

A finite-population revenue management model and a risk-ratio procedure for the joint estimation of population size and parameters

Kalyan Talluri *

February 2, 2009

Abstract

Many dynamic revenue management models divide the sale period into a finite number of periods T and assume, invoking a fine-enough grid of time, that each period sees at most one booking request. These Poisson-type assumptions restrict the variability of the demand in the model, but researchers and practitioners were willing to overlook this for the benefit of tractability of the models.

In this paper, we criticize this model from another angle. Estimating the discrete finite-period model poses problems of indeterminacy and non-robustness: Arbitrarily fixing T leads to arbitrary control values and on the other hand estimating T from data adds an additional layer of indeterminacy. To counter this, we first propose an alternate finite-population model that avoids this problem of fixing T and allows a wider range of demand distributions, while retaining the useful marginal-value properties of the finite-period model.

The finite-population model still requires jointly estimating market size and the parameters of the customer purchase model without observing no-purchases. Estimation of market-size when no-purchases are unobservable has rarely been attempted in the marketing or revenue management literature. Indeed, we point out that it is akin to the classical statistical problem of estimating the parameters of a binomial distribution with unknown population size and success probability, and hence likely to be challenging. However, when the purchase probabilities are given by a functional form such as a multinomial-logit model, we propose an estimation heuristic that exploits the specification of the functional form, the variety of the offer sets in a typical RM setting, and qualitative knowledge of arrival rates. Finally we perform simulations to show that the estimator is very promising in obtaining unbiased estimates of population size and the model parameters.

Key words. revenue management, estimation, multi-nomial logit, risk-ratio.

1 Introduction

Many published papers in the area of revenue management model time in discrete intervals with a finite number of periods and assume that there is at most one customer arrival per period. Some

*Kalyan Talluri, ICREA and UPF, Universitat Pompeu Fabra, Ramon Trias Fargas 25-27, 08005 Barcelona, Spain, email: kalyan.talluri@upf.edu

prominent examples are Lee and Hersh [21] and Talluri and van Ryzin [24, 25]¹. The rationale is that by taking the number of periods large enough we can assume that the physical time period is sufficiently small so the possibility of more than one arrival in that period can be discarded. These assumptions have the same flavor as the ones behind the Poisson distribution. Moreover, they make the resulting finite-period dynamic program amenable to analysis allowing us, say, to obtain monotone marginal value properties.

This discrete finite-period model however carries over some of the limitations of the Poisson demand assumption—namely, the demand distribution has a co-efficient of variation significantly lower than what is seen in data (1 compared to a c.v of 2 or 3 in most RM industries). This criticism of the model is well-known and practitioners either use non-dynamic-programming heuristics or adjust their dynamic program results somehow to account for it.

To this existing criticism we add another significant one from the point of view of estimating the model: the model says the number of periods T is fixed and large enough but gives no guidance on what it is. Different values of T give different control values which leads to an awkward gap between theory and implementation. Estimating T directly from data is also troublesome—it adds another source of indeterminacy to the estimation as we discuss shortly.

Rather than continuing developing theory using this finite-period model, we propose in this paper that we change the paradigm to a finite-population model, that effectively avoids choosing the number of periods. It also has the advantage of allowing arbitrary distributions (and arbitrary variance) in the aggregate demand. Estimating such a model still involves some indeterminacy and we propose a new method of resolving this using a least-squares criteria. Finally we show that some of the standard monotone properties of the dynamic program carry over to this model.

Our main contributions in this paper are:

- To propose a more realistic and flexible finite-population model for revenue management and show that many of the appealing properties of the finite-period model carry over to this finite-population model.
- Propose a new estimation heuristic, the log risk-ratio estimate, that, once we fix a functional form on the purchase probabilities such as the multinomial-logit model, and under reasonable assumptions on the data, does a remarkable job of jointly estimating the market-size and purchase probabilities.

The finite-population model facilitates estimation as follows: we can consider the population or market size $M = \lambda T$ where λ is the arrival rate in the finite period model. By posing the model in terms of M , we reduce it to estimating a single parameter M rather than the two parameters T and λ . Nevertheless, when no-purchases are unobservable, there is still a significant problem of jointly estimating the market-size and purchase probabilities, and we consider the risk-ratio based estimation heuristic that we propose as the main contribution of the paper. This problem comes up in many areas of marketing, supply-chain management and revenue management, and is of great practical importance. Biased or unstable estimates can lead to significantly less revenues, and as we point out in Section 3.1 evidence from the literature indicates that both bias and stability are likely to be serious problems.

Researchers have widely used the MNL model to model purchase probabilities, but by and large

¹There are many examples in other areas of Operations Management as well, but for concreteness, we take revenue management as our example area.

have avoided considering this joint-estimation problem by assuming there is sufficient data on total population ([11]) or some exogenous knowledge of market size ([19], [27]).

Our proposed solution does not cover all situations, and indeed, given the relationship with the binomial estimation problem with unknown population and success probability, it would be too much to expect a complete solution. We now enumerate the limitations of our procedure and the situations where it is applicable.

1. It is designed for situations where the customer’s consideration set (and the firm’s offer set) consists of at least two products, or situations where the consideration set is a single product but the firm knows a priori how to rank the products. An example of the latter is restriction-free revenue management where price is the only differentiating factor between the products and the co-efficient of the price variable is negative—an entirely reasonable proposition.
2. We make the assumption that demand arrives over time and the firm has some idea of which periods have similar demand rates. This can potentially restrict applying our estimation procedure, but there are many situations where it is reasonable to expect that the firm has some knowledge of arrival rates. Airlines for instance divide the booking period to data-collection-points or DCPs such that the arrival rates within each DCP are similar across different instances of the flight. Note that this is a qualitative requirement—the analyst or manager only needs to know that two periods have similar arrival rates, but need not have any knowledge of what the arrival rate exactly is. This is much easier to implement than requiring the analyst guess the market size or potential population.
3. We assume a single segment. In many RM situations it is more reasonable to assume that the population is made up of two or three distinct segments with different purchase behavior.

The proposed estimation heuristic performs very well in simulations within the confines of its stated assumptions. We believe further research extending the method to multi-segment populations holds the promise of a significant improvement in the estimation of RM models. In the last section we outline some possible directions such research might take.

2 Models

For concreteness we consider the finite-period multinomial-logit (MNL) customer choice model of revenue management ([25]). The MNL model is parsimonious, widely used in practice, and serves to illustrate the basic ideas in our estimation procedure with a concrete model.

2.1 The Finite-Period Model

Time is discrete and indexed by t , and the indices run backwards in time (e.g. smaller values of t represent later points in time). That is, time 0 represents the deadline for the sale of units. There are a total of T time periods.

In each period there is at most one arrival (the time grid is sufficiently fine, or alternately T is sufficiently large, to justify this assumption). The probability of arrival is denoted by λ , which we assume is the same for all time periods t .

There are n fare products and $N = \{1, \dots, n\}$ denotes the entire set of fare products. Each fare product $j \in N$ has an associated revenue r_j , and without loss of generality we index fare products so that $r_1 \geq r_2 \geq \dots r_n \geq 0$.

In each period t , the firm must choose a subset $S \subseteq N$ of fare products to offer. When the fares S are offered, the probability that a customer chooses class $j \in S$ is denoted $P_j(S)$ and we assume $P_j(S) = 0$ if $j \notin S$. We let $j = 0$ denote the no-purchase choice; that is, the event that the customer does not purchase any of the fares offered in S , so $P_0(S)$ denotes the no-purchase probability.

The probability that a sale of class j is made in period t is therefore $\lambda P_j(S)$, and the probability that no sale is made is $\lambda P_0(S) + (1 - \lambda)$. This last expression reflects the fact that having no sales in a period could be due either to no arrival at all or an arrival that does not purchase; as mentioned, this leads to an incomplete data problem in practice.

2.1.1 MNL model

$P_j(S)$ is defined for all possible subsets S and $j \in S$. Such generality clearly poses an estimation problem—to estimate these probabilities we would need an unreasonable amount of data. To keep the parameters small, we fix on a functional form for the choice probabilities.

The MNL is used widely in travel demand forecasting and marketing (see Ben-Akiva and Lerman [2].) In the MNL, consumers are utility maximizers and the utility of each choice is a random variable. Formally, the utility of each alternative j is assumed to be of the form

$$U_j = u_j + \xi_j$$

where u_j is the mean utility of choice j and ξ_j is an i.i.d., Gumbel random noise term with mean zero and scale parameter one for all j . (The Gumbel, or double-exponential, distribution with zero mean and scale parameter one is $F(x) = \exp(-e^{-(x+\gamma)})$, $x \geq 0$ where $\gamma \approx 0.577$ is Euler's constant.) Because utility is an ordinal measure, the assumption of zero mean and a scale parameter of one are without loss of generality (see Ben-Akiva and Lerman [2]). Similarly, there is a no-purchase option where the no-purchase utility is assumed to be

$$U_0 = u_0 + \xi_0,$$

where ξ_0 is also Gumbel with mean zero and scale parameter one. Again, since utility is ordinal, without loss of generality we can assume $u_0 = 0$.

Under this utility model, one can show (See Ben-Akiva and Lerman [2] for a derivation.) that the choice probabilities are given by

$$P_j(S) = \frac{e^{u_j}}{\sum_{i \in S} e^{u_i} + e^{u_0}}, \quad j \in S \text{ or } j = 0. \quad (1)$$

and zero otherwise. For notational convenience, we define “weights” $w_j = e^{u_j}$, $j = 0, 1, \dots, n$, so that the choice probabilities can be expressed as

$$P_j(S) = \frac{w_j}{\sum_{i \in S} w_i + 1}, \quad j \in S \text{ or } j = 0. \quad (2)$$

That is, the probability of choosing j is its *weight* w_j divided by the weights of all other choices in S . So the no-purchase weight is w_0 that we can normalize to 1.

2.1.2 Discrete-choice finite-period dynamic program

Let C denote the initial resource capacity, t the number of remaining periods (recall time is indexed backwards) and x the number of remaining inventory units. Define the value function $V_t(x)$ as the maximum expected revenue obtainable from periods $t, t-1, \dots, 1$ given that there are x inventory units remaining at time t . Then the Bellman equation for $V_t(x)$ is

$$\begin{aligned} V_t(x) &= \max_{S \subseteq N} \left\{ \sum_{j \in S} \lambda P_j(S) (r_j + V_{t-1}(x-1)) + (\lambda P_0(S) + 1 - \lambda) V_{t-1}(x) \right\} \\ &= \max_{S \subseteq N} \left\{ \sum_{j \in S} \lambda P_j(S) (r_j - \Delta V_{t-1}(x)) \right\} + V_{t-1}(x), \end{aligned} \quad (3)$$

where $\Delta V_{t-1}(x) = V_{t-1}(x) - V_{t-1}(x-1)$ denotes the marginal cost of capacity, and we have used the fact that for all S ,

$$\sum_{j \in S} P_j(S) + P_0(S) = 1.$$

The boundary conditions are

$$V_t(0) = 0, \quad t = 1, \dots, T \quad \text{and} \quad V_0(x) = 0, \quad x = 1, \dots, C. \quad (4)$$

2.1.3 Unobservability, Estimation and Indeterminacy

The parameters of the multi-nomial logit model are fairly straightforward to estimate if one can observe all transactions, including no-purchases. (See Section 5.5 of Ben-Akiva and Lerman [2]). However, the main complication arises when we cannot observe no-purchases, which is the case in most revenue management sale settings.

Talluri and van Ryzin [25] give a procedure to simultaneously estimate λ and the parameters of the MNL, the w 's, using the EM algorithm to correct for unobserved no-purchases. T is assumed fixed. Kök and Fisher [19], in the context of assortment planning, generalize this algorithm when there is a substitution matrix. Recently Vulcano, van Ryzin and Ratliff [27] propose an alternate method with Poisson aggregate demand between state changes (change in the subset of offerings). To overcome the complexity of the likelihood function, [27] propose estimating based on unobservable primary demand (essentially, an estimate of what the customer would have chosen if all products were offered).

All these procedures suffer a fundamental indeterminacy problem which can be described roughly as follows: if there are M customers who arrive to purchase the product but only m purchase it (the others say found it too expensive and walked away) then our model is $m = Mq$, where q is the purchase probability. We observe m but neither M nor q , and we are trying to estimate both M and q from data on m . Clearly, many different solutions are possible for a given m , such as a high M and a low q or a low M and a high q . There is really no way to resolve this observing only m . Unfortunately, different values of M and q would give different control values in our subsequent optimization problem, so this is not just a theoretical curiosity.

Talluri and van Ryzin [25] note this indeterminacy but do not propose a way to resolve it. Vulcano, van Ryzin and Ratliff [27] show that there can be an infinite number of parameterized solutions that can explain the data and propose using an exogenous parameter, representing potential

market, to resolve this extra degree of freedom, but the question of choosing this parameter becomes another issue.

2.1.4 Criticism of the finite-period model

Before one can implement the finite-period model, one needs to know T , λ , and the parameters of the choice model, in this case the w 's (or the utilities u 's as a function of the product characteristics such as price).

That the dynamic program gives different control values for different values of T should be obvious. For instance, say if capacity is one and there is just one period left, then the marginal value (bid-price) would be zero whereas if we half λ and make it a two period model, the bid-price would go positive.

The relationship between choice of T and estimation is a bit more subtle. At first glance, as we are working with the same T in the estimation as well as the optimization part, it might seem that one can just fix a large enough T , and changing T would not have an effect. In one sense this is true: the convergence criteria of the EM algorithm (see for instance Equation 9.61 on page 480 of [26]) show that, say, if we double T , halving λ and keeping the same MNL parameters satisfy the convergence and optimal solution criteria of the EM algorithm, as it should be. But the problem is that, with unobservable no-purchases, the log-likelihood function is not globally concave (which is another problem in practice) so the choice of T influences the local maximum the EM algorithm converges to. So one could possibly get different solutions for different choices of T .

One could of course make T endogenous to the model and let the estimation algorithm determine T , but this only introduces further, unnecessary, indeterminacy into the model. Going back to our illustration of indeterminacy in Section 2.1.3, the observed purchases m would then satisfy $m = \lambda T q$, so now we really have to estimate λ , T and q observing only m .

2.2 The Finite-Population Model

One way of getting around the problem of choosing T is to switch to a finite-population model. As an added benefit we also get more flexibility in modeling the demand distribution.

2.2.1 The Model

Let M be a random variable representing the total, potential, market size. M can be considered the forecast at the current point of time (the point of optimizing the controls) over the remaining sale period. Note that one can specify any distribution for M , allowing us to fit the variance in the observed data.

One of the hallmarks of revenue management industries is that demand arrives over the sale period. This has led to time-based models such as Poisson, and the concept of a booking-curve. The resulting dynamic program indexes over time (some earlier papers, Curry [12], Wollmer [28], Brumelle and McGill [9], index over classes assuming independent classes; Lee and Hersh [21] appears to be the first paper with a time-based model). Here we forgo doing this—the order of arrivals is immaterial to our optimization problem, just the total number of customers yet to come. The model can potentially be extended to multiple segments and incorporate knowledge of how they arrive

over time with the advantage that such an extension uses fewer time periods explicitly linked to our knowledge of the demand.

To summarize, we model the demand process for the usage of the resource for a specific day in the future is as follows:

1. A population size that is a draw from a known distribution M with unknown parameters.
2. The individual customers in the potential market arrive over the booking horizon, distributed according to some distribution. For simplicity of exposition, assume they are distributed uniformly from 0 to T .
3. Each customer considers the prospect of buying exactly one item of a subset of the n products. If a subset S is offered, the customer buys $j \in S$ with probability $P_j(S)$, and does not buy any product with a probability $P_0(S) = 1 - \sum_{j \in S} P_j(S)$. If he does not purchase, he disappears from the system. For simplicity we assume that the purchase probabilities follow a multinomial-logit model, but conceivably more complicated models of purchase can be substituted in.

2.2.2 The Dynamic Program

We make no distributional assumptions on M except it is bounded above by a number B . Given a distribution for M , we can define distributions for the random variables $M - m$ where m is any positive integer, $0 \leq m \leq B$. Let $V(x, M - m)$ be the value function given current capacity is x and the distribution of the future market size is given by the random variable $M - m$. $V(x, M - m | M \geq k)$ is the value function when the market size distribution is given by $M - m$ conditional on $M \geq k$.

The Bellman equation is given by

$$V(x, M) = \max_S P(M \geq 1) \left\{ \sum_{j \in S} P_j(S) (r_j + V(x-1, M-1 | M \geq 1)) + (1 - \sum_{j \in S} P_j(S)) V(x, M-1 | M \geq 1) \right\}, \quad (5)$$

with the boundary condition $V(0|\cdot) = V(\cdot|M - B) = 0$.

B in Equation 5 is typically not too large. For instance, for an airline flight with 200 seats the probability of a demand of 1000 would be so low that we would just bound the distribution at 1000.

The fact that we could eliminate time from Equation 5 is mainly because we are assuming just one type of customer. If there are more segments, and we know how their arrivals are correlated with time, we would have to bring back time as a state variable. However, we could take a very coarse breakup of time without affecting the value function too much. For instance in the independent class models of Curry [12], Wollmer [28], Brumelle and McGill [9], the number of time intervals correspond to the number of classes. If we know that the low-sensitivity segment arrives closer to the usage of the resource, breaking up the booking horizon into two or three periods might be sufficient. Notice arrival rates can differ for different periods but as long as the segment mix does not change we do not need more time periods. For simplicity of exposition we assume from now on that there is only one segment.

The difference between the recursions (5) and (3) is that the latter assumes we can optimize after every arrival *and* at any point in time, whereas the former optimizes only after every arrival. So, the finite-population model is restrictive in some sense. On the other hand, the finite-period model fixes the evolution of the demand over time (it has to, to figure out the optimization as a function

of time) and this leads to restrictions on the variance of the demand. The finite-population model can easily be extended to incorporate re-optimization and re-forecasting points in the future; such extensions are more appealing in our view because we can have a coarse grid where the choice of the grid is linked to physical time points where we actually plan to do such re-optimizations, rather than a vague parameter as in the finite-period model.

Finally we would like to mention that (5) is not entirely new—the implicit dynamic program considered by Dudey [14] is the duopoly version of (5) when M is deterministic. Also, the estimation procedure of Section 3 can be applied equally well to the finite-period model with a fixed T .

While Equation 5 does not make use of time, we need to model demand as arriving over time when estimating the parameters of the model, as explained further in Section 3.

Define the marginal value function $\Delta V(x|M) = V(x|M) - V(x-1|M)$. Akin to the finite-period model, the marginal values exhibit nice properties.

Proposition 1 $\Delta V(x|M) \leq \Delta V(x-1|M)$ and $\Delta V(x|M-1) \leq \Delta V(x|M)$.

Proof

See Appendix.

Q.E.D

Proposition 1 justifies the use of a threshold value control such as the bid-price control.

3 Estimation

In this section we informally define the estimation procedure using examples. First, we review existing literature for the estimation of models similar to ours.

3.1 Literature

The classical problem of estimation of the parameters of a binomial distribution $Bin(M, p)$ from K samples (representing counts of successes), X_1, X_2, \dots, X_K , with M and p fixed but unknown is a good place to start. It was considered by Fisher [16] and Haldane [17]. In theory the sample maximum $X_{(K)} \rightarrow M$, as observed by Fisher, and hence the problem is “solved”, but the number of samples required for the sample maximum to approach M makes this fact irrelevant for practice.

Our problem has a similar flavor. Consider the model as described in Section 2.2.1. If we fix the population size and fix an offer set, the number of purchases is a binomial random variable with a fixed purchase probability. We can only observe purchases, so this is akin to the binomial problem with M and p fixed but unknown. In our model, the population size is considered a random draw from a distribution and the offer sets change frequently, so neither M nor p is fixed. So, as it were, the binomial problem is a special case as far as M is concerned, but also a less restrictive one with regards to p , as it is allowed to be any number, while we impose a functional form on the values it can take.

Time plays a crucial role in our model. The population is supposed to be uniformly distributed

over the booking horizon², and we can exploit this, along with the fact that the purchase probabilities follow a functional form, to estimate population size.

Coming back to the literature on the binomial estimation problem, the problems of estimating M or p when both are unknown is considered notoriously difficult. Both maximum-likelihood (MLE) and method-of-moments (MME) estimators are found to be highly unstable, i.e., even the addition of a single success changes the estimate \hat{M} by a tremendous amount; moreover underestimation of M is pervasive to all the methods. So stabilized estimates are given in Olkin et. al. [23] and Carroll and Lombard [10]. DasGupta and Rubin [13] show theoretically that no function of just M or just p can produce unbiased estimates, and give two new estimates of M that compare favorably with the Carroll-Lombard estimator; Blumenthal and Dahiya [6] study small-sample and asymptotic properties. Some of these approaches are Bayesian in nature ([10]) requiring a prior on the form of p , that may be hard to apply in our context. The problem has further connections with the reliability literature ([1]), and similar estimation problems have been studied for other distributions ([7]).

Given the difficulty of the fixed M, p estimation problem, obtaining statistically good estimates for our problem, or even heuristic methods that perform well, seems like a daunting task. The Operations Management and Marketing literature where similar estimation problems crop up mostly sidestep the issue by exogenously fixing a market size. For instance in revenue management [27] assume that the firm knows the potential “market share”, akin to fixing M and estimating p in the binomial model. In the marketing literature, some examples of fixing *potential market* exogenously are [22] and [3].

Fixing a potential market makes sense if it is a stable market and we have at least aggregate market level data. Neither applies to our operational revenue management problem. The potential market changes depending on which day the resource is used—the variability across days of the observed demand ought to convince us of this. Nor do we have market-share data at the resource level. So, it is very important that estimates of the market size are endogenously determined.

Considering the evidence from literature, the performance of a joint estimation procedure—efficiency, consistency, unbiasedness—might be too hard to prove theoretically. But simulations can provide reasonable validation as the models are easy to simulate.

3.2 Risk-ratio

The statistical basis for the method is the following set of results ([15],[18], [20]) on the comparison of the ratio of two binomial proportions: Let $X_1 \sim Bin(M_1, p_1)$ and $X_2 \sim Bin(M_2, p_2)$ be two independent binomial random variables and let $R = p_1/p_2$, called the *risk-ratio*. If x_1 and x_2 are the realizations of X_1 and X_2 then $\hat{p}_1 = x_1/M_1$ and $\hat{p}_2 = x_2/M_2$ are estimates of p_1 and p_2 respectively, and $\hat{R} = \hat{p}_1/\hat{p}_2$ an estimate of R .

- The random variable $\hat{p}_1 - R\hat{p}_2$ is approximately normally distributed with approximate zero mean and approximate variance $\frac{\hat{p}_1(1-\hat{p}_1)}{M_1} + \frac{\hat{p}_2(1-\hat{p}_2)}{M_2}$.
- The random variable $\ln(\hat{R})$ is approximately normally distributed with approximate mean $\ln(R)$ and approximate variance $\frac{(1-\hat{p}_1)}{x_1} + \frac{(1-\hat{p}_2)}{x_2}$.

²In practice, we can relax the uniform distribution assumption, but just assume that we can divide the booking horizon *a priori* into periods where the population arrives uniformly within the period.

One could use either fact to build an estimation heuristic, but for computational reasons, to be explained shortly, we fix the logarithmic risk-ratio in our subsequent discussion. So $\ln(\hat{R})$ converges to $\log R$.

3.2.1 Logarithmic risk-ratio estimate

We describe the two situations where we can use our estimation procedure.

3.2.2 Case 1: customers consider multiple products

First consider the case where the customer’s consideration set consists of at least two products of the firm, and the firm offers at least two products, or at least there are enough periods where the firm had offered at least two products.

We first estimate the weights in the MNL model w_1, \dots, w_n corresponding to the n products, conditioned on the purchases. We can do this based only on observed purchases. The w ’s are functions of the product characteristics (typically of the form $w = e^{u_j}$, where $u_j = \beta_0 + \beta_1 y_{1j} + \dots$). The maximum-likelihood estimation of these weights, conditioned on purchase, is quite standard and we refer the reader to Ben-Akiva and Lerman [2], page 118.

We note that the maximum-likelihood function is globally concave (under the form of utilities mentioned above) and the estimators have nice properties—consistent, asymptotically normal, asymptotically efficient etc. So from now on we fix $\hat{w}_1, \dots, \hat{w}_n$.

The difficult part of the problem then is to estimate the no-purchase weight w_0 (with respect to the $\hat{w}_1, \dots, \hat{w}_n$) and the parameters of the distribution of M .

Assume that the potential demand arrives uniformly over time. (The adjustments required when this is not the case are rather straightforward, and the assumption simplifies the presentation.) Note that we do not require knowledge of the rate itself (which would make the whole problem trivial), but a breakdown of the booking horizon into periods with more or less uniform rates. In the airline industry it is common to break up the booking period into data collection points (DCPs) which roughly corresponds to this idea.

Suppose in a sample k we had offered a set S_1 for a certain period of time and set S_2 the rest. The expected number of purchases under each offer set is proportional to the market size \tilde{M}_k (the realization of the market size for instance k). So by taking a ratio of the number of purchases, we eliminate the unobserved \tilde{M}_k . Based on these ratios we estimate w_0 in the first step, and based on the estimate of w_0 , we form a sample specific estimate of \tilde{M}_k and thus estimate the parameters of the distribution of M . This procedure works if we have sufficiently rich data, with enough variation in the offer sets, as well as some knowledge of the arrival rates.

Consider an example with three products where demand comes uniformly over time, and for a proportion ρ of the booking period products 1 and 2 were offered, and for the remaining fraction $(1 - \rho)$, products 2 and 3 were offered. Let m_{ij} be the total number of sales that have been observed for product i during the period j . Let \tilde{M} represent the unobserved total number of customers who had arrived (i.e., the unknown realization of the random variable M).

We first make a modeling approximation and assume that the market population \tilde{M} is distributed deterministically over the booking period. So in the above example with three products, we assume

that the number of potential customers who show up in the first period when products 1 and 2 are offered are *exactly* $\rho\tilde{M}$, rather than being a (higher-level) draw from a binomial distribution with parameters ρ and \tilde{M} . This assumption allows us to treat the number of purchases in the two periods as drawn from two *independent* binomial distributions with population sizes $\rho\tilde{M}$ and $(1 - \rho)\tilde{M}$ (we fudge over integrality) that we now compare using the risk-ratio approximation.

Since we know ρ , the risk-ratio cancels out the unobservable \tilde{M} in the sample estimate $\ln(\hat{R})$ allowing us to calculate it without observing \tilde{M} , and at the same time, $\ln(R)$ conserves its dependence on w_0 if we have rich enough data. We illustrate this using our example.

As we mentioned earlier, we have the (correct) estimates of the weights w_1, w_2, w_3 , namely \hat{w}_1, \hat{w}_2 and \hat{w}_3 . From our model and assumption, we know that the first period has $\rho\tilde{M}$ arrivals, and the second period $(1 - \rho)\tilde{M}$ arrivals. Let p_{ij} be the probability of purchasing product i in period j given our offer sets. These are given by the multinomial logit model assumption, so for instance, $p_{11} = \frac{\hat{w}_1}{w_0 + \hat{w}_1 + \hat{w}_2}$ as $\{1, 2\}$ was the offer set in period 1.

The total probability of a purchase in period j be p_j , so $p_1 = p_{11} + p_{21}$ and $p_2 = p_{22} + p_{32}$, and the risk ratio R is a function of w_0 —just notice that because the offer sets are different in the two periods, the denominators in the ratio p_1/p_2 do not cancel out³—so we write it as $R(w_0)$ to emphasize this.

$$\hat{R} = \frac{\hat{p}_1}{\hat{p}_2} = \frac{(m_{11} + m_{21})/(\rho\tilde{M})}{(m_{22} + m_{32})/((1 - \rho)\tilde{M})} = \frac{(m_{11} + m_{21})/(\rho)}{(m_{22} + m_{32})/((1 - \rho))}. \quad (6)$$

So \hat{R} can be computed based only on observed purchases.

$\ln(\hat{R})$ is an observation from a normal distribution with mean $\ln(R(w_0))$. So we can extract one observation for each sample (in this two period example). So if there are K samples and n_k different offer sets in each sample, by taking ratios pairwise from these n_k periods in each sample we would have created a large number of observations of the risk-ratios from normal distributions.

By the law of large numbers, the sum of a sufficiently large sample of these observations should tend to the sum of means of these random variables, which, for most samples, would be a function of w_0 . We have complete freedom on which sub-sample we take. It is possible that the ratios cancel out, for instance taking pairs 1 and 2 and 2 and 1 would give us $\ln(\frac{p_1}{p_2}) + \ln(\frac{p_2}{p_1}) = 0$ serving no purpose. Our choice is guided by this and computational considerations: If we choose the ratios such that the numerator is always less than or equal to the denominator we are assured that the numerator and denominator do not cancel out, and also we minimize a convex objective function by the following fact.

Proposition 2 *Let the sum of the weights of the set of products be w_1 in period 1 and w_2 in period 2 and $w_1 < w_2$. Let w_0 be the no-purchase weight, and the purchase probability in period 1 be given by $p_1(w_0) = \frac{w_1}{w_1 + w_0}$ and in period 2 by $p_2(w_0) = \frac{w_2}{w_2 + w_0}$. Then $L = \ln(\frac{p_1(w_0)}{p_2(w_0)})$ is a convex function of w_0 .*

Proof

See Appendix.

Q.E.D

³There is a small possibility that $w_0 + \hat{w}_1 + \hat{w}_2 = w_0 + \hat{w}_2 + \hat{w}_3$ but by assuming the products are sufficiently different we ignore this from now on.

The reason for using the logarithm of the risk-ratio should be apparent now. We can take two periods with different lengths and base the statistic only on the risk-ratio of the purchase probabilities. The ratio of the lengths of the two periods separates out as a constant. For instance, taking the log of (6) gives the additive constant $\log \frac{(1-\rho)}{\rho}$.

Let \mathcal{F} represent a collection of feasible pairs. For sample k (an instance of the usage of the resource, such as all the arrivals for a single flight departure on a specific date), define $(i, j)_k$ if $(i, j) \in \mathcal{F}$ if the periods i and j are believed to have equal arrival rates. The duration of the two periods can differ. But as we pointed out, this just adds a constant, so it does not make a difference to the optimization problem that we define below (it changes the objective value, but not the convexity properties)—and to keep the notation simple we ignore this and assume the periods are of the same duration. Let \mathcal{F}_k be the collection of all feasible pairs $(i, j)_k$ in sample k .

Every pair in \mathcal{F} leads to a function of w_0 , $L_{ij}(w_0) = \ln\left(\frac{p_i(w_0)}{p_j(w_0)}\right)$ with an observed ratio $l_{ijk} = \ln\left(\frac{m_{ik}}{m_{jk}}\right)$ if $(i, j)_k \in \mathcal{F}_k$. Let n_{ij} be the number of times the pair i, j appear in the samples $n_{ij} = |\{(i, j)_k \in \mathcal{F}_k\}|$. We expect $L_{ij}(w_0)$ to be close to $\bar{l}_{ij} = \frac{\sum_{(i, j)_k \in \mathcal{F}_k} l_{ijk}}{n_{ij}}$.

Our estimate \hat{w}_0 minimizes the squared error:

$$\hat{w}_0 = \arg \min \left\{ \sum_{ij} (L_{ij}(w_0) - \bar{l}_{ij})^2 \right\}$$

which by Proposition 2 is a convex minimization problem.

All along we have ignored the possibility that one of the denominators can be zero. For instance, \hat{p}_1 or \hat{p}_2 could be 0 if there were no purchases in that period. We ignore such periods which also implies that we should rederive the statistics with binomial distributions conditional on at least one success—rather than getting into these messy technicalities, as a heuristic, we can still use the logarithmic risk-ratio method as is.

So the method completely avoids any EM-style iterations and subsequent problems of convergence. Moreover, as we pointed out earlier, evidence from the binomial distribution case suggests that the MLEs of the population size and probabilities are themselves likely to be poor predictors of the true values, so even if the EM algorithm finds the true MLEs they do not inspire much confidence.

Once we obtain $\hat{w}_0, \hat{w}_1, \dots, \hat{w}_n$, we use them to estimate the parameters of the market size distribution by first obtaining a point estimate of the market size for each instance. Given our model this is quite trivial. If m_S is the total number of observations during a period when the offer set was S , an estimate of the number of true arrivals \hat{M}_S would be

$$\hat{M}_S = m_S \left(\frac{\hat{w}_0 + \sum_{j \in S} \hat{w}_j}{\sum_{j \in S} \hat{w}_j} \right).$$

We replace the observed m_S for each instance of the data and forecast the total market size distribution based on the \hat{M}_S 's. The sum of the \hat{M}_S 's across all the state changes in an instance is the “market” demand for that sample. The parameters of the market demand distribution are then estimated. For instance, if we assume the demand comes from a normal distribution, we take the sample mean and sample standard deviation as our estimate of the mean and standard deviation respectively.

3.2.3 Case 2: customers consider a single product

If the customer's consideration set consists of a single product, or if the firm offers only one product at a time, then we cannot estimate the parameters of the MNL model based only on purchases. This situation can arise for instance in restriction-free revenue management (followed by many low-cost carriers) where a single price is quoted at any point in time.

We make the assumption that the firm can compare products and determine which product has a higher probability of purchase, without actually knowing the parameters of the model. In the restriction-free RM case, where price is the only attribute, it would amount to an innocuous assumption that customers are sensitive to price, so the higher price product has a lower probability of purchase. If the product has many other attributes this assumption might prove restrictive however.

To estimate the model, we fix the weight of the no-purchase option at 1. Let β be the vector of parameters for the products. Say product 1 is offered in period 1 and product 2 in period 2. The purchase probabilities are given by $p_1(\beta) = \frac{e^{\beta^\top x_1}}{e^{\beta^\top x_1} + 1}$ and $p_2(\beta) = \frac{e^{\beta^\top x_2}}{e^{\beta^\top x_2} + 1}$, and suppose we know that $p_1(\beta) \leq p_2(\beta)$.

The function $\log(\frac{p_1(\beta)}{p_2(\beta)})$ is in general not convex, but quasiconvex, based on the following fact:

Proposition 3 $\frac{p_1(\beta)}{p_2(\beta)}$ is quasiconvex in β when $p_1(\beta) \leq p_2(\beta)$.

Proof

See Appendix. Q.E.D

Sum of quasiconvex functions is not necessarily quasiconvex, so finding $\hat{\beta}$ by minimizing a squared error poses a computational problem.

Define \mathcal{F} , $(i, j)_k$ and \mathcal{F}_k exactly as in Case 1 but now assuming that in sample k , we can rank periods such that $p_i(\beta) \leq p_j(\beta)$ (without knowing β).

Every pair in \mathcal{F} leads to a function of β , $L_{ij}(\beta) = \ln(\frac{p_i(\beta)}{p_j(\beta)})$ with an observed ratio $l_{ijk} = \ln(\frac{m_{ik}}{m_{jk}})$ if $(i, j)_k \in \mathcal{F}_k$. Let n_{ij} be the number of times the pair i, j appear in the samples $n_{ij} = |\{(i, j)_k \in \mathcal{F}_k\}|$. As before, we expect $L_{ij}(\beta)$ to be close to $\bar{l}_{ij} = \frac{\sum_{(i, j)_k \in \mathcal{F}_k} l_{ijk}}{n_{ij}}$. The function $(L_{ij}(\beta) - \bar{l}_{ij})^2$ is quasiconvex, but $\sum_{ij} (L_{ij}(\beta) - \bar{l}_{ij})^2$ need not be; however, $\max_{ij} (L_{ij}(\beta) - \bar{l}_{ij})^2$ is quasiconvex.

Our estimate $\hat{\beta}$ minimizes the maximum error:

$$\hat{\beta} = \arg \min \{ \max_{ij} (L_{ij}(\beta) - \bar{l}_{ij})^2 \} \tag{7}$$

which involves minimizing a quasiconvex function.

4 Computational Results

The objective of this section is to test the effectiveness of the risk-ratio estimation procedure with MNL purchase probabilities. We use simulations to test how effectively the procedure recovers the

MNL parameters and unknown market size distribution.

We believe simulations, rather than estimating real-world data sets, are the right test for the method at this stage. First, we have no access to a RM data set with market size information (indeed, the definition of market itself is somewhat ambiguous) so it is unclear how one would verify the performance of the procedure. Second, even if one has a good idea of the true market-size it is not clear if the (unobservable) customer purchase behavior fits the MNL model (or some other functional form), and it is difficult to separate specification error from the estimation error.

So our controlled simulation experiments, while limited in scope, are designed to answer only the following question: If the customer purchase behavior follows the MNL model, how well does the risk-ratio method recover the true MNL parameters and the true market size distribution. We also keep the simulations very simple for two reasons: to illustrate the estimation procedure, and to make the results easily reproducible (using say Excel and a public and free software (BIOGEME)).

We perform two simulations corresponding to the situations described in Case 1 and Case 2. In both cases, the simulation methodology is simple. First, we generate a set of samples from a normal distribution representing market size. Each customer is randomly assigned to a booking period.

In the first experiment, there are three periods, each product has two attributes and customers consider both products if they are offered. The firm offers both products in the first period and one of the two products in the remaining two periods. Customers purchase according to the probabilities given by an MNL model with certain parameters. Only purchase information is available for the estimation procedure. We study how well the risk-ratio procedure recovers the MNL parameters and the moments of the market-size distribution.

For the second experiment corresponding to Case 2 we perform two simulations. In the first, there are two periods and each product has exactly one attribute (the price) and customers purchase only as a function of price. The firm offers a low-price product in the first period and raises the price in the second period. Since there is a single pair, the optimization with the min-max criterion coincides with the least-squares criterion. In the second simulation, we perform the estimation with three periods and prices and test the min-max criterion.

In both experiments customers are myopic—if they cannot purchase the product they disappear and find alternatives outside the firm’s purview. Also, as the primary assumption behind the risk-ratio estimate is that customers are distributed deterministically across the periods, we want to emphasize that in all the simulations below, customers are randomly assigned to each period. That is, first a customer population size is generated from the population distribution and each generated customer is randomly assigned to one of the periods with equal probability.

4.1 Multi-product consideration sets

The maximum likelihood function of the MNL distribution is not strictly concave and one can get multiple optimal solutions. To test the procedure we run the simulation on data where the variety in the offer sets is limited and the maximum-likelihood procedure gives estimates of the parameters which are quite different from the original—but with purchase probabilities nearly identical—and we see that the procedure still does a remarkable job of recovering the true purchase probabilities and the market size.

The market-size distribution is taken as Normal with mean μ of 125 and standard deviation σ of 20.

The products and their attributes, and the MNL co-efficients are given as in Table 1. The no-purchase weight is fixed at 0.25. The purchase probabilities depending on the offer set are given as

	Product 1	Product 2	True MNL Coefficient
attribute 1 (price)	70	200	$\beta_1 = -0.005$
attribute 2 (other)	30	5	$\beta_2 = -0.001$

Table 1: Two products, two attributes, their coefficients β under the MNL model.

in Table 2

	Product 1	Product 2
{1,2}	0.52608526	0.281593148
{1}	0.732294324	—
{2}	—	0.594185249

Table 2: Purchase probabilities under the MNL model with no-purchase weight fixed at 0.25.

A sample of the generated purchase data is given in Table 3 (in BIOGEME format).

ID	x11	x21	x12	x22	av1	av2	Choice
0	70.0	200.0	30.0	5.0	1	1	2
0	70.0	200.0	30.0	5.0	1	1	2
0	70.0	200.0	30.0	5.0	1	1	1
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots

Table 3: Choices offered, their attributes, and the choices made in period 1.

The MNL estimation is done using BIOGEME ([5],[4]), a freely available software for estimating extreme-value distributions. The MNL estimates conditioned on purchases use only the first period data. Based on 100 samples, and a total of 3411 customer purchases choosing one of the two alternatives during the first period, the summary statistics are given in Table 4.

Summary statistics	
Number of observations = 3411	
$\mathcal{L}(0)$	= -2364.325
$\mathcal{L}(\hat{\beta})$	= -2192.658
$-2[\mathcal{L}(0) - \mathcal{L}(\hat{\beta})]$	= 343.334
ρ^2	= 0.073
$\hat{\rho}^2$	= 0.072

Table 4: Summary statistics output of BIOGEME.

The co-efficient estimates are given in Table 5.

	True MNL Coefficient 2	Estimated MNL Coefficient
attribute 1 (price)	$\beta_1 = -0.005$	$\hat{\beta}_1 = -0.0269$
attribute 2 (other)	$\beta_2 = -0.001$	$\hat{\beta}_2 = -0.114$

Table 5: True and estimated co-efficients under the MNL model.

We use our risk-ratio procedure, using a simple search procedure along the single dimension. After scaling⁴ the product weights by 100 we obtain a no-purchase weight estimate \hat{w}_0 of 0.0018.

Although the parameters are significantly different (primarily because of our limited variety in the offer sets during the first period), the purchase probabilities are remarkably close, as seen in Table 6. Finally, based on these estimated probabilities, we estimate the market-size for each instance

	Product 1	Product 2
{1,2}	0.530418446	0.277734665
{1}	0.734381702	—
{2}	—	0.591451394

Table 6: Purchase probability estimates with $\hat{\beta}_1 = -0.0269, \hat{\beta}_2 = -0.114$ and $\hat{w}_0 = 0.0018$, to be compared with Table 2.

and an estimate of the parameters of the market-size distribution. Just taking the sample mean and standard-deviation of the market sizes, we get the following estimates: $\hat{\mu} = 123.94, \hat{\sigma} = 30.01$, reasonably close to the true values $\mu = 125, \sigma = 20$.

When the MNL estimates are close to the true estimates, say with a large variety of offer and consideration sets, the no-purchase weight estimate in turn turns out to be much closer in value to the true weight.

4.2 Single-product consideration sets

As we mentioned in Section 3.2.3 when customers consider a single product, the methodology used in the previous section—estimate the MNL parameters, and then the no-purchase weight—cannot be applied. So in this section we simulate a restriction-free revenue management situation where customers only consider the lowest-priced product. This situation also satisfies one of the conditions for quasiconvexity of the objective function that we minimize—it is safe to assume that the price-coefficient is negative so customers have a higher probability of purchasing a lower-priced product.

As in the previous section, we assume that the market-size distribution is Normal with mean μ of 125 and standard deviation σ of 20. We run two simulations for this case.

4.2.1 Two prices, two periods

There are exactly two periods and a single product is offered at a price of 50 in the first period and a price of 100 in the second period. The two periods are of equal length and customers arrive at the

⁴Because of the low values of the $\hat{\beta}$'s; the procedure is robust with respect to this scaling, and one gets good results multiplying by 10 or 1000 in this case.

same rate in each period.

Price is the only attribute influencing customer purchase and the price-sensitivity is captured by the negative co-efficient on the price attribute $\beta = -0.04$. We generate 100 observations of the market size, with customers within each instance distributed randomly over the two periods, and record the number of purchases at the price of 50 in the first period and 100 in the second period. The weight of the no-purchase option is fixed at 0.03 (unknown to the estimation).

The purchase probabilities with the above set of parameters are given as in Table 7.

Price	Purchase Probability
50	0.818550527
100	0.379083032

Table 7: True purchase probabilities for the restriction-free revenue management simulation with price coefficient $\beta = -0.04$ and no-purchase weight 0.03 in the MNL model.

In our estimation procedure we set the no-purchase weight as 1.0 and find a constant $\hat{\beta}_0$ and the co-efficient of the price attribute, $\hat{\beta}_1$ that minimizes Equation (7). This minimization can be done by any standard minimization procedure, and as the objective function is quasi-convex, we are assured of convergence to a global minima.

The estimates for this example turn out to be $\hat{\beta}_0 = 7.39$, $\hat{\beta}_1 = -0.076$, with purchase probabilities as given in Table 8, and population estimates $\hat{\mu} = 107$, $\hat{\sigma} = 18.46$.

Price	Purchase Probability
50	0.973
100	0.4476

Table 8: Estimated purchase probabilities for the restriction-free revenue management simulation to be compared with the true probabilities in Table 7.

4.2.2 Three prices, three periods

To test the min-max criterion, we need richer data with more feasible pairs. So we simulate a three period model with prices 50, 75 and 100 respectively in each period. As for the two period case, true co-efficient of price is $\beta = -0.04$ with a no-purchase weight of 0.03 (both unknown to the estimation), and the population is equally distributed amongst the three periods. The purchase probabilities with the above set of parameters are given as in Table 9.

The estimates for this example with the min-max criterion of Equation (7) using the three period pairs (1 vs 2 and 3; 2 vs 3) turn out to be $\hat{\beta}_0 = 3.8$, $\hat{\beta}_1 = -0.042$, with purchase probabilities as given in Table 10, and population estimates $\hat{\mu} = 124.21$, $\hat{\sigma} = 20.562$. The richer set of data leads to a striking recovery of the original parameters even with the min-max criterion.

Price	Purchase Probability
50	0.818550527
75	0.623999219
100	0.379083032

Table 9: True purchase probabilities for the restriction-free revenue management simulation with price coefficient $\beta = -0.04$ and no-purchase weight 0.03 in the MNL model.

Price	Purchase Probability
50	0.845534735
75	0.657010463
100	0.40131234

Table 10: Estimated purchase probabilities for the restriction-free revenue management simulation to be compared with the true probabilities in Table 9.

4.3 Extensions and further research

This paper presents a finite population model of RM and a log risk-ratio estimation heuristic that appears to be very effective in estimating the unobservable market size when customer’s purchase probabilities follow the MNL model. We also point out the considerable difficulty in this joint estimation of purchase probability parameters and market size by its connection to the binomial estimation problem.

While we consider the results very promising, considerable more research is needed to increase the robustness and scope of the method. First, the estimation idea can be applied to models other than MNL, perhaps at the risk of losing the convexity properties. Second, many revenue management applications exploit the presence of multiple segments in the population with distinct purchase characteristics and price sensitivities.

Extending the estimation procedure to a multi-segment population with unobservable mixture probabilities is perhaps the most important direction from a practice point of view. One idea is to combine the heuristic with an EM-style updating algorithm for estimating the mixture probabilities, as this is one of the standard applications of the EM algorithm.

5 Acknowledgments

Thanks to Garrett van Ryzin and Victor Martínez de Albeníz for helpful comments.

References

- [1] S. Basu and N. Ebrahimi. Bayesian capture-recapture methods for error detection and estimation of population size: heterogeneity and dependence. *Biometrika*, 88:269–279, 2001.
- [2] M. Ben-Akiva and S. Lerman. *Discrete-Choice Analysis: Theory and Application to Travel Demand*. MIT Press, Cambridge, MA, 1985.

- [3] D. Besanko, S. Gupta, and D. Jain. Logit demand estimation under competitive pricing behavior: An equilibrium framework. *Management Science*, 44:1533–1547, 1998.
- [4] M. Bierlaire. BIOGEME: A free package for the estimation of discrete choice models. In *Proceedings of the 3rd Swiss Transportation Research Conference*, Ascona, Switzerland, 2003.
- [5] M. Bierlaire. Estimation of discrete choice models with BIOGEME 1.7. biogeme.epfl.ch, August 2008.
- [6] S. Blumenthal and R. Dahiya. Estimating the binomial parameter n . *Journal of the American Statistical Association*, 76(376):903–909, 1981.
- [7] S. Blumenthal and R. Marcus. Estimating population size with exponential failure. *Journal of the American Statistical Association*, 70(352):913–922, 1975.
- [8] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, U.K, 2004.
- [9] S. L. Brumelle and J. I. McGill. Airline seat allocation with multiple nested fare classes. *Operations Research*, 41:127–137, 1993.
- [10] R. J. Carroll and F. Lombard. A note on n estimators for the binomial distribution. *Journal of the American Statistical Association*, 80(390):423–426, 1985.
- [11] P.K Chintagunta. A flexible aggregate logit demand model. Technical report, Graduate School of Business, University of Chicago, Chichago, IL, 2001.
- [12] R. E. Curry. Optimal airline seat allocation with fare classes nested by origins and destinations. *Transportation Science*, 24:193–204, 1990.
- [13] A. DasGupta and H. Rubin. Estimation of binomial parameters when both n , p are unknown. *Journal of Statistical Planning and Inference*, 130:391–404, 2005.
- [14] M. Dudey. Dynamic Edgeworth-Bertrand competition. *Quarterly Journal of Economics*, 107:1461–1477, November 1992.
- [15] E.C. Fieller. A fundamental formula in the statistics of biological assay, and some applications. *Quarterly Journal of Pharmacy and Pharmacology*, 17:117–123, 1944.
- [16] R. A. Fisher. The negative binomial distribution. *Annals of Eugenics*, 11:182–187, 1941.
- [17] J. B. S Haldane. The fitting of binomial distributions. *Annals of Eugenics*, 11:179–181, 1941.
- [18] D. Katz, J. Baptista, S. P. Azen, and M. C. Pike. Obtaining confidence intervals for the risk ratio in cohort studies. *Biometrics*, 34:469–474, 1978.
- [19] G. Kok and M. Fisher. Demand estimation and assortment optimization under substitution: methodology and application. *Operations Research*, 55:1001–10021, 2007.
- [20] P. A. R Koopman. Confidence intervals for the ratio of two binomial proportions. *Biometrics*, 40:513–517, 1984.
- [21] T. C. Lee and M. Hersh. A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Science*, 27:252–265, 1993.
- [22] A. Nevo. Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342, 2001.

- [23] I. Olkin, A. J. Petkau, and J. V. Zidek. A comparison of n estimators for the binomial distribution. *Journal of the American Statistical Association*, 76(392):637–642, 1981.
- [24] K. T. Talluri and G. J. van Ryzin. An analysis of bid-price controls for network revenue management. *Management Science*, 44:1577–1593, 1999.
- [25] K. T. Talluri and G. J. van Ryzin. Revenue management under a general discrete choice model of consumer behavior. *Management Science*, January 2004.
- [26] K. T. Talluri and G. J. van Ryzin. *The Theory and Practice of Revenue Management*. Kluwer, New York, NY, 2004.
- [27] G. Vulcano, G. van Ryzin, and R. Ratliff. Estimating primary demand for substitutable products from sales transaction data. Technical report, GSB, Columbia University, NY, NY, 2008.
- [28] R. D. Wollmer. An airline seat management model for a single leg route when lower fare classes book first. *Operations Research*, 40:26–37, 1992.

Proof of Proposition 1: The proof is essentially identical to that for the marginal values time-based discrete choice model of RM, so we keep it brief.

$$(i) \Delta_x V(x|M) \leq \Delta_x V(x-1|M)$$

Induction on m assuming $\Delta_x V(x-m|M) \leq \Delta_x V(x-m-1|M)$, $m \geq 1$. $V(x, M)$ can be rewritten as

$$V(x, M) = P(M \geq 1)V(x, M-1|M \geq 1) + \max_S P(M \geq 1) \left\{ \sum_{j \in S} P_j(S)(r_j - \Delta V(x, M-1|M \geq 1)) \right\} \quad (8)$$

Let $S^*(x)$ be the optimal offer set for the first customer in the calculation of $V(x, M)$.

$$\Delta V(x, M) = V(x, M) - V(x-1, M) \quad (9)$$

$$\begin{aligned} &\leq P(M \geq 1)[\Delta V(x, M-1|M \geq 1) + \sum_{j \in S^*(x-1)} P_j(S^*(x-1))(-\Delta V(x, M-1|M \geq 1) \\ &\quad + \Delta V(x-1, M-1|M \geq 1))] \end{aligned} \quad (10)$$

and

$$\Delta V(x-1, M) = V(x-1, M) - V(x-2, M) \quad (11)$$

$$\begin{aligned} &\geq P(M \geq 1)[\Delta V(x-1, M-1|M \geq 1) \\ &\quad + \sum_{j \in S^*(x-1)} P_j(S^*(x-1))(-\Delta V(x-1, M-1|M \geq 1) \\ &\quad + \Delta V(x-2, M-1|M \geq 1))] \end{aligned} \quad (12)$$

so it suffices to show the quantity (10) is less than or equal to the quantity (12). By induction, $\Delta V(x, M-1|M \geq 1) - \Delta V(x-1, M-1|M \geq 1) \leq 0$, using which,

$$\begin{aligned} \Delta V(x, M) - \Delta V(x-1, M) &= P(M \geq 1)[(\Delta V(x, M-1|M \geq 1) \\ &\quad - \Delta V(x-1, M-1|M \geq 1))] \end{aligned} \quad (13)$$

$$\begin{aligned}
& + \sum_{j \in S^*(x-1)} P_j(S^*(x-1))(-\Delta V(x, M-1|M \geq 1) \\
& + \Delta V(x-1, M-1|M \geq 1) + \Delta V(x-1, M-1|M \geq 1) \\
& - \Delta V(x-2, M-1|M \geq 1))] \\
= & P(M \geq 1)[(1 - \sum_{j \in S^*(x-1)} P_j(S^*(x-1)))(\Delta V(x, M-1|M \geq 1) \\
& - \Delta V(x-1, M-1|M \geq 1)) \\
& + \sum_{j \in S^*(x-1)} P_j(S^*(x-1))(\Delta V(x-1, M-1|M \geq 1) \\
& - \Delta V(x-2, M-1|M \geq 1))] \\
\leq & 0
\end{aligned}$$

(ii) $\Delta_m V(x|M-1) \leq \Delta_m V(x|M)$

Consider

$$\begin{aligned}
\Delta V(x, M) & = V(x, M) - V(x-1, M) & (14) \\
& \geq P(M \geq 1)[\Delta V(x, M-1|M \geq 1) \\
& \quad + \sum_{j \in S^*(x)} P_j(S^*(x-1))(-\Delta V(x, M-1|M \geq 1) \\
& \quad + \Delta V(x-1, M-1|M \geq 1))] \\
& = \Delta V(x, M-1) + P(M \geq 1) \sum_{j \in S^*(x-1)} P_j(S^*(x-1))(\Delta V(x-1, M-1|M \geq 1) \\
& \quad - \Delta V(x, M-1|M \geq 1)) \\
& \geq \Delta V(x, M-1)
\end{aligned}$$

as $\Delta V(x-1, M-1|M \geq 1) - \Delta V(x, M-1|M \geq 1) \geq 0$ from the proof of (i).

Proof of Proposition 2: Let the sum of the weights of the set of products be w_1 in period 1 and w_2 in period 2 and $w_1 < w_2$. w_0 is the no-purchase weight. We show $L = \ln(\frac{p_1(w_0)}{p_2(w_0)}) = \ln(\frac{w_1(w_2+w_0)}{w_2(w_1+w_0)}) = \ln(\frac{w_1}{w_2}) + \ln(\frac{w_2+w_0}{w_1+w_0})$ is convex in w_0 .

$$\frac{dL}{dw_0} = \frac{w_1 + w_0}{w_2 + w_0} \left(-\frac{w_2 + w_0}{(w_1 + w_0)^2} + \frac{1}{w_1 + w_0} \right) = \frac{w_1 - w_2}{(w_1 + w_0)(w_2 + w_0)}$$

and

$$\frac{d^2L}{dw_0^2} = \frac{-(w_1 - w_2)}{(w_1 + w_0)^2(w_2 + w_0)^2} (2w_0 + w_1 + w_2)$$

which is greater than 0 if $w_1 < w_2$.

Proof of Proposition 3: If the customer's consideration set consists of a single product, $\frac{p_1(\beta)}{p_2(\beta)}$ is quasi-convex in β (see Boyd and Vandenberghe [8], Chapter 3 for all the facts on convex and quasiconvex functions used here). Note that as $\log(\cdot)$ is an increasing function, this shows quasi-convexity of $\log(\frac{p_1(\beta)}{p_2(\beta)})$ as well.

We show that for any $1 > \alpha \geq 0$, the set

$$S = \{\beta \mid \frac{p_1(\beta)}{p_2(\beta)} \leq \alpha\}$$

is convex.

Let the product offered during period 1 be indexed by 1 and the product in period 2 be indexed by 2. Recall that we normalize the weight of the no-purchase option to 1. As only one product is offered, S can be written as

$$S = \{\beta \mid (1 - \alpha) \exp^{\beta^\top (x_1 + x_2)} \leq \alpha \exp^{\beta^\top x_2} - \exp^{\beta^\top x_1}\}$$

Let $\beta_1, \beta_2 \in S$, so

$$(1 - \alpha) \exp^{\beta_1^\top (x_1 + x_2)} \leq \alpha \exp^{\beta_1^\top x_2} - \exp^{\beta_1^\top x_1} \quad (15)$$

and

$$(1 - \alpha) \exp^{\beta_2^\top (x_1 + x_2)} \leq \alpha \exp^{\beta_2^\top x_2} - \exp^{\beta_2^\top x_1} \quad (16)$$

Consider the case $1 \geq \alpha$. Multiplying the two sides of the inequalities of (15) and (16), and expanding

$$(1 - \alpha)^2 \exp^{(\beta_1 + \beta_2)^\top (x_1 + x_2)} \leq \exp^{(\beta_1 + \beta_2)^\top x_1} - \alpha \exp^{\beta_1^\top x_2 + \beta_2^\top x_1} - \alpha \exp^{\beta_1^\top x_1 + \beta_2^\top x_2} + \alpha^2 \exp^{(\beta_1 + \beta_2)^\top x_2} \quad (17)$$

We next verify that

$$\begin{aligned} & \exp^{(\beta_1 + \beta_2)^\top x_1} - \alpha \exp^{\beta_1^\top x_2 + \beta_2^\top x_1} - \alpha \exp^{\beta_1^\top x_1 + \beta_2^\top x_2} + \alpha^2 \exp^{(\beta_1 + \beta_2)^\top x_2} \\ & \leq (\alpha \exp^{(\frac{\beta_1 + \beta_2}{2})^\top x_2} - \exp^{(\frac{\beta_1 + \beta_2}{2})^\top x_1})^2 \end{aligned} \quad (18)$$

which reduces to verifying that

$$2 \exp^{\beta_1^\top x_1 + \beta_2^\top x_2} \exp^{\beta_1^\top x_2 + \beta_2^\top x_1} \leq \exp^{\beta_1^\top x_1 + \beta_2^\top x_2} + \exp^{\beta_1^\top x_1 + \beta_2^\top x_2} \quad (19)$$

which is true by the AM-GM inequality. So for $\alpha \leq 1$,

$$(1 - \alpha) \exp^{(\frac{\beta_1 + \beta_2}{2})^\top (x_1 + x_2)} \leq (\alpha \exp^{(\frac{\beta_1 + \beta_2}{2})^\top x_2} - \exp^{(\frac{\beta_1 + \beta_2}{2})^\top x_1})$$

implying $\frac{\beta_1 + \beta_2}{2} \in S$, so S is a convex set.