four types the customer is: (1) Go-show before departure, (2) Go-show after departure, (3) No-show with refund, and (4) No-show without refund. Section 3 in the online companion will discuss in more detail how this data is used to estimate no-show and go-show models.

We have the data described above for both 2011 and 2012. The 2011 data were used to calibrate the models that form the input of the deterministic optimization problem in Section 3.2. The 2012 data were used to calibrate the models that form the input of the simulation of the stochastic optimization problem in Section 3.1. As described in Section 6, the simulation was used to test how well the booking controls based on models estimated with 2011 data perform under 2012 customer choice behavior.

5. Estimated Models

In this section, we discuss the estimation of models that are used as input for the stochastic optimization problem presented in Section 3.1 and for the corresponding deterministic optimization problem presented in Section 3.2. We obtain these estimates using the data sets described in Section 4. Recall that the input parameters to our optimization models are arrival rates, competitor assortment selection probabilities, booking choice probabilities, booking-to-ticketing probabilities, and go-show/no-show probabilities. Since the arrival rates and these probabilities all depend on the booking time (and many other factors), the input parameters are really input functions of time. We assume that these functions are piecewise constant.

In Section 5.1, we estimate arrival rates. In Section 5.2, we estimate customers' booking choice probabilities with three discrete choice models. In Section 5.3, we estimate competitors' assortment selection probabilities. In Section 5.4, we estimate booking-to-ticketing conversion probabilities. We estimate on-show and go-show probabilities in the online companion.

5.1. Arrival Model

Our available data include the booking data (including the detailed booking time) of the whole market, but not data about customers who enquired about ticket prices but who did not book. Since airline XX is only one of three airlines in the market, and as we argued in the introduction, most of the customers in this market appear to be business customers who are not very price sensitive, the assortment provided by XX may not significantly affect the arrival rate to the whole market. In addition, in this market all potential customers have the opportunity to book a ticket and then not buy the ticket, and these data are available to us. Therefore, we think that most potential customers in the market make bookings, and are thus recorded in the available data, and hence we use the booking data of the market as the arrival data.

The arrival rate $\lambda_k(t)$ for each channel k and each hour t before departure is calculated as the average of the 39 observed arrivals in 2011 for the same channel and hour before departure. Figure 2 shows the hourly arrival rates for Monday flights. For confidentiality, the arrival data are scaled by a factor. As we can see from the plot, most bookings (77.5% in this example) arrive within one week of the departure day. The difference between weekday arrivals and weekend arrivals and the fact that most of the bookings happen during the work day (with a noticeable dip during the lunch hour) also suggest that most of the customers are business customers.

5.2. Booking Choice Models

In this section we introduce three customer choice models to determine the choice probabilities $\mathbb{P}_{j:A}(k, t)$ introduced in Section 3.1. These models express the choice probabilities as a function of attributes that are functions of observed factors, including customer-specific factors such as booking channel and booking time, as well as alternative-specific factors such as price, change fee and cancellation fee, frequent flyer mileage gain, departure times, etc.

Different choice models assume the availability of different amount customer information. We now give a brief overview. Each customer who makes a booking request is associated with three types of data. The first type of data is observed by the airline before the airline decides which products to offer the customer. These data are called the customer class. Both the model of customer behavior as well as the booking control may depend on the customer class. An important example of such a

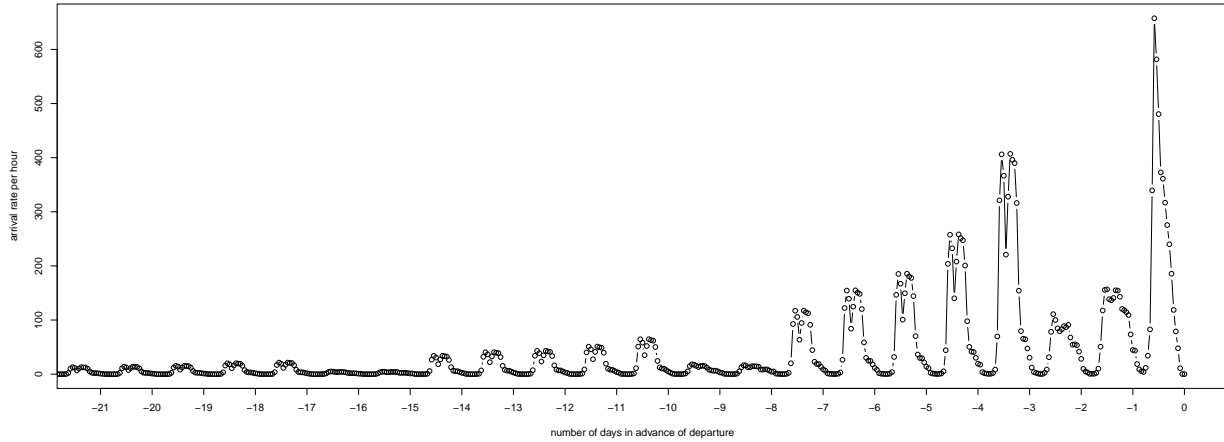


Figure 2 Hourly arrival rates for Monday flights

type of data is the following: The sales channel used by a customer is observed by an airline when the customer makes a request, and is therefore part of the customer class. As another example, in some cases the home base of the customer can be derived from the customer's request: if the customer requests a booking from A to B and B to A , then the airline considers the home base of the customer to be A . In the network setting, both origin A and destination B are data that are revealed by the customer when the request is made and that form part of the customer class. A third example is the customer's frequent flyer data, if the customer reveals a frequent flyer number when the request is made. For the setting in this paper involving a single domestic market, with no available frequent flyer data, the most important data in the customer class are the sales channel used, and therefore we do not introduce notation for the customer classes besides the notation for the sales channels. The second type of data associated with a customer request is the data that are not observed before the airline decides which products to offer the customer, but that are observed after the airline decides which products to offer the customer. The model of customer behavior may depend on these data, but the booking control may not depend on these data. An example is the customer's frequent flyer data, if the customer reveals a frequent flyer number after the airline makes an offer. In our study no such data were available. The third type of data is not observed by the airline. This includes a quantity called the customer type, as well as a quantity called the customer's utility for each alternative, both of which are discussed in more detail at the beginning of Section 5.2.3. Neither the model of customer behavior nor the booking control may depend on customer type data, but as described later, we estimate probability distributions for customer types using observed data.

In Sections 5.2.1-5.2.3, we introduce three customer choice models. The factors and attributes that are common to the different booking choice models are discussed in Section 5.2.4. Then the estimation results are compared and discussed in Section 5.2.6. Recall from Section 3.1 that for a customer n from channel k_n at booking time t_n , $A_{k_n}(t_n)$ denotes the list of assortments offered by all airlines to the customer. For ease of notation, let $A_n = A_{k_n}(t_n)$. We refer to A_n as the choice set of customer n . Note that it has to hold that the customer's actual choice j_n must be in A_n . For all three models, the attributes are indexed by $m \in \{1, 2, \dots, M\}$.

5.2.1. The Multinomial Logit (MNL) Model One of the most popular discrete choice models is the multinomial logit model. In this model, there is only one customer type. It is one of the shortcomings of the MNL model, but it contributes to the tractability of the MNL model. Let $x_{n,j,m}$ denote the value of attribute m for customer n and alternative j , and let β_m denote the weight of attribute m . Let $\beta := (\beta_1, \dots, \beta_M) \in \mathbb{R}^M$ denote the parameter vector, and let $x_{n,j} := (x_{n,j,1}, \dots, x_{n,j,M}) \in \mathbb{R}^M$ denote the attribute vector for customer n and alternative j . Then $v_{n,j} := \beta^\top x_{n,j}$ is called the systematic utility of customer n for alternative j . Since there is only one customer type, β is the same for all customers. Variation in preferences across the population is represented by i.i.d. additive Gumbel distributed random utilities $\xi_{n,j}$ for each customer n and alternative j ; for more detail of the MNL model, see for example Ben-Akiva and Lerman (1985) and Train (2003). Then the probability that customer n chooses alternative $j \in \cup A_n$ is given by

$$\mathbb{P}_{j:A_n}(k_n, t_n) = \frac{\exp(v_{n,j})}{\sum_{j' \in \cup A_n} \exp(v_{n,j'})} = \frac{\exp(\beta^\top x_{n,j})}{\sum_{j' \in \cup A_n} \exp(\beta^\top x_{n,j'})}. \quad (1)$$

Another shortcoming of the MNL model is that it assumes that all customers n choose from the assortments $A_{i,n}$ offered by airline i . In the context of airline demand (and many other applications), different customers consider different sets of alternatives, and the true choice sets considered by customers are not observed (but some data related to the consideration sets may be observed). For example, different customers consider different sets of departure times to be reasonable for their purposes. Some customers are flexible and may consider all flights in a wide time window, whereas other customers have tight schedules and want to depart as close as possible to a specific time. These time preferences are not observed.

The following modeler's selection of $A_{i,n}$ was suggested in Vulcano et al. (2010): Given that customer n booked a ticket from A to B for a flight departing on a particular day, $A_{i,n}$ is the set of all flights of airline i traveling from A to B on the same day. We used the same selection of $A_{i,n}$ for the MNL model results discussed in Section 5.2.6. However, the following intuitive argument suggests that such a selection may produce biased parameter estimates. Suppose that

the price of an alternative is an important attribute of the alternative. More specifically, suppose that each customer chooses the cheapest ticket for a flight that departs in the customer's interested time window. Thus customers are quite price sensitive, with attention restricted to a subset of alternatives. Now suppose that flights departing at different times of the day have different cheapest available fares (which is often the case). In a data set of bookings, a significant fraction of customers do not choose one of the cheapest tickets over all flights departing on the particular day (because none of the cheapest tickets were for a flight departing in the customers' time windows). If it is assumed that each customer chooses from the set of all flights on the same day, then it appears that customers are not very price sensitive, and as a result the estimated coefficients of price attributes will be biased. As shown in Section 5.2.6, our results were consistent with this intuition. Next we discuss a number of models that attempt to incorporate heterogeneity in customer preferences.

5.2.2. The Nested Logit (NL) model In the nested logit model, similar to the MNL model, there is only one customer type. However, the NL model can capture preferences for subsets of alternatives in a way that the MNL model cannot. Specifically, in the nested logit model, the set of alternatives is partitioned into subsets called nests, indexed by $l \in \{1, 2, \dots, L\}$. For example, different nests contain tickets for flights departing during different time windows. Correspondingly, for each customer n , $\cup A_n$ is partitioned into L nests denoted with $A_{n,l}$. In the NL model, variation in preferences across the population is represented by additive random utilities $\xi_{n,j} + \xi_{n,l}$ for each customer n and alternative $j \in A_{n,l}$, where $\xi_{n,j}$ are i.i.d. Gumbel and $\xi_{n,l}$ are i.i.d. Gumbel. Note that in this model, different alternatives in the same nest have correlated utilities. For example, if the value of $\xi_{n,l}$ is large, then customer n has a preference for alternatives in nest l , and vice versa. Thus, by choosing different nests to contain tickets for flights departing during different time windows, the NL model can capture random preferences for different departure times. For example, suppose that the nests represent departure time windows. A customer with a strong preference for departure time window $l = 1$ has a large value of $\xi_{n,l}$ for $l = 1$ and small values of $\xi_{n,l}$ for other values of l . A customer with a preference for either departure time window $l = 1$ or $l = 2$ but a strong dislike for other departure time windows has large values of $\xi_{n,l}$ for $l = 1$ and $l = 2$ and small values of $\xi_{n,l}$ for other values of l . A restriction of the NL model is that the set of alternatives has to be partitioned, for example, the NL model cannot capture a setting in which some customers prefer departure times between t_1 and t_3 , some prefer departure times between t_2 and t_4 , and some prefer departure times between t_3 and t_5 , where $t_1 < t_2 < t_3 < t_4 < t_5$. For more detail of the NL model, see for example Ben-Akiva and Lerman (1985) and Train (2003). The systematic utility of customer n for alternative $j \in A_{n,l}$ is given by $v_{n,j} := \beta^\top x_{n,j} / \alpha_l$, where $\alpha_l \in [0, 1/\alpha]$ is the parameter

that represents the variation of preferences for alternatives in $A_{n,l}$, and $\alpha > 0$ is a scaling factor. Then the probability that customer n chooses alternative $j \in A_{n,l}$ is given by

$$\begin{aligned}\mathbb{P}_{j:A_n}(k_n, t_n) &= \frac{\exp(v_{n,j})}{\sum_{j' \in A_{n,l}} \exp(v_{n,j'})} \frac{\exp(\alpha \alpha_l \bar{v}_{n,l})}{\sum_{l'=1}^L \exp(\alpha \alpha_{l'} \bar{v}_{n,l'})} \\ &= \frac{\exp(\beta^\top x_{n,j} / \alpha_l)}{\sum_{j' \in A_{n,l}} \exp(\beta^\top x_{n,j'} / \alpha_l)} \frac{\exp(\alpha \alpha_l \bar{v}_{n,l})}{\sum_{l'=1}^L \exp(\alpha \alpha_{l'} \bar{v}_{n,l'})},\end{aligned}$$

where

$$\bar{v}_{n,l} := \ln \left(\sum_{j \in A_{n,l}} \exp(\beta^\top x_{n,j} / \alpha_l) \right), \quad \forall l \in \{1, \dots, L\}.$$

5.2.3. The Mixed Logit (ML) model The attribute coefficients β introduced in Section 5.2.1 reflect the tastes of customers in evaluating attributes. A natural way for a choice model to capture taste variation is to allow variation in the values of β . Thus, in the mixed logit model there are multiple customer types, and different customer types have different values of the parameter vector β . We assume customers who request bookings are classified into a number of types. Let \mathcal{T} denote the set of customer types. The type of a customer is not observed by the airline. We assume that a customer who arrives at time t using channel k has type $\tau \in \mathcal{T}$ with probability $\pi_{k,\tau}(t)$. The customer's type is determined before the customer observes the airlines' assortments, and thus $\pi_{k,\tau}(t)$ is independent of the assortments. A classical example of customer type is the following: A customer is either a "business customer" or a "leisure customer". These type labels are somewhat misleading, since the customer's booking behavior determines the customer's type rather than the purpose of the customer's trip. For example, a business customer has a desired departure time, and wants a fully refundable ticket for the flight with departure time closest to the desired departure time, irrespective of the prices of other alternatives. On the other hand, a leisure customer has a more subtle trade-off among alternatives, taking into account prices, departure times, and other attributes. Another important example of customer type is the following: A customer may not like all products in P_i equally, even if all were offered, but rather prefers flights that depart between a time window $[t_1, t_2]$ (e.g., all morning flights between $t_1 = 07:00\text{am}$ and $t_2 = 11:00\text{am}$) to the flights with departure times more separated from $[t_1, t_2]$ (e.g., late night flights between $21:00\text{pm}$ and $00:00\text{am}$). Then the type τ of the customer determines the preferences of the customer in evaluating departure times.

The type distribution $\pi_{k,\tau}(t)$ determines the distribution of β ; for the ML model we write $\pi_{k,\beta}(t)$. The distribution $\pi_{k,\beta}(t)$ may be discrete or continuous, and is usually selected from a parameterized family. An advantage of the ML is that one can approximate different true choice sets considered by different customer types by including random alternative-specific coefficients β_j for alternatives j ,

where a value of $\beta_j < -M$ for large M in effect removes alternative j from the customer type's true choice set. For example, to model departure time preferences, one can partition the departure times into a number of time windows, and include coefficients β_w for time windows w . A large mean of β_w corresponds to time windows w that are on average more popular, a large variance of β_w corresponds to time windows w that some customers strongly like and other customers strongly dislike, and a large positive covariance of β_w and $\beta_{w'}$ corresponds to pairs of time windows w and w' with similar preferences — some customers like both and other customers dislike both. For more detail of the ML model, see for example Train (2003). The systematic utility of customer n for alternative j is given by $v_{n,j} := \beta^\top x_{n,j}$, and is random (given the vector $x_{n,j}$ of attribute values) with distribution determined by $\pi_{k_n, \beta}(t_n)$. Then the probability that customer n chooses alternative $j \in \cup A_n$ is given by

$$\mathbb{P}_{j:A_n}(k_n, t_n) = \mathbb{E}_{\pi_{k_n, \beta}(t_n)} \left[\frac{\exp(v_{n,j})}{\sum_{j' \in \cup A_n} \exp(v_{n,j'})} \right] = \mathbb{E}_{\pi_{k_n, \beta}(t_n)} \left[\frac{\exp(\beta^\top x_{n,j})}{\sum_{j' \in \cup A_n} \exp(\beta^\top x_{n,j'})} \right]. \quad (2)$$

5.2.4. Factors and Attributes Each of the booking choice models described above expresses the choice probabilities $\mathbb{P}_{j:A_n}(k_n, t_n)$ as a function of attribute values $x_{n,j,m}$, where $x_{n,j,m}$ denotes the value of attribute m for customer n and alternative j . Each of these attributes is a function of the observed values of several variables called factors. Table 1 lists the alternative-specific factors (1–5) and the customer-specific factors (6–7) on which we obtained data and that we used to encode attributes and estimate the discrete choice models.

Table 1 Alternative-specific and customer-specific factors used to encode attributes and estimate the discrete choice models.

	Factor	Description
1	Ticket price	the ticket fare, e.g., \$1350
2	Departure time	the time when a flight takes off, e.g., 09:00
3	Ticket change fee	the fee charged for changing to another flight, e.g., \$75
4	Mileage gain	the mileage credits earned by a customer if the customer buys the ticket, e.g., 1140 points
5	Carrier	the airline that sells tickets, e.g., XX, YY, ZZ
6	Booking time	the date, hour, minute at which the booking was made, e.g., Tuesday 2011-06-07 09:20
7	Booking channel	the channel via which a ticket is booked, e.g., airline website, call center

We discuss the use of the factors in Table 1 in the online companion. Due to a lack of good data, other studies such as Coldren et al. (2003) and Vulcano et al. (2010) estimated a single price coefficient for all customers. However, data on customer-specific factors such as booking time

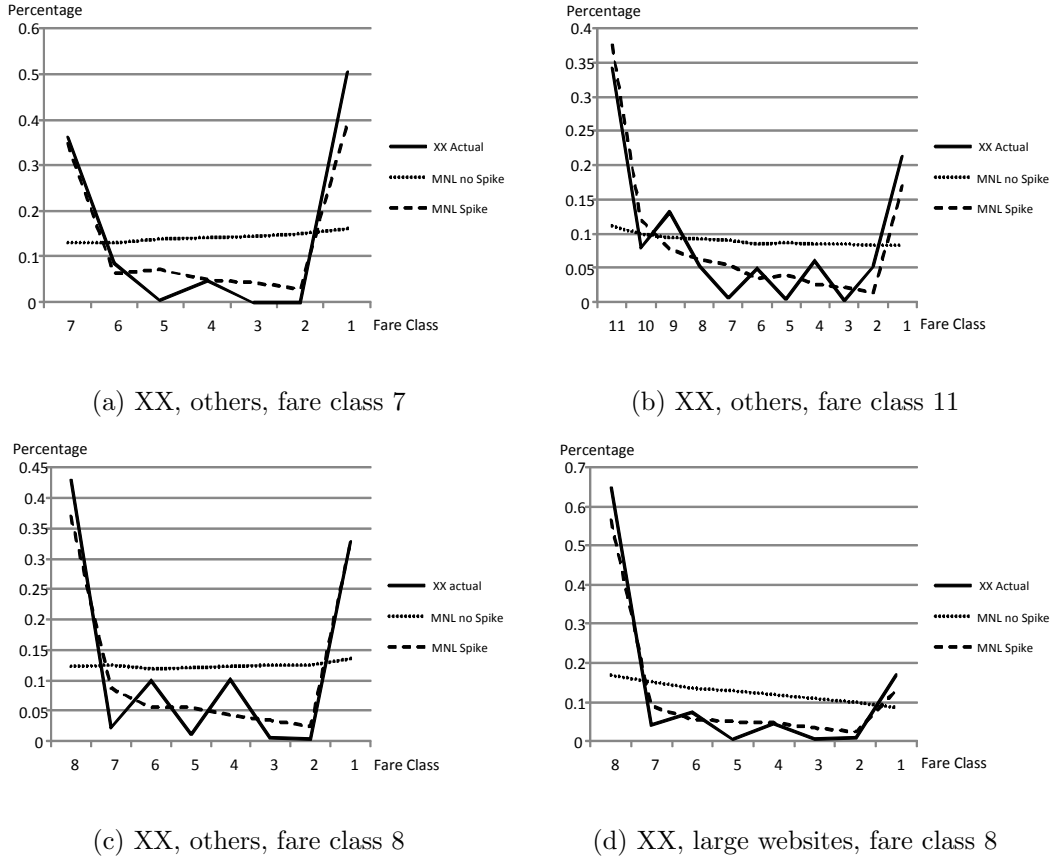


Figure 3 Comparison of distributions of XX bookings over fare classes of a flight via a particular booking channel when a specific fare class is the cheapest available one for the same flight and channel

and booking channel allow us to study the effect of these factors on price sensitivity. Thus, in each of the estimated models, for each combination of the days-until-departure, booking time-of-day, booking weekday/weekend, and booking channel factors described above, there is a separate attribute $x_{n,j,m}$ that is equal to the product of the price of alternative j and the indicator for the specific combination, and a corresponding price sensitivity coefficient β_m .

Each curve marked “XX actual” on Figure 3 represents the distribution (in percentage) of customers who booked a more expensive ticket for some fare class of a XX flight through a particular channel when a specific fare class is in actual the cheapest available one for the same flight and booking channel. The curve is plotted based on the actual booking data for airline XX. For example, Figure 3a shows that when fare class 7 is the cheapest available fare class for a flight, about 35% of customers who make a booking for that flight book in fare class 7, and about 50% of customers who make a booking for that flight book in fare class 1. It can be seen that a large portion of customers either book the most expensive tickets (from fare class 1) or the cheapest tickets no matter what fare class is the cheapest available one. We call this phenomenon “spikes”. Figure 3

indicates that a good choice model should be able to capture spikes. Meanwhile, the distribution of booked tickets over the fare classes is dependent on what fare class is the cheapest available one (as indicated by Figures 3a and 3b) and what channel is used by customers (as indicated by Figures 3c and 3d). The phenomenon of spikes, to some extent, reflects the competition dynamics between fare classes for each combination of carrier and booking channel. There are some other researchers who have been aware of the problem. As emphasized by Coldren and Koppelman (2005), there may exist prominent competition dynamics between fare classes within a flight and cross all the flights on the specific departure day. Coldren and Koppelman (2005) did not model any of these competition dynamics due to limited data. However with the good data, we are able to capture spikes by using appropriate attributes. To this end, we consider the following two combinations of factors as attributes: (i) the combination of carrier, the most expensive fare class (i.e., fare class 1 in our project) and booking channel and (ii) the combination of carrier, fare class, channel and whether or not the fare class is the cheapest available one for a flight.

For each of combination (i), there is a separate attribute $x_{n,j,m}$ that is equal to an indicator for the combination and a coefficient parameter β_m . Note that carrier XX has observations from all the five channels but carriers YY and ZZ only have observations from three channels 1, 2 and 5. We also consider a separate indicator for each combination (ii) and an attribute parameter for the indicator.

To further illustrate the importance of capturing spikes in a choice model, Figure 3 also shows the distributions of bookings over fare classes, which are plotted based on the booking data simulated by using the MNL model with combinations (i) and (ii) and the data predicted by the same choice model without the combinations. The curves obtained by using the NL and ML models are very close to the curves shown on Figure 3 and are thus omitted. It can be observed that the model capturing spikes behaves well enough to describe the booking behaviors of customers.

5.2.5. Description of Estimated Choice Models In this section we describe in detail various choice models that were estimated from the dataset of 2011. There are 89 price sensitivity parameters, 11 parameters for each of combinations (i), 102 parameters for each of combinations (ii), two parameters for carriers, one parameter for mileage gain and one for cancel fee, which gives 206 common parameters to all the four choice models. Combinations (i) and (ii) were introduced at the end previous section. For all the following three choice models, we take departure times in time window [21:00, 07:00) as the base case.

For the MNL model, there are 14 parameters to be estimated for attributes of departure times, which gives 220 parameters to be estimated.

For the NL model, we partition all the alternatives in A_n , where $n \in \mathcal{N}$, into $L = 3$ nests. Nest $l = 1$ contains the alternatives with departure times 1-4, nest $l = 2$ contains the alternatives with departure times 5-10, and nest $l = 3$ consists of the alternatives with departure times 11-15. We set the scaling factor $\alpha = 10^{-4}$ and need to estimate dissimilarity factors α_l , $l = 1, 2, 3$, 206 common parameters and 14 parameters for departure times.

For the ML model, to capture the variation of customer preferences in evaluating departure times, we consider a random parameter vector $\gamma_n := (\gamma_{n,1}, \gamma_{n,2}, \dots, \gamma_{n,15})$ for departure times for each customer $n \in \mathcal{N}$, where $\{\gamma_n\}_{n \in \mathcal{N}}$ is a sequence of i.i.d. Gaussian vectors with mean vector $\mu \in \mathbb{R}^{15}$ and covariance matrix $\Sigma \in \mathbb{R}^{15 \times 15}$. We can represent $\gamma_n = \mu + \sigma \xi_n$, where $\xi_n \in \mathbb{R}^{15}$ is a standard Gaussian vector and $\sigma \in \mathbb{R}^{15 \times 15}$ is the lower-triangular Cholesky factor such that $\Sigma = \sigma \sigma^\top$. Let $\gamma \in \mathbb{R}^{206}$ denote the 206 common parameters that are assumed to be deterministic and the same across the customer population \mathcal{N} . Let $x_{n,j} \in \mathbb{R}^{206}$ denote the vector of attribute values except departure times and $y_{n,j} \in \mathbb{R}^{14}$ denote the vector of attribute values for departure times for alternative $j \in A_n$. The systematic utility is written as $v_{n,j} = \gamma^\top x_{n,j} + \gamma_n^\top y_{n,j} = \gamma^\top x_{n,j} + \mu^\top y_{n,j} + \xi_n^\top \sigma^\top y_{n,j}$. We need to estimate (γ, μ, σ) by maximizing the simulated log-likelihood function

$$\max_{\gamma, \mu, \sigma} \frac{1}{|\mathcal{N}|} \sum_{n \in \mathcal{N}} \ln \left(\frac{1}{B} \sum_{i=1}^B \frac{\exp(\gamma^\top x_{n,j_n} + \mu^\top y_{n,j_n} + (\xi_n^i)^\top \sigma^\top y_{n,j_n})}{\sum_{j \in A_n} \exp(\gamma^\top x_{n,j} + \mu^\top y_{n,j} + (\xi_n^i)^\top \sigma^\top y_{n,j})} \right),$$

where B is the Monte Carlo integration sample size and $\{\xi_n^i\}_{i=1}^B$ is a sequence of i.i.d. standard Gaussian variates and independent for each n .

5.2.6. Estimation Results We describe in this section the estimation results for the three choice models, including the price coefficient parameters that reflect customers price sensitivity as well as the parameters for departure times that reflect the popularity of different departure time windows.

Price Sensitivity. To better understand the importance of modeling the competition dynamics within and across flights (or spikes), we estimate the models with combinations (i) and (ii) and those without the combinations. The first result we learn from the estimation is that all the price coefficients are negative for the MNL model with the combinations, which is consistent with our intuition that the more expensive the ticket is, the less customers who will choose it. The MNL model without the combinations, however, have some unexpected positive price coefficients. The same phenomenon is meanwhile observed in the estimation results for the NL and ML models. It is thus more reasonable to use the choice models with combinations (i) and (ii) and our discussions are all based on the results for the choice models with the two combinations.

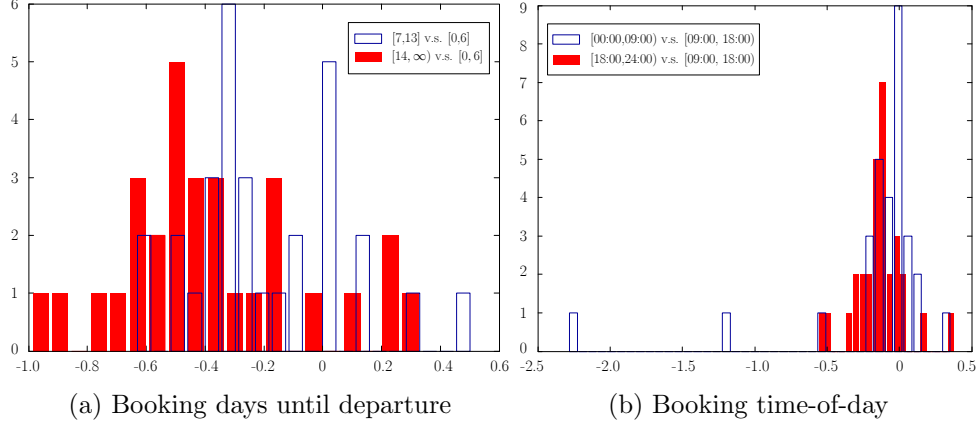


Figure 4 Histograms of relative price coefficients for booking days until departure and booking time-of-day

As expected, the estimation results show that customers' price sensitivity is dependent on booking days until departure, booking time-of-day, booking day-of-week and channel. To illustrate the dependences, we choose not to list all the 90 estimated price coefficients, but instead use the histogram of relative price coefficients for each of the four factors. Figure 4a shows the the histogram of price coefficients for booking days until departure $[7, 13]$ and $[14, \infty]$ relative to the price coefficients for $[0, 6]$. A relative price coefficient is computed as the scaled difference between a price coefficient for $[7, 13]$ or $[14, \infty)$ and a price coefficient for the base case $[0, 6]$ provided that the two prices have the same values in booking time-of-day, booking day-of-week and channel. For example, the price coefficient of $([7, 13], [00:00, 09:00), \text{weekday}, \text{others})$ and that of $([0, 6], [00:00, 09:00), \text{weekday}, \text{others})$ are -7.2397 and -4.8802 , respectively. As a result, the price coefficient of $([7, 13], [00:00, 09:00), \text{weekday}, \text{others})$ relative to $([0, 6], [00:00, 09:00), \text{weekday}, \text{others})$ is $\frac{-7.2397 + 4.8802}{7.2397} = -0.3259$. As Figure 4a shows, most of the bars are located at the left of zero, indicating a trend that the customers who book earlier are more price sensitive than the customers who book later. Figure 4b shows the histogram of relative price coefficients for booking time-of-day $[00:00, 09:00)$ and $[18:00, 24:00)$ to the base case $[09:00, 18:00)$. As Figure 4b indicates, the customers who book after and before work are more price sensitive than the customers who book during work hours.

Figure 5a shows the histogram of price coefficients for weekend relative to the price coefficients for weekday. Again, most of the bars are located at the left of zero, which indicate that the customers who book in the weekend are more price sensitive than those who make bookings on weekdays. The results are consistent with our intuition. Figure 5b shows the histograms of price coefficients for channels, others, airline website, large and small websites, relative to the price coefficients for

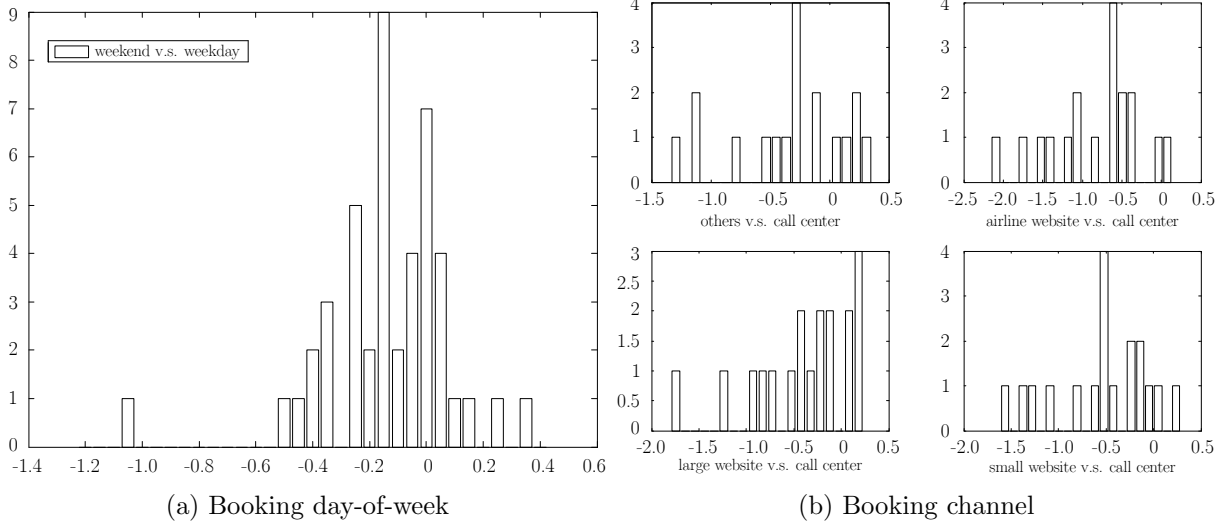


Figure 5 Histograms of relative price coefficients for booking day-of-week and channel

call center. It can be observed that the customers who book through the four channels are more price sensitive than the customers who book via call center.

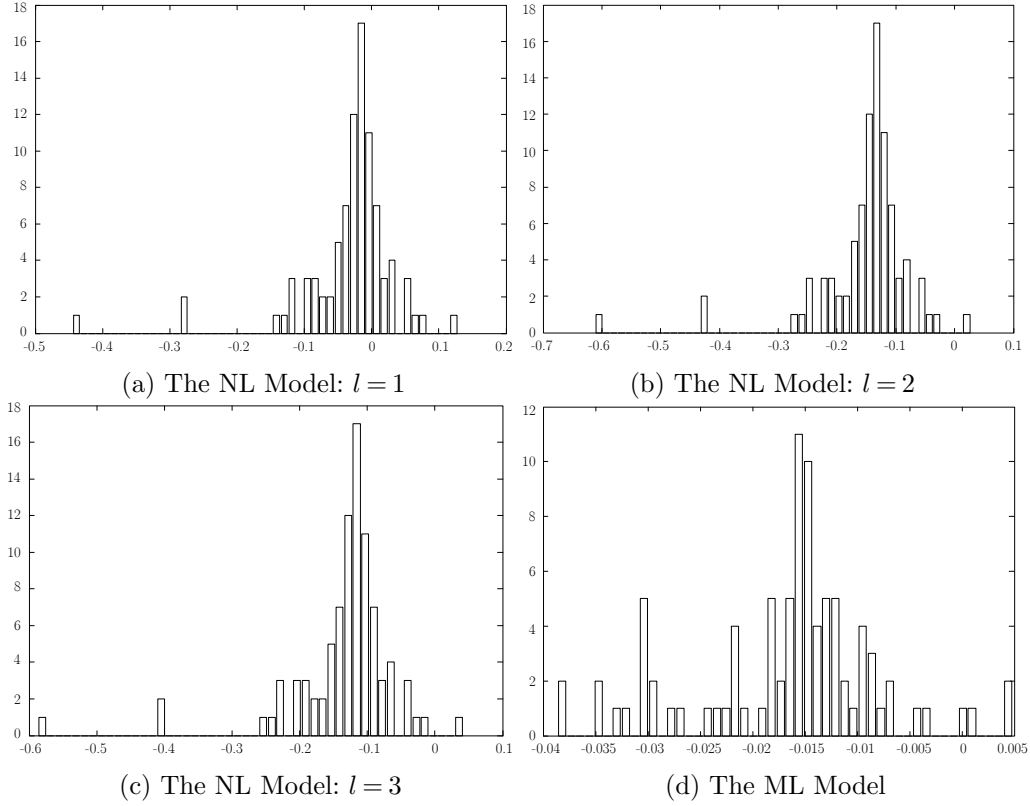


Figure 6 Histograms of relative price coefficients for the NL and ML models

Figure 6 shows the histograms of price coefficients estimated for the NL and MNL models relative

to the price coefficients for the MNL model. The relative price coefficients for the ML model are computed as the scaled difference between the price coefficients for the ML model and the price coefficients for the MNL. For the NL model, the relative price coefficients are computed as the scaled differences between $\beta^{\text{NL}}/\alpha_l$ and β^{MNL} , where $l = 1, 2, 3$, and β^{NL} and β^{MNL} are the estimated vectors of price coefficients for the NL and MNL models, respectively. Figure 6 shows that the price coefficients for the NL and ML models are more negative than the price coefficients for the MNL model (most bars are located at the left of zero), which indicates that the NL and ML models can better capture the variation in preferences of customers in departure times.

Departure Time Popularity. Table 2 shows the parameter coefficients for departure times with $[21:00, 07:00)$ as the base case. The first column shows the index of each departure time window. As Table 2 shows, the MNL, NL and ML model estimation results indicate that the flights departing during time windows 1,3,9 are the most popular among all the flights.

Table 2 Departure time coefficients for the MNL, NL and ML models

Index	Time window	MNL	NL	ML
1	[07:00, 08:00)	1.09650	3.02554	1.13200
2	[08:00, 09:00)	0.81833	2.84397	0.86406
3	[09:00, 10:00)	0.99187	3.03922	1.04250
4	[10:00, 11:00)	0.94432	3.02132	0.99750
5	[11:00, 12:00)	0.71482	3.07456	0.83717
6	[12:00, 13:00)	0.62374	2.94590	0.75148
7	[13:00, 14:00)	0.43684	2.75186	0.56319
8	[14:00, 15:00)	0.79139	3.17591	0.93265
9	[15:00, 16:00)	0.93135	3.33490	1.05930
10	[16:00, 17:00)	0.91611	3.29911	1.05450
11	[17:00, 18:00)	0.81184	0.78556	0.94122
12	[18:00, 19:00)	0.57137	0.64439	0.62031
13	[19:00, 20:00)	0.76000	0.74791	0.86083
14	[20:00, 21:00)	0.40656	0.42899	0.43594

Figure 7 shows a decreasing trend between correlation coefficients of departure time windows and the distances between indices of departure time windows for the estimated ML model. As the figure indicates, as two flights depart with a bigger time gap, the less correlated the two departure times are. In other words, customers who choose a particular flight would prefer the flights with departure times closer to the chosen flight to those with departure times more separated from the chosen one.

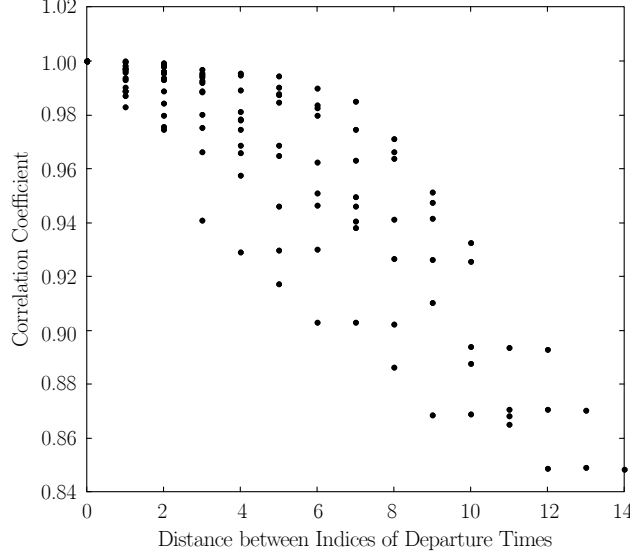


Figure 7 Correlation coefficient v.s. distance between indices of departure times for the ML model

5.3. Competitor Model

The booking choice models discussed in Section 5.2 enables us to compute choice probabilities of the form $\mathbb{P}_{j:A_i \cup A_{-i}}(k, t)$. In the optimization problem, A_i is determined by the decision variables of airline i , but A_{-i} is forecasted by airline i using a probability distribution $\hat{\mathbb{P}}_{i,k}(t; A_{-i})$.

First consider the easier case: Suppose that channel k is exclusive to airline i , such as the web site or call center of airline i . In that case $\hat{\mathbb{P}}_{i,k}(t; \emptyset) = 1$ for all t .

Next consider the case in which channel k is shared among the airlines, such as third party web sites. In that case, airline i uses the empirical distribution based on observed data of competitors' assortments for departures on a particular day of the week offered in channel k at times t in the booking horizon for $\hat{\mathbb{P}}_{i,k}(t; A_{-i})$. For example, airline XX has daily snapshots of competitors' assortments offered in channel k for 39 weeks in 2011 (April to December). Thus, for any given departure day of the week (for example Monday departures), and for any number of days before departure (for example the Friday 3 days before departure), airline XX has 39 observations in the data set of the assortments that airlines YY and ZZ offered at this time t in the booking horizon (for example, 39 Fridays in the data set on which airlines YY and ZZ offer assortments in channel k for departure on the following Monday). Hence $\hat{\mathbb{P}}_{1,k}(t; A) = 1/39$ for each shared channel $k \in K_1$, each $t \in [0, T]$, and each competitors' assortment $A \in \mathcal{A}_{-1,k}(t)$ observed in the data set for channel k and time t . Note that airline XX uses the joint empirical distribution for the assortments of airlines YY and ZZ, because one airline's offered assortment may influence the other airline's offered assortment, especially when one airline (YY) is larger than the other (ZZ), as is the case in the considered market.

Therefore the probability $\bar{\mathbb{P}}_{j:A_i}(k, t)$ that a customer who uses shared channel k at time t in the booking horizon chooses alternative $j \in A_i$ is given by

$$\bar{\mathbb{P}}_{j:A_i}(k, t) = \sum_{A_{-i} \in \mathcal{A}_{-i,k}(t)} \hat{\mathbb{P}}_{i,k}(t; A_{-i}) \mathbb{P}_{j:A_i \cup A_{-i}}(k, t).$$

5.4. Booking-to-Ticketing Conversion Model

When customers finally chooses a ticket out of the consideration set of choices and make the booking, they do not need to provide the payment information immediately. For each booking, there is a time limit before which the price at the booking will be guaranteed. If the customers do not make the payment before the time limit, the booking will be cancelled automatically without any charge to the customers. When a booking finally gets paid by the customer, we call it a ticketing.

In 2011 and 2012, the booking-to-ticketing ratios are 50.79% and 51.22%, respectively. These low ratios imply that our model has to distinguish between booking and ticketing. To model the booking-to-ticketing conversion, we will use the binary logistic regression model. In the following, we show the relationship between the booking-to-ticketing ratio and various booking attributes.

Figure 8a plots the booking-to-ticketing ratios for different fare classes in 2011 and 2012. The fare classes are sorted by their prices in descending order. It can be seen from the figure that the booking-to-ticketing varies across different fare classes. Generally the customers who book the cheaper tickets (on the right of the figure) are less likely to purchase their tickets finally than those who book the expensive tickets.

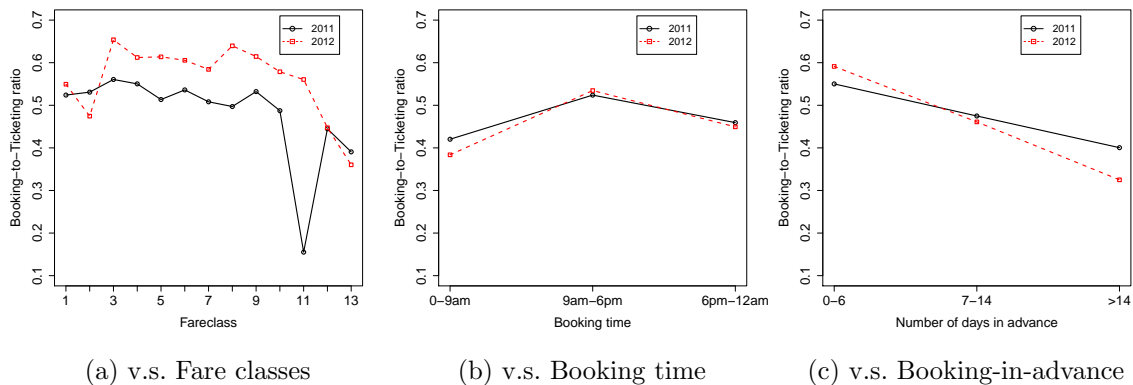


Figure 8 Booking-to-Ticketing percentage v.s. different factors

Figure 8b plots the booking-to-ticketing ratios for different booking time sections in year 2011 and 2012. The figure clearly shows that the customers who book the tickets during the day (from

9am to 6pm) are more likely to purchase the tickets than customer who book the ticketing during in the early morning or late evening. This is intuitively correct because customers booking the tickets during the day are probably business travelers and are more determined to purchase the ticket once they book it.

Figure 8c plots the booking-to-ticketing ratios versus number of days in advance of the booking. Clearly, customers who book the tickets earlier are less likely to purchase the tickets finally. This may be because there are more leisure customers than business customers to book tickets early and leisure customers change their ideas more often.

Other factors that are significant to booking-to-ticketing behavior include: month of the booking time, weekday of the booking time, number of days in advance of the booking time, and the booking channel. All these attributes are treated as nominal variables in the binary logistic regression.

The details of estimation model and its effectiveness is presented in the online companion of the paper.

6. Simulation

In this section, we discuss the results of simulations that approximate the objective of the stochastic optimization problem described in Section 3.1. For the study, we selected a typical Monday in 2012 and considered bookings for all the flights from A to B scheduled to depart on the selected Monday. The simulation used the following models:

1. For the market booking arrival times, the simulation used the actual customer booking times for the flights on the selected Monday.
2. For airline XX assortments, we used the following 7 booking policies:
 - (a) 2012 actual policy: The actual assortments offered by airline XX during the booking horizon for the flights on the selected Monday.
 - (b) 2011 MNL policy: the fluid booking policy obtained by calibrating all models with 2011 data and then solving the deterministic problem with a MNL model of booking choice behavior.
 - (c) 2011 NL policy: the fluid booking policy obtained by calibrating all models with 2011 data and then solving the deterministic problem with a NL model of booking choice behavior.
 - (d) 2011 ML policy: the fluid booking policy obtained by calibrating all models with 2011 data and then solving the deterministic problem with a ML model of booking choice behavior.
 - (e) 2012 MNL policy: the fluid booking policy obtained by calibrating all models with 2012 data and then solving the deterministic problem with a MNL model of booking choice behavior.
 - (f) 2012 NL policy: the fluid booking policy obtained by calibrating all models with 2012 data and then solving the deterministic problem with a NL model of booking choice behavior.

(g) 2012 ML policy: the fluid booking policy obtained by calibrating all models with 2012 data and then solving the deterministic problem with a ML model of booking choice behavior.

3. For airline YY and ZZ assortments, the simulation used the actual assortments offered by airlines YY and ZZ during the booking horizon for the flights on the selected Monday.

4. For booking choice behavior, the simulation used the three choice models described in Section 5.2:

- (a) MNL model calibrated with 2012 data (2012 MNL behavior),
- (b) NL model calibrated with 2012 data (2012 NL behavior),
- (c) ML model calibrated with 2012 data (2012 ML behavior).

5. For booking-to-ticketing conversion, the simulation used the model described in Section 5.4, calibrated with 2012 data.

6. For no-show and go-show behavior, the simulation uses the model described in Section 3 in the online companion, calibrated with 2012 data.

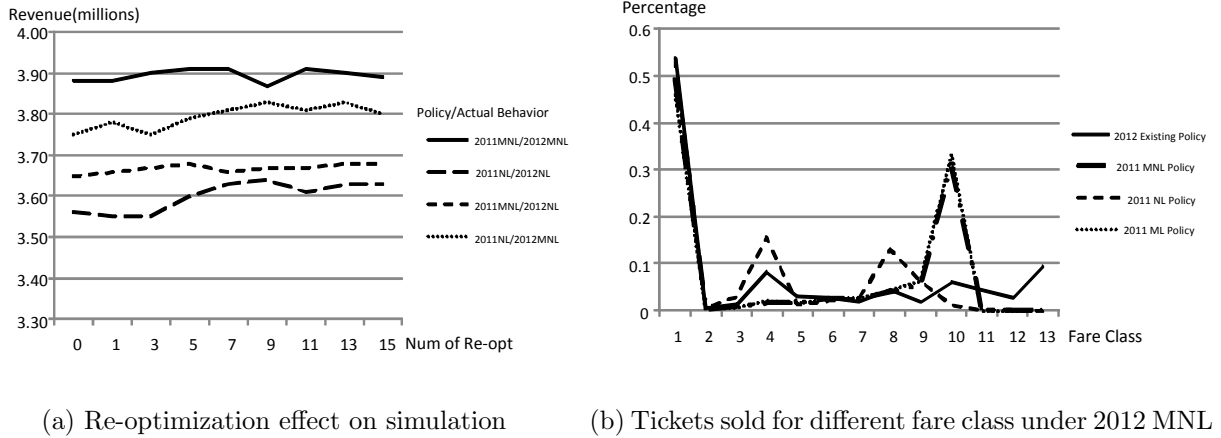


Figure 9 Simulation results

Figure 9a shows the effect of resolving the deterministic optimization problem a number of times during the booking horizon for various combinations of policy and booking choice behavior. Re-optimization seems to increase revenue with a small amount, but the incremental benefit seems insignificant after a few times. In the simulation runs, we re-optimized daily during the last week before departure for all choice models except for the mixed logit model.

Table 3 gives the simulated sample average revenues of airline XX when the airline either offers the assortments that it actually offered during the booking horizon for the flights of the selected Monday (2012 actual), or the airline uses one of six policies based on models calibrated with

Table 3 Simulated sample average revenues (in millions) of airline XX when the airline offers assortments using either its observed policy (2012 actual) or one of three policies based on solving a deterministic optimization problem using models calibrated with 2011 and 2012 data (corresponding to different rows in the table), and with customers making booking choices using one of three choice models (corresponding to different columns in the table).

	MNL	NL	ML		MNL	NL	ML
2012 actual	3.68	3.66	3.67	2012 actual	3.68	3.66	3.67
2011 MNL	3.90	3.67	3.88	2012 MNL	4.07	3.82	4.04
2011 NL	3.81	3.61	3.76	2012 NL	3.76	3.76	3.72
2011 ML	3.86	3.64	3.86	2012 ML	4.05	3.79	4.04

respectively 2011 or 2012 data, a MNL, NL, or ML model of booking choice behavior, and solving a deterministic optimization problem. The sample averages reported in both Table 3 were obtained with a sufficient number of replications such that the sample averages of revenue differences between the MNL, NL, or ML policies and the actual policy was greater than 3 times the standard deviations of the sample averages of these revenue differences. The only exception was the case of the 2011 MNL policy under 2012 NL booking choice behavior.

Note that under MNL or ML booking choice behavior, all other policies perform significantly better than the actual policy. However, under NL booking choice behavior, the 2011 NL and 2011 ML policies perform worse than the actual policy, the 2011 MNL policy performs about the same as the actual policy, and all the 2012 policies perform better than the actual policy. Thus all 2012 policies surpass the the actual policy under all booking choice models considered. Also, in all but two settings do the 2012 policies perform better than the corresponding 2011 policies. In addition, the MNL policy performs better than the ML policy which in turn performs better than the NL policy in all cases. That suggests that in order to obtain higher revenue, it helps to use more up-to-date models to generate the policy, even if the booking choice model is structurally incorrect and simpler than the actual choice behavior. We were surprised that the policies based on the simpler MNL model outperform the policies based on the more sophisticated models representing the more intuitively appealing customer choice behavior, even when the simulated behavior was based on the more sophisticated models.

Also, recall that most customers in this market seem to be business customers, and that most customers in this market book quite close to departure time when typically only expensive tickets are available (often only the highest fare class is available). Consequently, most tickets are sold at the maximum price. Given this situation, it is to be expected that the airline’s revenue in this particular market cannot be increased by much. Even when our policies perform better than the actual policy, we can surpass their revenue by only a little.

Figure 9b shows the distribution of fare classes sold by airline XX under MNL booking choice behavior. Each line represents a different policy used by airline XX in the simulation. We can see that more than 50% of tickets are sold at full price (fare class 1). Under our policy more tickets are sold for fare class 10 compared with the actual policy. Our policies rarely offer fare classes cheaper than class 10. Under the actual policy quite a number of tickets are sold for fare classes 11, 12 and 13. It suggests that airline XX can increase revenue by making the cheapest 3 fare classes available less often.

References

- Alstrup, J., S.-E. Andersson, S. Boas, O. B. G. Madsen, R. V. V. Vidal. 1989. Booking control increases profit at Scandinavian Airlines. *Interfaces* **19**(4) 10–19.
- Ben-Akiva, M., S. R. Lerman. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. The MIT Press, Cambridge, MA.
- Coldren, G. M., F. S. Koppelman. 2005. Modeling the competition among air-travel itinerary shares: GEV model development. *Transportation Research Part A: Policy and Practice* **39**(4) 345–365.
- Coldren, G. M., F. S. Koppelman, K. Kasturirangan, A. Mukherjee. 2003. Modeling aggregate air-travel itinerary shares: Logit model development at a major US airline. *Journal of Air Transport Management* **9**(6) 361–369.
- Davis, J., G. Gallego, H. Topaloglu. 2013a. Assortment optimization under variants of the nested logit model. Working paper, School of Operations Research and Information Engineering, Cornell University.
- Davis, J., G. Gallego, H. Topaloglu. 2013b. Assortment planning under the multinomial logit model with totally unimodular constraint structures. Working paper, School of Operations Research and Information Engineering, Cornell University.
- Gallego, G., H. Topaloglu. 2013. Constrained assortment optimization for the nested logit model. Working paper, School of Operations Research and Information Engineering, Cornell University.
- Greene, W. H., D. A. Hensher. 2003. A latent class model for discrete choice analysis: Contrasts with mixed logit. *Transportation Research Part B: Methodological* **37**(8) 681–698.
- Li, G., P. Rusmevichientong, H. Topaloglu. 2013. The d -level nested logit model: Assortment and price optimization problems. Working paper, School of Operations Research and Information Engineering, Cornell University.
- Littlewood, K. 1972. Forecasting and control of passenger bookings. *AGIFORS 12th Annual Symposium Proceedings*. Nathanya, Israel, 95–117.
- Rusmevichientong, P., D. Shmoys, C. Tong, H. Topaloglu. 2013. Assortment optimization under the multinomial logit model with random choice parameters. Working paper, School of Operations Research and Information Engineering, Cornell University.

-
- Rusmevichientong, P., H. Topaloglu. 2012. Robust assortment optimization in revenue management under the multinomial logit choice model. *Operations Research* **60**(4) 865–882.
- Slager, B., L. Kapteijns. 2004. Implementation of cargo revenue management at KLM. *Journal of Revenue and Pricing Management* **3**(1) 80–90.
- Smith, B. C., D. P. Günther, B. V. Rao, R. M. Ratliff. 2001. E-commerce and operations research in airline planning, marketing, and distribution. *Interfaces* **31**(2) 37–55.
- Smith, B. C., J. F. Leimkuhler, R. M. Darrow. 1992. Yield management at american airlines. *Interfaces* **22**(1) 8–31.
- Talluri, K. T., G. J. van Ryzin. 2004. Revenue management under a general discrete choice model of consumer behavior. *Management Science* **50**(1) 15–33.
- Train, K. E. 2003. *Discrete Choice Methods with Simulation*. Cambridge University Press, Cambridge, UK.
- Vulcano, G., G. Van Ryzin, W. Chahr. 2010. Choice-based revenue management: An empirical study of estimation and optimization. *Manufacturing & Service Operations Management* **12**(3) 371–392.
- Zhang, Dan, William L. Cooper. 2005. Revenue management for parallel flights with customer-choice behavior. *Operations Research* **53**(3) 415–431.