COLUMBIA UNIVERSITY
Center for Pricing and Revenue Management

WORKING PAPER SERIES: NO. 2010-3

A maximum entropy joint demand estimation and capacity control policy

Costis Maglaras,
Columbia University

Serkan Eren,
C12 Capital Management

2010

A maximum entropy joint demand estimation and capacity control policy

Serkan Eren Costis Maglaras*

October 21, 2010

Abstract

We propose a tractable, data-driven demand estimation procedure based on the use maximum entropy (ME) distributions, and apply it to a stochastic capacity control problem motivated from airline revenue management. Specifically, we study the two fare-class "Littlewood" problem in a setting where the firm has access to only potentially censored sales observations. We propose a heuristic that iteratively fits an ME distribution to all observed sales data, and in each iteration selects a protection level based on the estimated distribution. When the underlying demand distribution is discrete, we show that the sequence of protection levels converges to the optimal one almost surely, and that the ME demand forecast converges to the true demand distribution for all values below the optimal protection level. That is, the proposed heuristic avoids the "spiral down" effect, making it attractive for problems of joint forecasting and revenue optimization problems in the presence of censored observations.

Keywords: Revenue management, censored demand, uncensoring, spiral down, maximum entropy distributions, stochastic approximations.

*Columbia Business School, 3022 Broadway, NY, NY 10027. (se2027@columbia.edu, c.maglaras@gsb.columbia.edu)

A maximum entropy joint demand estimation and capacity control policy

October 21, 2010

Abstract

We propose a tractable, data-driven demand estimation procedure based on the use maximum entropy (ME) distributions, and apply it to a stochastic capacity control problem motivated from airline revenue management. Specifically, we study the two fare-class "Littlewood" problem in a setting where the firm has access to only potentially censored sales observations. We propose a heuristic that iteratively fits an ME distribution to all observed sales data, and in each iteration selects a protection level based on the estimated distribution. When the underlying demand distribution is discrete, we show that the sequence of protection levels converges to the optimal one almost surely, and that the ME demand forecast converges to the true demand distribution for all values below the optimal protection level. That is, the proposed heuristic avoids the "spiral down" effect, making it attractive for problems of joint forecasting and revenue optimization problems in the presence of censored observations.

Keywords: Revenue management, censored demand, uncensoring, spiral down, maximum entropy distributions, stochastic approximations.

1. Introduction

One of the key challenges of revenue management systems is to accurately forecast demand when one has only access to observed sales data that may be censored. It is well known in the area that common uncensoring techniques and their interaction with the iterative application of forecasting and revenue optimization routines may prevent these systems from making optimal decisions in a dynamic setting; c.f., Boyd et al. (2001) and Cooper et al. (2006). In this paper, we propose a tractable and intuitive approach for incorporating and uncensoring sales data into the demand forecast based on maximum entropy (**ME**) distributions that leads to asymptotically optimal control decisions.

A prototypical problem where the above effect has been observed is that of dynamic airline capacity allocation decisions. In its simplest form this problem is described as follows: an airline has a fixed capacity for a flight to sell to the market; there is a low-fare and a high-fare class, and low-fare demand is realized before the high-fare demand; the key decision is to select how many units of capacity to "reserve" for the high-fare demand (i.e., make them unavailable for the low-fare demand that gets realized first) so as to maximize the total expected revenue per flight. The manager does not have accurate demand information, and uses the sales observations in each flight to update the respective demand forecasts for the two fare classes. Demand observations may be censored, when the low-fare demand depletes the capacity that is made available to it, or when the high-fare demand depletes the remaining capacity for the flight; in both cases the manager does not know how much extra demand could have been realized in each of these two classes if there was extra capacity to be allocated. As Boyd et al. (2001) highlighted and later on Cooper et al. (2006) demonstrated analytically, many common forecasting and demand uncensoring methods generate a sequence of forecasts and protection levels that "spirals down" to a suboptimal level. There are two underlying issues: a) the interpretation of censored demand data, and b) the interaction between control and forecasting, and specifically that the choice of a control at any given iteration (flight) serves the joint purpose of revenue optimization and demand learning. A simple illustration of the first issue is the following: suppose that at a particular flight, the manager has 50 units of capacity available for the high-fare demand, and that all of this capacity ends up being sold. What was the true high-fare demand for this flight? Was it 50? Was it more? By how much? A naive approach is to treat the demand as being exactly 50, but this would lead to an underestimation of the true demand, since the actual observation was the event $\{\text{High-fare demand} \geq 50\}$. There are numerous other heuristics that try to reallocate this sales observation to some other demand level that is greater or equal to 50, but, as Cooper et al. (2006) show, many of them do not achieve the desired result.

This paper describes a demand forecasting algorithm based on ideas from ME distributions that can readily incorporate censored sales data, which correspond to fractile observations of the form shown above. The proposed demand forecasting algorithm leads to control decisions (in the form of capacity protection levels) that converge to the optimal ones for the actual underlying demand distribution.

Background on Maximum Entropy distributions: The entropy of a random variable X with probability mass function p_j for all j on some some support set $\mathcal{J}$ is defined by

$$H(X) := -\sum_{j \in \mathcal{J}} p_j \ln p_j;$$

it is also common to use the base 2 logarithm in the above definition. Entropy is non-negative and is a concave function of the probabilities p_j. Entropy is a measure of *average uncertainty* or *disorder* or *randomness* of the random variable. It is also a measure *descriptive complexity* of the random variable, i.e., how much information one needs to describe it. As a concept, entropy is of central focus in the area of information theory, and plays important roles in communications theory, physics, computer science, probability theory, statistics, and economics. The book by Cover and Thomas (1991) offers a thorough introduction to the topic of information theory and explores its connections and the abovementioned fields.

Entropy has also played a central role in estimation theory. In particular, maximum entropy distributions are a useful and intuitive tool in fitting unknown distributions to partial information about the underlying random variables. The most celebrated example comes from statistical mechanics, where Maxwell and Boltzmann showed that the distribution of velocities in a gas at a given temperature is the maximum entropy distribution that corresponds to the temperature constraint that itself fixes the variance of the distribution. In this setting, the maximum entropy solution arises naturally as the correct underlying distribution. In other settings, such as the one that is motivating this study, one may have access to partial information about the underlying distribution, for example specifications of the moments of the distribution, of its fractiles, etc. The decision maker is faced with the question of fitting a model that satisfies these specifications, and in such contexts the maximum entropy criterion provides an approach for how to do that. This approach was advocated to be used in a broad context by Jaynes (1982).

Given a set of specifications of the form $\sum_j r_i(j)p_j = b_i$ for appropriate choices for the functions $r_i(\cdot)$, the canonical ME estimation problem is:

$$\max_p \left\{ -\sum_{j\in\mathcal{J}} p_j \ln p_j \ : \ \sum_{j\in\mathcal{J}} r_i(j)p_j = b_i \text{ for } i = 1,\dots,m, \quad \sum_{j\in\mathcal{J}} p_j = 1, \quad p \geq 0 \right\}, \tag{1}$$

and its solution is

$$p_j^* = e^{\lambda_0 + \sum_{i=1}^m \lambda_i\, r_i(j)} \quad j \in \mathcal{J}, \tag{2}$$

where λ_i is the Lagrange multiplier associated with the linear constraint $\sum_j r_i(j)p_j = b_i$, and λ_0 is the normalization constant such that $\sum_{j\in\mathcal{J}} p_j^* = 1$. Inequality constraints of the form $\sum_j r_i(j)p_j \leq b_i$ can also be added in the above formulation, which is a tractable concave maximization problem that can be solved efficiently even in large problem instances. Fractile constraints that are of particular interest to this paper can be added by setting $r_i(j) = \mathbf{1}_{\{j\geq k\}}$, where $\mathbf{1}_{\{j\geq k\}}$ is the indicator function which is equal to 1 for all $j \geq k$, and k is the position for which the decision maker has fractile information. For continuous distributions the summations

Constraint	ME Distribution
Range $= S = [a, b]$	$U\ [a, b]$
Mean $= \mu$	$\exp(1/\mu)$
Mean $= \mu$, variance $= \sigma^2$	$N(\mu, \sigma^2)$

Table 1: Common constraints and the associated ME distributions.

in the objective function and the constraints are replaced by integration. Some commonly used distributions are the ME distributions that correspond to appropriate constraints (see Table 1 for examples of continuous maximum entropy distributions that admit discretized analogues). ME distributions are of the parametric form given in (2). The parametric degree of the ME distribution depends on the specifications that one starts with. This modeling flexibility is in stark contrast with the common approach of fixing a priori the parametric form of the distribution, e.g., uniform, exponential, gamma, etc., and then searching for the best possible match from within that family; it is easy to see that the latter approach may not even satisfy the problem specifications and introduce significant model selection bias, which is not the case when fitting the ME distribution.

Proposed solution: Returning to the motivating problem for this paper, the key issue faced by the firm is that of building a good demand forecast using sales data that may be censored. As explained earlier, censored observations correspond to fractile information. The policy proposed in this paper is to form a demand forecast by fitting a ME distribution to the observed sales data, which can be formulated as a set of fractile conditions that the demand distribution needs to satisfy. The resulting ME demand distribution provides a tractable and intuitive way of "unconstraining" the sales observations. The firm then uses the resulting demand forecast as if it is the "true" demand distribution, and accordingly computes its protection level for the next flight. A new sales observation is recorded, and the process is repeated. It is worth noting that the proposed policy is "passive" or "myopic" in that it does not choose its controls in a way that jointly optimizes immediate expected revenues with the ability to learn the demand distribution and therefore extract high revenues in future flights. This is done for two reasons: practical considerations, and analytic tractability.

The main analytical contribution of this paper is to establish that the sequence of protection levels generated by the above approach converges to the optimal control that the firm would select if it knew the true underlying demand distribution. The latter is defined by an appropriate critical fractile of the demand distribution, which depends on the relative magnitude of the high and low prices offered by the firm. The intuition behind the convergence result is fairly simple. Suppose that after the first k observations, the firm is underestimating the critical fractile of the demand distribution, e.g., it thinks it corresponds to the point where the demand is equal to 30 when the correct fractile position is at the point where the demand is equal to 35. Then, firm will protect 30 units of capacity for the high-fare demand stream, and will sell-out with a higher probability than it is optimal. In this case, the ME demand forecast will reallocate some of the censored sales observation to higher demand points, driving the critical fractile point towards a higher value. A similar argument applies for the case where the firm is overestimating the critical fractile. In addition to providing a method for asymptotically computing the optimal protection level, the ME forecast is also shown to converge to the correct demand distribution at all points below the critical fractile. For ease of exposition, we assume the low-fare class demand is ample. However, we briefly illustrate in §5, without proof, how the results generalize to any low-fare demand distribution; while this distinction does not affect the optimal protection level, it does affect the learning algorithm since in some instances the unprotected capacity will not be sold out to the low fare demand, and as such the firm will be able to sell more than the original protection level to the high fare demand. We also provide an an interesting extension that allow for the low-fare customers to "buy up" and purchase the high-fare product when the low-fare product is sold out.

Finally and importantly, although the results of the paper are related to a specific motivating problem, the approach of using maximum entropy distributions to uncensor sales observations, and to incorporate other information on the underlying distribution that one may have is of broader interest. In the concluding section we give a brief illustration how it could be used in the context of fitting a willingness-to-pay distribution that can then be used for making pricing decisions.

Literature survey: This paper is directly motivated by the observation in Boyd et al. (2001) and the analysis in Cooper et al. (2006). Another paper by Weatherford (2002) has also showed that most heuristic uncensoring techniques do not avoid the "spiral-down" effect. van Ryzin and McGill (2000) propose an adaptive, Robins-Monro, algorithm that controls via appropriate feedback signals the protection levels depending on whether the allocated capacity for some fare class (or set of nested fare classes) is sold out or not; it increases the protection level when the capacity is sold out, and decreases it otherwise. They prove that the proposed algorithm converges to the optimal protection levels, and as such avoids the spiral down phenomenon. A closely related paper is by Kunnumkal and Topaloglu (2009).

There are a set of related papers that develop adaptive inventory ordering policies for the newsvendor problem when the demand distribution is unknown; see, e.g., the papers by Burnetas and Smith (2000), Godfrey and Powell (2001) and Huh and Rusmevichientong (2008). The above sets of papers take the approach of directly adjusting the protection levels or the inventory ordering decisions, bypassing the demand estimation step that is central in our approach. In contrast, our paper provides an explicit algorithm for demand estimation based on censored observations, which is then applied to the airline capacity control problem, but may of more general interest.

The closest paper to our work is the one by Huh et al. (2009) that proposes to use the demand estimation procedure based on the non-parametric Kaplan-Meier estimator, and then select the control based on that estimated demand distribution. They analyze their approach in the context of the newsvendor model, and prove that their inventory ordering decisions converge to the optimal newsvendor quanitity defined by a critical fractile. They also show that their demand estimation procedure asymptotically characterizes correctly the unknown demand distribution up to the critical fractile position. The results in their paper mirror many of our findings, but the demand estimation procedures based on the Kaplan-Meier estimator and Maximum Entropy distributions, respectively, are quite different in their structure the type of information that they can incorporate in their formulation, such as additional information about the moments of the distribution, and their potential applications.

A subtle point that underlies our work as well as that in Kunnumkal and Topaloglu (2009), Huh et al. (2009) is what happens when the realized sales volume is equal to the firm's capacity, which could happen if the demand was equal to the capacity, or if the demand was larger than the allowed capacity and the resulting demand observation was censored. The above papers assume that the seller can observe whether the demand observation was censored, i.e., when the "sales = capacity" the seller can differentiate between the two cases "demand = capacity" and "demand > capacity." This assumption can be fairly restrictive in many operational settings. Without this assumption, the above algorithms will converge to the "optimal capacity - 1." While this difference may be of small significance in practical settings, it does raise the question of whether one can do better. One possible way of circumventing this assumption while achieving the desired learning, is to allow the algorithm to "experiment" around what seems to be the optimal capacity level. This will allow the seller to observe some uncensored observations at the optimal capacity level and as such be able to make a crisp determination of the demand distribution up to the optimal capacity level and correctly deduce the value of the optimal capacity control. We motivate one possible choice for the degree of experimentation, which seems to lead to fast learning and convergence to the correct optimal control. Given the non-stationary nature of demand in many real applications, this seems like a reasonable and practical solution. We do not address the issue of what is the minimal level of experimentation that would still allow us to deduce the optimal control, but refer to a recent paper by Besbes and Muharremoglu (2010) that address this point for a related model in the context of an appropriate asymptotic analysis.

The analytical results of this paper make use of standard results from adaptive algorithms and stochastic approximations. The particular references that are useful for our work are Kushner and Yin (2003) and Benveniste et al. (1990), and the key result that we use in our paper is reproduced (without proof) in Appendix A.

One of the few papers that deal with "maximum entropy" in revenue management literature belongs to Bilegan et al. (2004) who simply formulate a dual geometric program for the convex ME problem for capacity allocation and demonstrate how to solve it in a short paper. To the best of our knowledge the operations management and revenue management literatures have not explored the use of ME techniques to approximate unknown demand or willingness-to-pay distributions.

Finally, we conclude this section by listing a few references that are partially related to our work. First, there is a significant body of literature that studies capacity control or newsvendor problems with uncertain demand distributions using some form of a worst case criterion. Examples in this area include the papers by Gallego and Moon (1993), Bertsimas and Thiele (2004), Perakis and Roels (2006), and Ball and Queyranne (2009). The above papers do not involve learning. Second, there is a growing literature in joint learning and price optimization, which is somewhat related to the motivating problem and the demand estimation procedure of this paper. Incorporation of partial information is typically done in a Bayesian setting under some parametric assumptions for the willingness to pay distribution and using conjugate pairs of distributions to maintain tractability; see, e.g., Lobo and Boyd (2003), Aviv and Pazgal (2005), Araman and Caldentey (2005), Farias and Van Roy (2006), and the references therein. Assuming a parametric family of distributions for the unknown demand runs the risk of model mis-specification due to the arbitrariness of that assumption. A non-parametric approach that is asymptotically optimal is due to Besbes and Zeevi (2009).

2. Single-resource capacity control with two fare-classes

We study a repeated version of Littlewood's two-period capacity allocation problem, where the distributions for the two classes of potential demand are unknown, but where the seller can try to learn the demand distributions from (potentially censored) sales observations. We first describe the static version of Littlewood's problem under full demand information, and then proceed to pose the repeated version of this problem with no prior demand information.

2.1 Littlewood's model: Full information, static benchmark

A firm has C identical units of a good to sell over two time periods to two demand classes indexed by $i = 1, 2$. The class-2 demand, denoted by D_2, arrives first and pays a price of p_2, followed by the class-1 demand, denoted by D_1, which pays $p_1 > p_2$. The salvage value is assumed without loss of generality to be 0. The two demands are discrete random variables that are independent of each other, and independent of any capacity control decisions made by the system manager, drawn from some distributions F_i for $i = 1, 2$. The firm controls whether to accept or reject each class-i request for one unit of its capacity, and its objective is to allocate the available capacity to the two demand streams described above so as to maximize its total expected revenue over the entire selling horizon. It is well known that the structure of the firm's optimal policy takes the form of a threshold, or protection level, denoted by L, which sets the number of units of capacity to be reserved for the high-fare class demand, D_1. That is, class 2 demand requests are accepted as long as it the remaining capacity left for period 1 for the high-fare demand stream is greater than L, and are rejected otherwise. In summary, the firm's problem is to choose the protection level L to maximize its expected revenue

$$\max_{0 \leq L \leq C} \mathbf{E} \left[p_1 \min(D_1,\ \max(C - D_2, L)) + p_2 \min(D_2, C - L) \right], \tag{3}$$

where the expectation is taken with respect to the two demand distributions. The term $\min(D_2, C - L)$ is the sales for the low-fare class, which arrives first; and consequently, the high-fare class sales is the minimum of demand D_1 and the remaining number of seats $C - \min(D_2, C - L) = \max(C - D_2, L)$. If D_1 and D_2 were continuous random variables, then the optimal protection-level L^* would be given by the following equality

$$p_1\, \mathbb{P}(D_1 \geq L^*) = p_2 \ \text{ if and only if } \ F_1(L^*) = \gamma := 1 - p_2/p_1. \tag{4}$$

This condition is commonly referred to as Littlewood's rule. The left hand side of the above expression equates the marginal expected revenues from an immediate sale at price \$$p_2$ versus a potential sale in the next period at the higher price \$$p_1$. For discrete demand distributions, the optimal protection level satisfies

$$p_2 < p_1\, \mathbb{P}(D_1 \geq L^*) \text{ and } p_2 \geq p_1\, \mathbb{P}(D_1 \geq L^* + 1) \;\Leftrightarrow\; \gamma > F_1(L^* - 1) \text{ and } \gamma \leq F_1(L^*) \tag{5}$$

That is optimal protection level is given by

$$L^* = \inf\{L \;:\; F_1(L) \geq \gamma\}, \tag{6}$$

Equation (5) highlights an additional challenge for the discrete version of the problem that was highlighted in the introduction: for a given protection level L, the sales event $D_1 \geq L$ provides a censored observation, and yields information for the underlying high fare demand only up to $L - 1$. As a consequence, to identify the optimal level L^*, either one needs to sample at level $L^* + 1$, or assume that one can distinguish events $D_1 = L^*$ and $D_1 > L^*$ when high fare sales occur at L^*. Going forward we will assume we cannot distinguish events $D_1 = L$ and $D_1 > L$ for any given protection level L, and consequently that $D_1 \geq L$ is a censored observation. To overcome this challenge and learn the fractile information at L^* so as to be able to correctly identify it as the optimal control, we will sample at $L + 1$ with probability q defined below. We find this to be a more natural approach that is closer to real business practice, e.g. airlines occasionally open up their protection levels to better sample underlying demand distribution. Specifically, we propose and analyze a randomized policy, whereby the seller will protect L^* units with probability $1 - q^*$ and protect $L^* + 1$ with probability q^* where q^* chosen to satisfy $(1 - q^*)\; F_1(L^* - 1) + q^*\; F_1(L^*) = \gamma$.

One motivation for the above choice of the randomization parameter q^* is as follows: If the support of the demand distribution as well as the optimal L^* are large and the cumulative distribution F_1 is well behaved, then it is plausible that the objective in (3) is well approximated by a linear function around L^*, in which case q^* would be the optimal randomization parameter. This argument is, of course, heuristic, and indeed one could use any value for the randomization parameter q greater than 0 and less than 1 and still guarantee the convergence result that we will derive in the next section.

2.2 Repeated Littlewood's problem with unknown demand distributions

The model analyzed in this paper is a repeated version of Littlewood's problem in settings where a) the distribution of the high-fare demand, D_1, $F_1(\cdot)$ is unknown, and b) the seller can estimate the unknown distribution based on sales observations; the distribution for the low-fare demand D_2 may also be unknown, but as explained above is not needed for the characterization of the optimal protection level L^*. In broad terms, the seller can estimate the high-fare demand distribution given past sales observations. The goal is to describe an estimation procedure and an associated control policy that will converge to the optimal protection level L^*. The core of the problem is that the sequence of protection levels affect the sequence of observations, and thus the resulting estimation output. As was highlighted in the introduction this may lead to sub-optimal estimation and control outcomes.

In more detail, we consider a sequence of instances of the two-period Littlewood problem defined above, which we index by $k = 1, 2, \ldots$. In each instance k, the seller applies the capacity control L^k, the realized demands are D_i^k, and the realized sales are given by $S_2^k = \min(C - L^k, D_2^k)$ and $S_1^k = \min(L^k + x^k, D_1^k)$, for the low-fare and high-fare demand stream respectively, and where $x^k = (C - L^k - D_2^k)^+$ is the unused capacity from the low-fare class. The realized revenue is $\min(C - L^k, D_2^k) \cdot p_2 + \min(L^k + x^k, D_1^k) \cdot p_1$.

We make the following assumptions on the demand distributions:

Assumption 1: $D_2^k \geq C$ with probability 1 for all k.

Assumption 2: Let S denote the size of the support of F_1, and $\pi_j = P(D_1 = j)$ for $j = 0, 1, \ldots, S-1$. Then, $\pi_j > \epsilon$ for some $\epsilon > 0$ for $j = 0, \ldots, L^*$ and $S \geq 1/\epsilon + 1$.

It is well established that lower price demand does not effect the choice of optimal protection level in Littlewood's setup, so one would expect a similar result in this setting. Assumption 1 corresponds to the most aggressive censoring of high fare sales, and proving our result under this conservative assumption allows us to ease notation in formal exposition below and to discuss its extension to general case later briefly in §5. Under assumption 1, $x^k = 0$ and the realized revenue is $(C - L^k) \cdot p_2 + \min(L^k, D_1^k) \cdot p_1$, for all k. Assumption 2 is an optional mild technical condition that can always be satisfied by selecting S to be sufficiently large, and is used to prove Proposition 3 in the appendix.

The observation history is $\{(L^1, S_1^1, S_2^1),\ (L^2, S_1^2, S_2^2), \ldots, (L^k, S_1^k, S_2^k)\}$. Under assumption

1, $S_2^k = C - L^k$ for all k, and thus all information is captured in the sequence $\mathcal{I}^k := \{(L^i, S_1^i)\}_{i=1}^k$. As discussed previously, we assume that the event $D_1^k = L^k$, which results in $S_1^k = L^k$, provides a censored observation.

Problem formulation: For all $k \geq 1$, given the information set $\mathcal{I}^k$, find a control $L^{k+1} : \mathcal{I}^k \to [0, S]$ for $k \geq 1$, such that $L^k \to L^*$ almost surely as $k \to \infty$, for L^* identified in (6).

Once convergence of the protection levels L^k to L^* has been established, one could switch to a more refined criterion that studies some measure of revenue loss from that obtained under L^*, or some full information benchmark where the firm would observe the demand realization as opposed to the potentially censored sales observations. We will not pursue this in this paper.

3. Proposed policy based on Maximum Entropy distributions

The structure of the proposed solution is motivated from what is typically observed in practice, for example in the airline industry, where a two-step procedure is adopted: a) build some type of a forecast for D_1^k based on $\mathcal{I}^k$, and b) compute a protection level L^k given that forecast. Let $F_{\mathcal{I}^k}$ denote the estimated high-fare class demand distribution after the first k observations. The type of policies that we will consider are "passive" or "myopic" in the sense that at every point in time they select the protection level L^{k+1} as if $F_{\mathcal{I}^k}$ was the correct demand distribution. This essentially reduces the joint estimation and control problem posed above to one of estimation of a critical fractile of a demand distribution based on censored observations.

The demand forecasting procedure we propose makes use of a more aggregated form of the observed information, which is independent of the sequence of the various sales observations. Specifically, given $\mathcal{I}^k$, the uncensored and censored information recorded thus far is summarized by the vectors $K^k \in \mathbb{R}^S$ and $J^k \in \mathbb{R}^S$, where

$$K_j^k = \# \text{ of uncensored observations at position } j, \text{ and,}$$
$$J_j^k = \# \text{ of censored observations at position } j;$$

clearly $\sum_j (K_j^k + J_j^k) = k$ for all k. We will summarize this aggregated observation history by $\theta^k \in \mathbb{R}^{S \times 2}$ defined as follows

$$\theta^k := (\kappa^k, \zeta^k) := (K^k/k, J^k/k); \tag{7}$$

κ_j^k and ζ_j^k are frequencies of uncensored and censored observations at position j, respectively.

The proposed policy fits a maximum entropy (ME) distribution to the observation history, which in itself provides a systematic way in which to "re-allocate" the censored sales observations into possible (higher) demand realizations. The intuition behind this policy is that censored observations offer fractile information that can be readily incorporated in picking the ME distribution that best fits the sales observations θ^k.

Let $\eta_j^k = \kappa_j^k + \zeta_j^k$ be the frequency of observations at j if one does not distinguish between censored and uncensored observations. Let $p \in \mathbb{R}_+^S$, $z \in \mathbb{R}_+^{S \times S}$, where p_j is the probability assigned to observing a demand realization in position j, and z_{ij} denotes the probability mass allocated to position j due to censored observations in position $i \leq j$. The ME distribution that corresponds to the observation vector θ^k is computed as follows:

$$\max_{p,\, z} \quad -\sum_j p_j \ \ln p_j \tag{8}$$

$$s.t. \quad p_j = \kappa_j^k + \sum_{i \leq j} z_{ij}, \quad \forall j \tag{9}$$

$$\sum_{j \geq i} z_{ij} = \zeta_i^k, \quad \forall i \tag{10}$$

$$z_{ij} = 0, \quad \forall i < j, \qquad z_{ij} \geq 0, \quad \forall i, j \tag{11}$$

$$\sum_j p_j = 1 \ , \tag{12}$$

This is simplified in the following result (the proof is given in the Appendix B). Specifically, if we define the auxiliary vector $\tilde{\kappa} \in \mathbb{R}_+^S$ as follows: $\tilde{\kappa}_j^k = \kappa_j^k$ for $j = 0 \ldots S-2$, and $\tilde{\kappa}_{S-1}^k = \eta_{S-1}$, then (8)-(12) can be reduced to:

$$\min_p \quad \sum_j p_j \ \ln p_j \tag{13}$$

$$s.t. \quad p_j \geq \tilde{\kappa}_j^k, \quad \forall j \tag{14}$$

$$\sum_{i \geq j} p_i \geq \sum_{i \geq j} \eta_i^k \quad \text{if} \ \ \zeta_j^k > 0 \tag{15}$$

$$\sum_j p_j = 1. \tag{16}$$

The algorithm we propose can be summarized as follows:

Algorithm 1: Maximum entropy capacity allocation for two fare-classes

1. At each observation k, update the vector $\theta^k := (\kappa^k, \zeta^k)$ according to (7)

2. Given θ^k, compute the ME probability mass function p_{θ^k} through (13)-(16); denote the corresponding distribution function as $F_{\theta^k}(\cdot)$.

3. Set $L(\theta^k) = \min\{L \mid F_{\theta^k}(L) \geq \gamma\}$.

4. Implement $L(\theta^k)$ with probability $1-q(\theta^k)$, and $L(\theta^k)+1$ with probability $q(\theta^k)$, where $q(\theta^k)$ is the unique solution to

$$(1 - q(\theta^k))\ F_{\theta^k}(L(\theta^k) - 1) + q(\theta^k)\ F_{\theta^k}(L(\theta^k)) = \gamma\ . \tag{17}$$

5. Observe new sales in period $k+1$ and go to step 1.

To ease exposition going forward, we will use the shorthand notation $L^k := L_{\theta^k} := L(\theta^k)$ and $q^k := q_{\theta^k} := q(\theta^k)$ depending on the context. The "passive" or "myopic" structure of the proposed policy is reflected in steps 3 and 4 above that treat the high-fare distribution estimate $F_{\theta^k}(\cdot)$ as if it is the correct demand distribution in every iteration.

4. Convergence analysis of the ME capacity allocation policy

This section proves that Algorithm 1 yields a sequence of controls $\{L^k\}$ that converges to the optimal level L^*. In addition to correctly identifying the γ-fractile of the high-fare distribution, the ME demand estimation approach will correctly approximate the entire high-fare class demand distribution up to L^*. As a byproduct of our approach one can also show that the estimates of the probabilities that the high-fare demand will be equal to j, denoted by $p_j^k = P_{F_{\theta^k}}(D_1 = j)$, converge to the correct probabilities π_j for all $j \leq L^*$; i.e., the forecasting procedure based on ME distributions "learns" the demand distribution correctly below L^*.

4.1 Preliminaries

Let $\pi \in [0,1]^S$ represent the probability mass distribution of the actual high-fare class demand. Through the ME algorithm, the vector (K^k, J^k) evolves recursively as

$$(K^{k+1}, J^{k+1}) = (K^k, J^k) + (W^{k+1}, Q^{k+1}) \tag{18}$$

where $W^{k+1}, Q^{k+1} \in \mathbb{R}^S$ are random vectors satisfying

$$\mathbf{P}(W_j^{k+1} = 1) = \begin{cases} \pi_j & \text{for } j < L^k \\ q^k \ \pi_{L^k} & \text{for } j = L^k \\ 0 & \text{otherwise} \end{cases}, \quad \mathbf{P}(Q_j^{k+1} = 1) = \begin{cases} (1-q^k) \ \sum_{i \geq L^k} \pi_i & \text{for } j = L^k \\ q^k \ \sum_{i > L^k} \pi_i & \text{for } \ j = L^k + 1 \\ 0 & \text{otherwise,} \end{cases}$$

for $j = 0 \ldots S-1$; i.e., vectors W^{k+1} and Q^{k+1} track the realization of uncensored and censored observations at step $k+1$. Recall that $\theta^k = (K^k, J^k)/k$ and define

$$f(\theta^k) := (\mathbf{E}(W^{k+1}), \mathbf{E}(Q^{k+1})), \quad \text{and} \quad g(\theta^k) := f(\theta^k) - \theta^k, \tag{19}$$

where the expectation is taken with respect to the (unknown) true demand distribution for the high-fare class demand so that

$$\mathbf{E}(W_j^{k+1}) = \begin{cases} \pi_j & \text{for } j < L^k \\ q^k \ \pi_{L^k} & \text{for } j = L^k \\ 0 & \text{otherwise} \end{cases}, \ \mathbf{E}(Q_j^{k+1}) = \begin{cases} (1-q^k) \ \sum_{i \geq L^k} \pi_i & \text{for } j = L^k \\ q^k \ \sum_{i > L^k} \pi_i & \text{for } j = L^k + 1 \\ 0 & \text{otherwise.} \end{cases} \tag{20}$$

Dividing both sides of (18) by $1/(k+1)$, we get that

$$\begin{aligned} \theta^{k+1} &= \frac{k \ \theta^k}{k+1} + \frac{(W^{k+1}, Q^{k+1})}{k+1} \\ &= \frac{k \ \theta^k}{k+1} + \frac{(W^{k+1}, Q^{k+1}) - f(\theta^k)}{k+1} + \frac{f(\theta^k)}{k+1} \\ &= \theta^k + \frac{(W^{k+1}, Q^{k+1}) - f(\theta^k)}{k+1} + \frac{f(\theta^k) - \theta^k}{k+1} \\ &= \theta^k + \frac{1}{k+1} \ \delta M^k + \frac{1}{k+1} \ g(\theta^k), \end{aligned} \tag{21}$$

where $\delta M^k := (W^{k+1}, Q^{k+1}) - f(\theta^k)$ are bounded martingale differences for each k, and $g(\theta^k) = f(\theta^k) - \theta^k$ is the deterministic function governing the mean drift of the process. The martingale properties of δM^k follows from the definition of $f(\theta^k)$, and the fact that the components of both (W^{k+1}, Q^{k+1}) and $f(\theta^k)$ are bounded by 1.

Equation (21) describes an adaptive algorithm, whose asymptotic properties can be analyzed using standard techniques from stochastic approximations. We will follow the so called "ODE (Ordinary Differential Equation) approach" (see Kushner and Yin (2003)) that relies on an asymptotic analysis of appropriate continuous approximations to the process $\{\theta^k\}$. By establishing convergence of θ^k to an appropriate limit point, we will be able to conclude that a) the ME probability estimates for different demand realizations converge to the correct probabilities π_j for all $j \leq L^*$, and b) that the protection level computed using Algorithm 1 converges to L^*. The main intuition behind it is that the effect of random martingale difference terms δM^k vanishes as k gets larger, and the process can be approximated accurately by the limit paths of the continuous ODE $\dot{\theta}(t) = g(\theta(t))$, where $\theta(t)$ is a continuous approximation to the discrete process θ^k, and $\dot{\theta}(t)$ is the time derivative of this approximation. That is, intuitively, as $1/(k+1) \to 0$, θ^k changes slowly, and, in the absence of the δM^k terms, equation (21) yields roughly

$$g(\theta^k) = \frac{\theta^{k+1} - \theta^k}{1/(k+1)} \approx \dot{\theta} \mid_{\theta^k} \quad .$$

The continuous approximations to the discrete process $\{\theta^k\}$ help make the above intuition exact. We refer the reader to Kushner and Yin (2003) for a rigorous development of this approach. Our argument will use Theorem 2.1 (Ch. 5, pg. 127) from Kushner and Yin (2003), which we state for the special case of bounded $\{\theta^k\}$ process with bounded martingale differences in the Appendix A together with its necessary conditions (A.4.3.1), and (A.5.2.1)-(A.5.2.6). Below we provide the outline of our argument for establishing $\lim_{k\to\infty} L^k = L^*$:

1. Section 4.2 studies the asymptotic behavior of the continuous, deterministic dynamical system governed by the ODE $\dot{\theta}(t) = g(\theta(t))$.

2. Section 4.3 verifies that the conditions (A.4.3.1), and (A.5.2.1)-(A.5.2.5) needed by Kushner and Yin's theorem are satisfied by the process $\{\theta^k\}$ and the ODE, and invokes their result to complete our proof.

4.2 Analysis of the ODE

Denote the domain of problem (13) as $\mathcal{D}(\theta) \in \mathbb{R}^S$ for any given $\theta = [\kappa, \zeta]$. Let $p_\theta \in \mathcal{D}(\theta)$ be the optimal solution of problem (13) for a given θ, and let $[\kappa_\theta, \zeta_\theta]$ denote the corresponding vectors of reallocated uncensored and censored observation frequencies at this solution such that $\kappa_\theta = \kappa$ and $p_{\theta,j} = \kappa_{\theta,j} + \zeta_{\theta,j}$ for all j. Hence, p_θ is the probability mass function of the ME distribution implied by any given θ through problem (13), and $F_\theta(\cdot)$ is the corresponding cumulative distribution function. Hence, the corresponding protection level produced by the ME algorithm is $L(\theta) := \min\{L \mid F_\theta(L) \geq \gamma\}$, and q_θ is the associated randomization probability specified in (17).

We start with some preliminary lemmas that establish some of the necessary structural properties of the ODE $\dot{\theta}(t) = g(\theta(t))$; their proofs are given in Appendix B.

Lemma 1 *$\mathcal{D}(\theta)$ is a continuous correspondence.*

Lemma 2 *The maximum entropy distribution computed in Algorithm 1, denoted by p_θ, is continuous in θ.*

Lemma 3 *The function $g(\theta)$ defined in (19) is continuous in θ.*

As mentioned above, the limit paths of the continuous (mean direction) ODE $\dot{\theta}(t) = g(\theta(t))$ and the process $\{\theta^k\}$ show similar asymptotic behavior. In fact, if the process $\{\theta^k\}$ converges to a unique equilibrium point, say θ_s, this point would be the unique *stationary* solution of the continuous ODE equation $\dot{\theta}(t) = g(\theta(t))$, i.e. $\theta(t) = \theta_s$, $\forall t \geq 0$, under some regularity conditions. However, in our case, the process $\{\theta^k\}$ converges to a stable set Θ identified below rather than a unique point. The ODE can still have a unique stationary solution in this case, and it turns out it has indeed one in our setup. However, this does not mean that θ_s is the only limit point of the ODE *over all initial conditions*. The set of limit paths of the ODE over all initial conditions is a subset of Θ in this case. Therefore, identifying this unique stationary solution is not essential in our case, as a consequence, we present it in Propositions 2 and 3 in the appendix, proofs of which are lengthy and omitted for brevity.

We next characterize the candidate limit set Θ of the process $\{\theta^k\}$; and show that it is globally asymptotically stable[1] for the ODE $\dot{\theta}(t) = g(\theta(t))$. This will be needed to prove convergence of $\{\theta^k\}$ further below. The proof is given in Appendix B. We use the candidate Lyapunov function

$$V(\theta) = \sum_{j<L^*} ||\kappa_j - \pi_j|| + (\gamma - \sum_{j<L^*} \pi_j - \kappa_{L^*})^+, \tag{22}$$

where $||x|| := \sqrt{x^2}$ is the L^2 norm on $\mathbb{R}$, and $x^+ = \max\{x, 0\}$ is the positive part of x. This function is continuous everywhere and continuously differentiable almost everywhere, and the left and right partial derivatives are strictly negative at breakpoints of the gradient for each component of θ.[2]. Observe that for all θ such that $V(\theta) = 0$, $L(\theta) = L^*$.

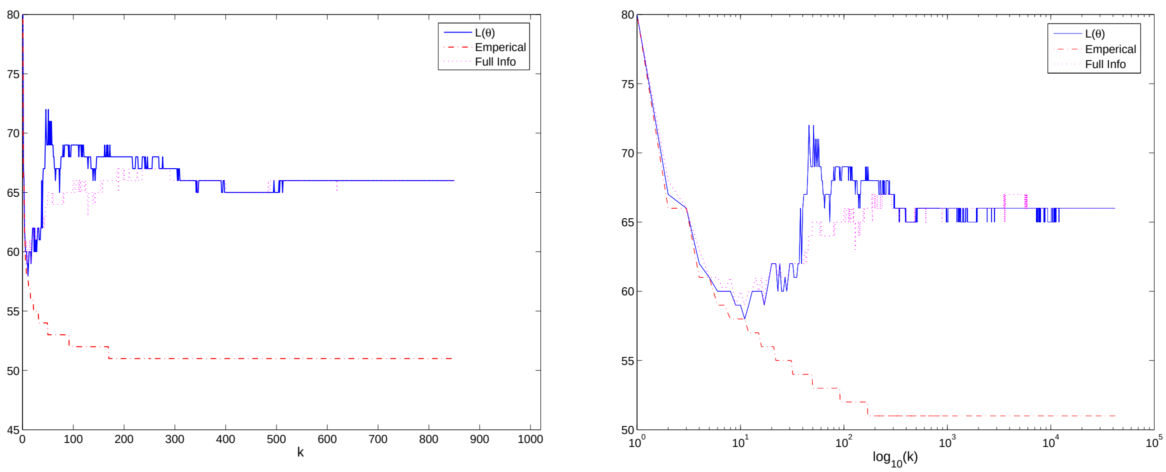

(a) Linear scale

(b) Logarithmic scale

Figure 1: Protection levels produced by the ME algorithm, the empirical distribution, and the uncensored actual demand histogram at each iteration are compared. For the example, $p_2/p_1 = 1 - \gamma = 0.5$, $S = 200$, $D_1 \sim U[50, 80]$.

Proposition 1 *Under assumptions 1 and 2, the set* $\Theta := \{\theta \mid \kappa_j = \pi_j, \ \forall j < L^*, \quad \gamma -$

[1]Stability theory studies properties of solutions for dynamic systems expressed through differential equations. Let $N_\delta(\theta)$ be a δ-neighborhood of θ. In broad terms, a set $A \subset H$ is *locally stable* if for each $\delta > 0$ there is a $\delta_1 > 0$ such that all trajectories $\theta(t)$ of ODE $\dot{\theta}(t) = g(\theta(t))$ starting in $N_{\delta_1}(A)$ never leave $N_\delta(A)$. If the trajectories ultimately go to A, then A is *locally asymptotically stable in the sense of Lyapunov.* If this holds for all initial conditions then the set A is said to be *globally asymptotically stable* for the ODE. See Khalil (1996) for detailed explanations of these stability concepts.

[2]Lyapunov functions act as potential (or penalty) functions for the state of dynamical systems. Lyapunov's second stability theorem states that if there is function $V : \mathbb{R}^S \to \mathbb{R}$ satisfying $V(\theta(t)) \geq 0$ with equality if and only if $\theta(t) \in A$ and $\dot{V}(\theta(t)) \leq 0$ with equality if and only if $\theta(t) \in A$ then A is globally asymptotically stable in the sense of Lyapunov for the system governed by ODE $\dot{\theta}(t) = g(\theta(t))$ where $\dot{V}(\theta(t))$ denotes the time derivative of the function $V(\theta(t))$. The intuition behind Lyapunov's theorem is that the existence of such a function implies that the potential of the system must strictly decrease in time eventually reaching to 0.

$\sum_{j<L^*} \pi_j \leq \kappa_{L^*}\}$ *is globally asymptotically stable for the ODE* $\dot{\theta}(t) = g(\theta(t))$.

Note that under assumptions 1 and 2 we have $L_\theta = L^*$ for $\forall \theta \in \Theta$ as

$$\sum_{j<L^*} p_{\theta,j} = \sum_{j<L^*} \pi_j < \gamma \quad \text{and} \quad \sum_{j\leq L^*} p_{\theta,j} \geq \sum_{j<L^*} \pi_j + \kappa_{L^*} \geq \gamma, \tag{23}$$

satisfying the optimality condition for the protection level for discrete demand distributions stated in (5).

4.3 Proof of convergence

Finally, having established a candidate limit set that is asymptotically stable in which the desired result for the protection level is guaranteed, we show that the process $\{\theta^k\}$ indeed converges to this set Θ by using a theorem provided by Kushner and Yin (2003).

Theorem 1 *Under assumptions 1 and 2,* $\lim_{k\to\infty} L^k = L^*$ *almost surely.*

Proof: The proof follows from Kushner and Yin's (Kushner and Yin, 2003) Theorem 2.1 (Ch. 5, pg. 127) for identifying limits of adaptive algorithms making use of continuous time ordinary differential equations (ODE). This theorem and its conditions, (A.4.3.1) and (A.5.2.1)-(A.5.2.5), are supplied in Appendix A.

In our case the process $\{\theta^k\}$ is bounded in the hyperrectangle $H := [0,1]^{S\times 2}$ with probability 1 by construction; hence condition (A.4.3.1) is satisfied. For the ME algorithm, we have that $\beta^k = 0$ and $\epsilon^k = 1/k$, and $\mathbf{E}[Y^k \mid \theta^0,\ Y^i,\ i < k] = g(\theta^k) = f(\theta^k) - \theta^k$. Consequently, conditions (A.5.2.2), (A.5.2.4) and (A.5.2.5) are trivially satisfied. We also satisfy condition (A.5.2.3) as shown in Lemma 3 above. Condition (A.5.2.1) is satisfied as $\sup_k \mathbf{E}\,|Y^k|^2 \leq S < \infty$ by construction. Also, we have shown in Proposition 1 that the set $A_H := \Theta$ is globally asymptotically stable stable in the sense of Lyapunov for $\dot{\theta}(t) = g(\theta(t))$. Consequently, applying Theorem 2.1 of Kushner and Yin, we conclude that limit points of the process $\{\theta^k\}$ are in Θ with probability 1. However, as shown before in equation (23) above, we have that $L_\theta = L^*$ for all $\theta \in \Theta$. $\square$

In addition, Theorem 1 and Proposition 1 together yield the following result, which states that the ME algorithm also asymptotically provides the correct values of the underlying discrete demand distribution up to level L^*. (This follows from the convergence of θ to a limit set Θ, in which κ_j^k must be converging to π_j.)

Corollary 1 $p_j^k \to \pi_j$ *for all* $j < L^*$.

Figure 1 provides an example of the ME algorithm and the spiral-down effect. As illustrated, the protection levels attained by the empirical distribution of observations spiral down as predicted by Cooper et al. (2006). The protection levels L^k provided by the ME algorithm converge to the correct level. Also, observe that the controls obtained by accumulating the true (uncensored) demand observations, which corresponds to the (first) best case for the firm, seem to converge almost at the same rate as those provided by the ME algorithm. The uncensored full information controls seem to under protect for the high-fare class demand while approaching the optimal level; whereas the protection levels $\{L^k\}$ over protect.

5. Extensions

5.1 Stochastic low-fare demand

A straightforward extension would allow for the low-fare demand to be drawn according to probability mass function π_j^2 for $j = 0, \ldots, S_2 - 1$ for some support $S_2 > 0$. In this case, given a protection level L^k, the amount of unused capacity made available to the high-fare demand is $L^k + \max(C - L^k - D_1^k, 0)$; i.e., if $D_1^k < C - L^k$, then the available capacity to the high-fare demand is greater than L^k, allowed the firm to collect uncesnored demand observations even for positions $j > L^k$ when $D_1^k < C - j$. Intuitively, this should be helpful in the demand forecasting procedure.

Algorithm 1 is still applicable in this setting. The only difference is that the function $g(\theta)$ in the convergence proof needs to be adjusted; this has a consequence on the evolution of the state vector θ^k. Note that an uncensored observation at level $j \geq L^k = L_{\theta^k}$ can occur only when low-fare demand is sufficiently low, i.e., $D_2^k < C - j$. The resulting vector of uncensored mass in equation (20) becomes

$$\mathbf{E}(W_j^{k+1}) = \begin{cases} \pi_j & \text{for } j < L_{\theta^k} \\ \pi_{L_{\theta^k}}(q_{\theta^k} + (1-q_{\theta^k})\sum_{i<C-L_{\theta^k}} \pi_i^2) & \text{for } j = L_{\theta^k} \\ \pi_j \sum_{i<C-j} \pi_i^2 & \text{otherwise} \end{cases} . \tag{24}$$

The structural results of Lemmas 1, 2, and 3 remain valid. In fact, we can show that Proposition 1 is also valid in this case with the same set Θ and using the same Lyapunov function in equation (31). The only difference in the proof, due to equation (24) above, is that the function $g(\theta)$ in equation (33) now satisfies

$$g(\theta) = \begin{cases} \pi_j - \kappa_j & \text{for } j < L_\theta \\ \pi_{L_\theta}(q_\theta + (1-q_\theta)\sum_{i<C-L_\theta} \pi_i^2) - \kappa_{L_\theta} & \text{for } j = L_\theta \\ \pi_j \sum_{i<C-j} \pi_i^2 - \kappa_j & \text{otherwise} \end{cases} . \tag{25}$$

Given Proposition[3] 1, Theorem 1 also applies to the case with stochastic low-fare demand. This result is in line with the fact that the Littlewood's formula in (4) is independent of the distribution of the low-fare demand.

5.2 Model with buy-up

In practical settings, if the low-fare capacity is depleted, low-fare customers may still be willing to "buy-up" and purchase the high-fare product. A model commonly studied within the capacity allocation setting is the so called "buy-up" model. The sequence of events is similar in that the low-fare class demand D_2^k arrives before the high-fare class demand D_1^k. However, if the capacity allocated for the low-fare class customers runs out, i.e. $C - L^k < D_2^k$, some of these customers are assumed to upgrade their demand and request high-fare seats. A common assumption in literature is that each such customer is likely to upgrade with some constant probability α. Hence, at each iteration there is an additional source possible demand for high-fare seats that

[3]The function $g(\theta)$ has extra terms only for $j \geq L_\theta$. The gradient of the Lyapunov function in (32) is equal to 0 for components $j < L^*$. Using this observation, one can extend the proof of Proposition 1 for the case with $L_\theta > L^*$. Also, for $L_\theta < L^*$, the inequalities (34) and (35) remain valid and the subsequent steps of proof are unchanged. Finally, for $L_\theta = L^*$, equation (44) takes the form $\dot{V}(\theta) = -\sum_{j<L^*} ||\kappa_j - \pi_j|| + \kappa_{L^*} - \pi_{L_\theta}(q_\theta + (1-q_\theta)\sum_{i<C-L_\theta} \pi_i^2)$, which is less than the value of the right hand side in equation (44) for any given θ, and hence the drift inequalities remain valid in this case as well.

has a binomial distribution of $X^k = \min(\text{Bin}((D_2^k - C + L^k)^+, \alpha),\ L^k)$. The capacity that is made available to high-fare customers is equal to $L^k + (C - L^k - D_2^k)^+ - X^k$, where $(C - L^k - D_2^k)^+$ is the unused low-fare capacity and $(C - L^k - D_2^k)^+ > 0$ implies that $X^k = 0$.

The sequence of observations are the following: a) $\min(D_2^k, C - L^k)$ low-fare sales; b) X^k buy-ups due to excess low-fare demand, c) $\min(D_1^k, L^k + (C - L^k - D_2^k)^+ - X^k)$ high-fare sales. To illustrate the application of the ME approach in this setting we will work in the simplest setting where the firm knows α but does not know the discrete demand distributions for two fare-classes F_i.

It is relatively easy to show that the optimal protection level L^* is defined through the following condition

$$L^* = \inf\{L \ :\ \mathbb{P}(D_2 + X(L) \leq L) \geq \gamma\}, \tag{26}$$

for $X(L) = \min(\text{Bin}((D_2 - C + L)^+, \alpha),\ L)$. To simplify the subsequent exposition we will assume that the buy-ups will not consume all of the protected capacity for the high-fare class. In this case, $X^k = \text{Bin}((D_2^k - C + L^k)^+, \alpha)$. The firm makes direct observations of high-fare sales that can be embedded in the ME demand estimation procedure proposed in §3 and refined in §5.1. The low-fare demand observations give rise to low-fare sales and high-fare buy-ups. Whenever the low-fare sales equal $C - L^k$, i.e., the capacity made available to be sold at the lower price p_2 is depleted, the firm has to use the observed number of buy-ups given by X^k, which may potentially be equal to zero, to infer probabilistically what are the possible values of $D_2^k \geq C - L^k$ that could have given rise to the observation of $C - L^k$ low-fare sales followed by X^k buy-ups. This is done as follows. Let $b^k = C - L^k$, and $n = D_2^k - b^k$. Then, $X^k \sim \text{Bin}(n,\ \alpha)$. Then,

$$\begin{aligned} \mathbb{P}(D_2^k = d \mid X^k, \alpha) &= \frac{\mathbb{P}(D_2^k = d, X^k, \alpha)}{\mathbb{P}(X^k, \alpha)} \\ &= \frac{\mathbb{P}(D_2^k = d, X^k, \alpha)}{\sum_{d \geq C - L^k} \mathbb{P}(\text{Bin}(d - b^k, \alpha) = X^k)}. \end{aligned}$$

The latter expression can be easily evaluated, and thus the firm can allocate each observation of X^k buy-ups (which may be potentially zero) to different D_2^k levels with appropriate probabilities. The result can produce an empirical distribution for D_2.[4]

[4]The above calculation assumes that α is sufficiently small such $X^k < L^k$. In this case the resulting D_2^k "observations" are never censored. If we allow for the possibility that X^k may consume L^k units of capacity, then one could apply the ME algorithm to unconstrain the sales observations.

We will not offer the proof that the above procedure produces demand forecasts and protection levels that converge to the optimal choice L^*, but just offer a numerical illustration of sample paths of protection levels generated by the algorithm for different values of α. As illustrated by Figure 2, the algorithm seems to be performing well, and as expected the protection levels are non-decreasing in α.

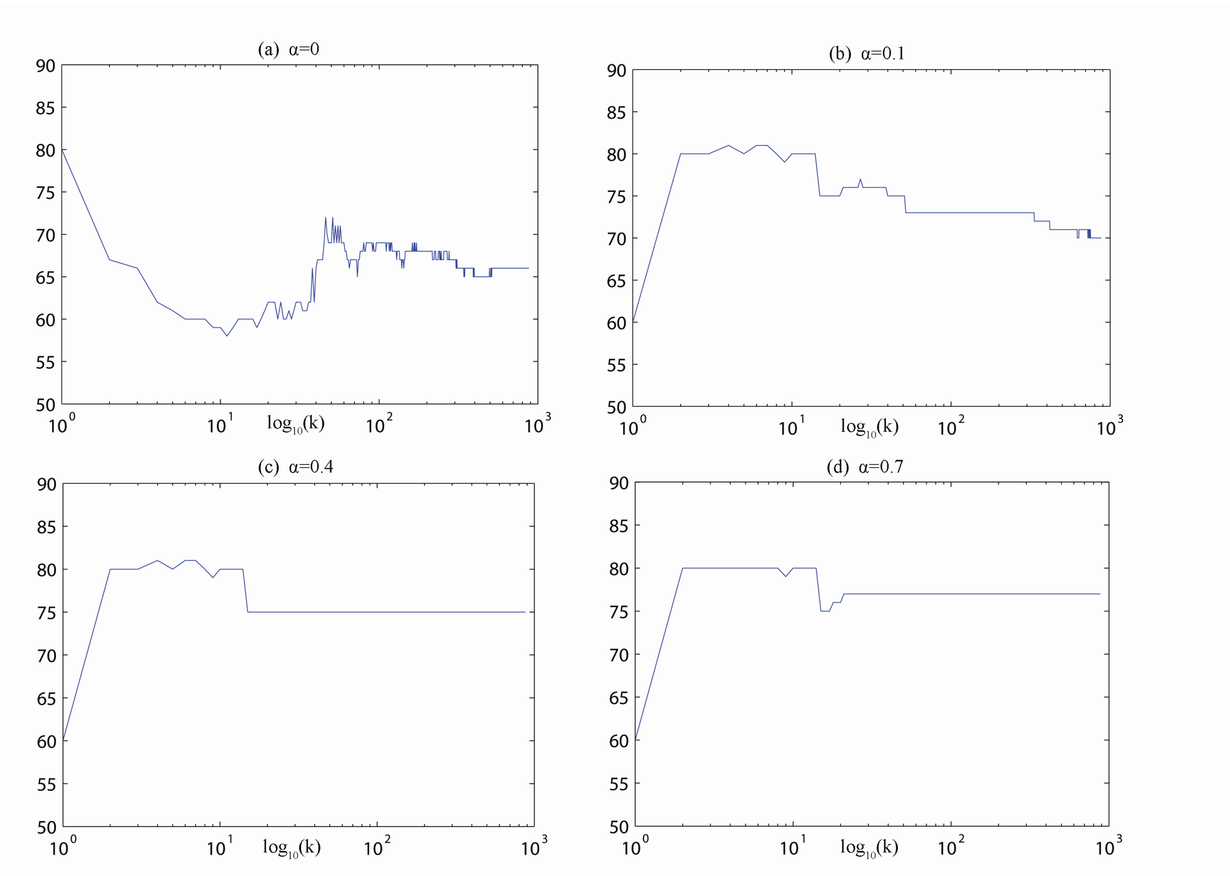

Figure 2: Protection levels produced by the ME algorithm with "buy-up" from low-fare class. $D_1 \sim U[51, 80]$, $D_2 \sim U[101, 200]$, $C = 200$, $\gamma = 0.5$. Buy-up probabilities for each case are given by $\alpha = 0.1$, 0.4 and 0.7 respectively.

6. Concluding remarks

The two main contributions of the paper are the following: first, to demonstrate how Maximum Entropy distributions can offer an intuitive way to unconstrain censored observations of a random variable of interest -in our setting a demand distribution; second, show how ME distributions can be used successfully in the context of forecasting-optimization loops in a way that converges to optimal control decisions even when one starts with no information about the underlying demand distribution.

Specifically, censored information corresponds to fractile information on the demand distribution that can be readily incorporated in the calculation of the ME distribution. Other types of information that could be incorporated in that forecasting step could be upper and lower bounds on the mean of the unknown distribution, information about its second moment, specific information about the probability of specific events, etc. In the context of capacity control of the type studied in this paper, these could be due to side information available to the forecaster or "expert" assessments to be added in the forecast.

A similar approach based on the use ME distributions may be applicable in many other settings. One example arises in the context of estimating a willingness-to-pay distribution to be used in pricing and product design decisions, where the seller may have past sales observations at different price points (fractile information), moment conditions ("expert" assessment), price sensitivity and price elasticity conditions (extracted from limited price experimentation and marketing surveys), etc.[5] Such disparate and partial information is hard to incorporate in many commonly used parametric families of distributions, such as the uniform, exponential, logit, and the normal. In contrast, the ME distribution provides a tractable and intuitive way to incorporate and exploit this information in demand modeling and optimization of pricing and product design decisions. [6]

A. Theorem 2.1 of Kushner and Yin (2003), (Ch. 5, pg. 127)

Kushner and Yin consider an adaptive process $\{\theta^k\}$ on some compact set $H \in \mathbb{R}^n$ that follows the equation $\theta^{k+1} = \theta^k + \epsilon^k Y^k + \epsilon^k Z^k$, where ϵ^k is the step size, $\epsilon^k Z^k$ are correction terms that take the process back to the nearest point in the set H when $\theta^k + \epsilon^k Y^k$ is out of this set, and Y^k are random variables satisfying the conditions given below.

The compact set H and the corresponding correction terms are allowed to take one of several

[5]For example, price sensitivity information at a point j would specify the probability p_j that the willingness-to-pay of a typical customer is equal to j. An elasticity measurement ϵ at j is equivalent to the linear constraint $\epsilon(\sum_{k\geq j} p_k) = p_j \times j$. Inequality constraints on the price sensitivity and/or the elasticity measurements are also easy to incorporate as linear inequality constraints on the probabilities p_j.

[6]It is also possible to formulate ME distribution estimates even when the measurements that the decision maker is considering are noisy, by allowing the constraints in (1) to be violated but striving to keep the degree of violation small.

specific forms in their Theorem 2.1. The $\{\theta^k\}$ process in our problem satisfies a much simpler structure in that it is bounded with probability 1 on a hyperrectangle. Hence, no correction terms $\epsilon^k Z^k$ are necessary, and the simplest specification for set H, which is stated below, is sufficient. Therefore, we adopt and state the theorem below in a simpler form as it applies to our setting, with only the conditions that are required for our specific structure.

- (A.4.3.1) H is a hyperrectangle, i.e., $\exists\ a_i < b_i\ i = 1, \ldots, n$ such that $H = \{\theta\ :\ a_i \leq \theta \leq b_i,\ \forall i\}$.
- (A.5.2.1) $\sup_k \mathbf{E}\,|Y^k|^2 < \infty$
- (A.5.2.2) There is a measurable function $g(\cdot)$ of θ and random variables β^k such that $\mathbf{E}[Y^k \mid \theta^0,\ Y^i,\ i < k] = g(\theta^k) + \beta^k$.
- (A.5.2.3) $g(\cdot)$ is continuous.
- (A.5.2.4) $\sum_i (\epsilon^i)^2 < \infty$
- (A.5.2.5) $\sum_i (\epsilon^i)\ |\beta^i| < \infty$ w.p.1.

Theorem 2.1 of Kushner and Yin (2003) *If $\{\theta^k\}$ is bounded with probability one, then the process converges with probability one to the set of limit trajectories of the mean limit ODE $\dot{\theta}(t) = g(\theta(t))$. If $A_H \subset H$ is a set that is locally asymptotically stable in the sense of Lyapunov for $\dot{\theta}(t) = g(\theta(t))$, and θ^k is in some compact set in the domain of attraction of A_H infinitely often with probability $\geq \rho$, then the process $\theta^k \rightarrow A_H$ with at least probability ρ.*

B. Proofs

Proof that (8)-(12) is equivalent to (13)-(16): Denote the feasible set for problem (8)-(12) as P_1, and similarly, the feasible region for problem (13)-(16) as P_2. We need to show that i) $\forall (p, z) \in P_1,\ p \in P_2$; and ii)$\forall p \in P_2,\ \exists z \in \mathbb{R}_+^{S^2}$ such that $(p, z) \in P_1$.

i) We first show for each $(p, z) \in P_1$, we have $p \in P_2$. Given any $(p, z) \in P_1$, using (9), we get that

$$p_j = \kappa_j^k + \sum_{i \le j} z_{ij} \ge \kappa_j^k = \tilde{\kappa}_j^k, \quad j = 0 \dots S-2, \quad \text{and}$$
$$p_{S-1} = \kappa_{S-1}^k + z_{S-1\,S-1} = \kappa_{S-1}^k + \zeta_{S-1}^k = \eta_{S-1}^k = \tilde{\kappa}_{S-1}^k,$$

hence, p satisfies the first set of constraints (14) in P_2.

Also, using constraints (9) and (10) in P_1, we have that

$$\begin{aligned} \sum_{i \ge j} p_i &= \sum_{i \ge j} \kappa_i^k + \sum_{i \ge j} \sum_{m \le i} z_{mi} = \sum_{i \ge j} \kappa_i^k + \left(\sum_{m<j} \sum_{i \ge j} z_{mi} + \sum_{m \ge j} \sum_{i \ge m} z_{mi} \right) \\ &= \sum_{i \ge j} \kappa_i^k + \left(\sum_{m<j} \sum_{i \ge j} z_{mi} + \sum_{m \ge j} \zeta_m^k \right) \ge \sum_{i \ge j} \eta_i^k, \end{aligned} \tag{27}$$

which shows that p satisfies (15) in P_2. As $\sum_j \; p_j = 1$, the last constraint (16) also obviously holds, and therefore, we have $p \in P_2$.

ii) Next, we show that for all $p \in P_2$, there exists a z such that $(p, z) \in P_1$. Given any $p \in P_2$, define $d_j = p_j - \kappa_j^k$ for all j. Observe $\sum_j d_j = \sum_j p_j - \sum_j \kappa_j^k = 1 - \sum_j \kappa_j^k = \sum_j \zeta_j^k$. Also, note that constraints (14) and (15) imply $\sum_{i \ge j} p_i \ge \sum_{i \ge j} \eta_i^k$ for all j. Therefore, we have that $\sum_{i<j} p_i \le \sum_{i<j} \eta_i^k$, and hence, $\sum_{i<j} d_i \le \sum_{i<j} \zeta_i^k$. Now, let us define a transportation network flow problem as follows: there are S origin nodes each of which has supply ζ_j^k for $j = 0 \dots S-1$, and S destination nodes each of which has demand d_j for $j = 0 \dots S-1$. The variables, z_{ij} denote the flow from origin node i to destination node j for all i, j. We impose an upper bound of zero on flows whenever $i < j$. We minimize the cost $c\, z$ where c is any vector in $R_+^{S \times 2}$. That is we solve the problem

$$\min_z \left\{ c\, z \;\mid\; \sum_{i \le j} z_{ij} = d_j \;\forall j, \quad \sum_{j \ge i} z_{ij} = \zeta_i^k \;\forall i, \quad z_{ij} = 0 \;\forall i < j, \quad z_{ij} \ge 0 \right\} . \tag{28}$$

As $\sum_{i<j} d_i \le \sum_{i<j} \zeta_i^k$, i.e., the cumulative demand is less than the supply and therefore can be met, and as $\sum_j d_j = \sum_j \zeta_j^k$, i.e., the transportation problem is balanced, the above problem is feasible and bounded for all $c \in R_+^{S \times 2}$. For any feasible solution z to the above transportation problem, the corresponding vector $(p, z) \in P_1$ by construction. $\square$

Proof of Lemma 1: We first show $\mathcal{D}(\theta)$ is upper semi-continuous. Denote the universal space of all possible parameters as Θ^U. Consider a generic open set V of the form $V = \{p \mid p_j > \kappa_j - \epsilon_j,\ \forall j,\ \sum_{i \geq j} p_i > \sum_{i \geq j} \kappa_i + \zeta_i - \delta_j \text{ if } \zeta_j > 0\ ,\ \ \sum_j\ p_j = 1\}$, so that $\mathcal{D}(\theta) \subseteq V$ for any $\epsilon_j,\ \delta_j \geq 0$. Now, for any $\epsilon_j,\ \delta_j \geq 0$ and V, define the open set $U = \{p \mid p_j > \kappa_j - \frac{\epsilon_j}{K_1},\ \forall j,\ \sum_{i \geq j} p_i > \sum_{i \geq j} \kappa_i + \zeta_i - \frac{\delta_j}{K_2} \text{ if } \zeta_j > 0\ ,\ \ \sum_j\ p_j = 1\}$, where $K_1,\ K_2 > 1$ are sufficiently large numbers. Then, if $\theta' = [\kappa', \zeta'] \in U \cap \Theta^U$, we have that $\kappa_j' > \kappa_j - \epsilon_j$ and $\kappa_i' + \zeta_i' > \kappa_i + \zeta_i - \delta_j$, which yields $\mathcal{D}(\theta') \subseteq V$. Therefore, $\mathcal{D}(\theta)$ is upper semi-continuous at $\forall \theta \in \Theta^U$.

Next we show that $\mathcal{D}(\theta)$ is also lower semi-continuous. Fix some $\theta = [\kappa, \zeta] \in \Theta^U$, and let V be an open set satisfying $V \cap \mathcal{D}(\theta) \neq \emptyset$, and let $p \in V \cap \mathcal{D}(\theta)$. As V is open, there exists some $\delta > 0$, satisfying $\bar{p} = [\delta,\ 0, \ldots, 0, -\delta] + p \in V$ as well.

Define the "ϵ-neighborhood" of θ as $N_\epsilon(\theta) = \{x \mid ||x - \theta|| < \epsilon\}$, where $|| \cdot ||$ is the L^2 norm. Now by contradiction suppose that there is no neighborhood of θ such that $V \cap \mathcal{D}(\theta') \neq \emptyset$ for all θ' in the neighborhood. Let $\{\epsilon_n\} \to 0$ be a sequence of positive reals, and pick some $\theta^n \in N_{\epsilon_n}(\theta)$ such that $V \cap \mathcal{D}(\theta^n) = \emptyset$. Note that we can find such θ^n by the contradictory assumption. Then, using definitions of V, $\bar{p}$ and $N_{\epsilon_n}(\theta)$, we have that $\bar{p}_j - \kappa_j^n \to p_j - \kappa_j > 0$ and $\sum_{i \geq j} \bar{p}_i - \sum_{i \geq j} \kappa_i^n + \zeta_i^n \to \sum_{i \geq j} \bar{p}_i - \sum_{i \geq j} \kappa_i + \zeta_i > 0$. Consequently, we have that $\bar{p} \in \mathcal{D}(\theta^n)$ for some large n, which yields a contradiction as $\bar{p} \in V$ and $V \cap \mathcal{D}(\theta^n) = \emptyset$. This completes the proof of the lemma. □

Proof of Lemma 2: The optimal solution in problem (13) is $p_\theta \in \mathcal{D}(\theta)$ for any given θ . Note that the objective function $\sum_j p_j \ln p_j$ is strictly convex in p. As shown in Lemma 1, $\mathcal{D}(\theta)$ is a continuous correspondence, which is also easily seen to be convex and compact valued. Then, the result follows from the "The Maximum Theorem under Convexity" (see, e.g., Sundaram (1996), Theorem 9.17.3), which states that under these conditions p_θ is a continuous function in θ.□

Proof of Lemma 3: $f(\theta) = (\mathbf{E}(W^{k+1}), \mathbf{E}(Q^{k+1}))$ satisfies

$$\mathbf{E}(W_j^{k+1}) = \begin{cases} \pi_j & \text{for } \ j < L_\theta \\ q_\theta\ \pi_{L_\theta} & \text{for } \ j = L_\theta \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad \mathbf{E}(Q_j^{k+1}) = \begin{cases} q_\theta\ \sum_{i \geq L_\theta} \pi_i & \text{for } \ j = L_\theta \\ (1 - q_\theta)\ \sum_{i > L_\theta} \pi_i & \text{for } \ j = L_\theta + 1 \\ 0 & \text{otherwise .} \end{cases} \tag{29}$$

By equation (17), the randomization probability q_θ for the ME algorithm satisfies

$$q_\theta = \frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{p_{\theta,L_\theta}}. \tag{30}$$

As p_θ is continuous in θ, q_θ is continuous whenever L_θ does not change value. Note that, in the case the value of L_θ changes, which can do so only by plus or minus one, q_θ changes value between 0 and 1; but, the vectors $\mathbf{E}(W^{k+1})$ and $\mathbf{E}(Q^{k+1})$ are still continuous as their definitions are anchored on the value of L_θ. Hence, $f(\theta)$ and $g(\theta)$ are continuous in θ. $\square$

Proposition 2 *Under assumptions 1 and 2, $\dot{\theta}(t) = g(\theta(t))$ has a unique stationary solution point θ_s which satisfies:*

$$L(\theta_s) = L^*, \quad \textit{and} \quad q(\theta_s) = \frac{(\gamma - \sum_{i=0}^{L^*-1} \pi_i)(\min\{S, S_w\} - L^* + 1)}{1 - \sum_{i=0}^{L^*-1} \pi_i}, \quad \textit{where}$$

$$S_w = \frac{1 - \sum_{i=0}^{L^*-1} \pi_i}{\sqrt{\pi_{L^*}(\gamma - \sum_{i=0}^{L^*-1} \pi_i)}} + L^* - 1.$$

Proposition 3 $\theta_s \in \Theta$.

Proof of Proposition 1: Define $\theta(t) = (\kappa(t), \zeta(t))$ and consider the candidate Lyapunov function $V : \mathbf{R}^{S\times 2} \to \mathbf{R}_+$, where

$$V(\theta(t)) = \sum_{j<L^*} ||\kappa_j(t) - \pi_j|| + (\gamma - \sum_{j<L^*} \pi_j - \kappa_{L^*}(t))^+. \tag{31}$$

It is easy to see that V is continuous everywhere and continuously differentiable almost everywhere with respect to the vector $\theta \in \mathbf{R}^{S\times 2}$. Observe that $\Theta = \{\theta \mid V(\theta) = 0\}$, and that for all $\theta \in \Theta$, we have $L_\theta = L^*$.

In order to establish the stability of the ODE $\dot{\theta}(t) = g(\theta(t))$ it suffices to show that $dV(\theta(t))/dt < 0$ for all t such that $\theta(t) \notin \Theta$ and that $dV(\theta(t))/dt = V(\theta(t)) = 0$ for all t such that $\theta(t) \in \Theta$. In the sequel, we will drop the time argument and denote $dV(\theta(t))/dt$ as $\dot{V}(\theta)$. It suffices to show that $\dot{V}(\theta) = \nabla V(\theta)' g(\theta) < 0$ for all $\theta \notin \Theta$, and that $\dot{V}(\theta) = \nabla V(\theta)' g(\theta) = V(\theta) = 0$ for all $\theta \in \Theta$.

Let $\text{sign}(x) = 1$ if $x > 0$, $\text{sign}(x) = 0$ if $x = 0$ and $\text{sign}(x) = -1$ otherwise. Then, the gradient of the Lyapunov function can be written as follows:

$$\nabla V(\theta) = [(\text{sign}(\kappa_0 - \pi_0), \ldots, \text{sign}(\kappa_{L^*-1} - \pi_{L^*-1}), \ \ -\text{sign}(\gamma - \sum_{j<L^*} \pi_j - \kappa_{L^*}), \ 0, \ldots, 0),$$
$$(0, \ldots, 0)].$$

Alternatively, using the fact that the Lyapunov function contains no ζ_j term, and hence, ignoring the last S components with value 0, with a slight abuse of notation, we can rewrite $\nabla V(\theta)$ and $g(\theta)$ as

$$\nabla V(\theta) = \left(\text{sign}(\kappa_0 - \pi_0), \ldots, \text{sign}(\kappa_{L^*-1} - \pi_{L^*-1}), \ \ -\mathbb{I}(\gamma - \sum_{j<L^*} \pi_j - \kappa_{L^*} > 0), \ 0, \ldots, 0 \right) \tag{32}$$

$$g(\theta) = ((\pi_0 - \kappa_0), \ldots, (\pi_{L_\theta - 1} - \kappa_{L_\theta - 1}), \ (q_\theta \, \pi_{L_\theta} - \kappa_{L_\theta}), \ \ -\kappa_{L_\theta+1}, \ldots, -\kappa_{S-1}) \ .$$

The analysis of the derivative $\dot{V}(\theta)$ will be split in three cases.

Case 1: $L_\theta > L^*$. In this case $\dot{V}(\theta)$ satisfies

$$\dot{V}(\theta) = -\sum_{j<L^*} ||\kappa_j - \pi_j|| + \mathbb{I}(\gamma - \sum_{j<L^*} \pi_j > \kappa_{L^*})(\kappa_{L^*} - \pi_{L^*}) \ .$$

First, observe that $\dot{V}(\theta)$ is less than or equal to 0 since clearly $-\sum_{j<L^*} ||\kappa_j - \pi_j|| \leq 0$, and the term $\mathbb{I}(\gamma - \sum_{j<L^*} \pi_j > \kappa_{L^*})(\kappa_{L^*} - \pi_{L^*})$ is also non-positive as $\gamma - \sum_{j<L^*} \pi_j \leq \pi_{L^*}$ by the definition of L^*. Second, note that $\dot{V}(\theta) = 0$ would imply $\sum_{j \leq L^*} p_{\theta,j} = \sum_{j<L^*} \pi_j + p_{\theta,L_\theta} \geq \sum_{j<L^*} \pi_j + \kappa_{L^*} \geq \gamma$, which contradicts the assumption that $L_\theta > L^*$. Consequently, we have that

$$\dot{V}(\theta) < 0 \quad \text{for } \theta \text{ with } L_\theta > L^* \ . \tag{33}$$

Case 2: $L_\theta < L^*$**.** In this case $\dot{V}(\theta)$ satisfies

$$\begin{aligned}
\dot{V}(\theta) = & - \sum_{j<L_\theta} ||\kappa_j - \pi_j|| + \mathrm{sign}(\kappa_{L_\theta} - \pi_{L_\theta})\,(q_\theta \pi_{L_\theta} - \kappa_{L_\theta}) - \sum_{L_\theta<j<L^*} \mathrm{sign}(\kappa_j - \pi_j)\,\kappa_j \\
& + \mathbb{I}(\gamma - \sum_{j<L^*} \pi_j > \kappa_{L^*})\,\kappa_{L^*} \\
= & \ \Gamma(\theta) - \sum_{L_\theta<j<L^*} \mathrm{sign}(\kappa_j - \pi_j)\,\kappa_j + \mathbb{I}(\gamma - \sum_{j<L^*} \pi_j > \kappa_{L^*})\,\kappa_{L^*} \\
\leq & \ \Gamma(\theta) + \sum_{L_\theta<j<L^*} \pi_j + \mathbb{I}(\gamma - \sum_{j<L^*} \pi_j > \kappa_{L^*})\,\kappa_{L^*} \qquad (34) \\
< & \ \Gamma(\theta) + \sum_{L_\theta<j<L^*} \pi_j + \gamma - \sum_{j\leq L_\theta} \pi_j \quad = \quad \Gamma(\theta) + \gamma - \sum_{j\leq L_\theta} \pi_j \ . \qquad (35)
\end{aligned}$$

where the function $\Gamma(\theta)$ is defined as

$$\Gamma(\theta) := - \sum_{j<L_\theta} ||\kappa_j - \pi_j|| + \mathrm{sign}(\kappa_{L_\theta} - \pi_{L_\theta})\,(q_\theta \pi_{L_\theta} - \kappa_{L_\theta}) \ . \qquad (36)$$

Define the maximum amount of mass that can be reallocated by the ME algorithm at position L_θ for any fixed θ as $p^u(\theta) := \min\{\frac{1-\sum_{i<L_\theta}\kappa_i}{S-L_\theta}, \frac{1}{S}\}$, which is less than π_{L_θ} as by assumption 2 $S > 1/\epsilon > 1/\pi_{L_\theta}$.

Then, we have that $p_{\theta,L_\theta} = \max\{\kappa_{L_\theta},\ p^u(\theta)\}$, and we can rewrite $\Gamma(\theta)$ as

$$\begin{aligned}
\Gamma(\theta) = & - \sum_{j<L_\theta} ||\kappa_j - \pi_j|| + \mathrm{sign}(\kappa_{L_\theta} - \pi_{L_\theta})\,(q_\theta \pi_{L_\theta} - \kappa_{L_\theta}) \\
= & - \sum_{j<L_\theta} ||\kappa_j - \pi_j|| + \mathrm{sign}(\kappa_{L_\theta} - \pi_{L_\theta}) \left(\frac{\gamma - \sum_{i<L_\theta} p_{\theta,i}}{p_{\theta,L_\theta}} \pi_{L_\theta} - \kappa_{L_\theta} \right) \\
= & - \sum_{j<L_\theta} ||\kappa_j - \pi_j|| + \mathbb{I}(\kappa_{L_\theta} \geq \pi_{L_\theta}) \left(\frac{\gamma - \sum_{i<L_\theta} p_{\theta,i}}{\kappa_{L_\theta}} \pi_{L_\theta} - \kappa_{L_\theta} \right) \\
& + \mathbb{I}(p^u(\theta) < \kappa_{L_\theta} < \pi_{L_\theta}) \left(\kappa_{L_\theta} - \frac{\gamma - \sum_{i<L_\theta} p_{\theta,i}}{\kappa_{L_\theta}} \pi_{L_\theta} \right) \\
& + \mathbb{I}(\kappa_{L_\theta} \leq p^u(\theta)) \left(\kappa_{L_\theta} - \frac{\gamma - \sum_{i<L_\theta} p_{\theta,i}}{p^u(\theta)} \pi_{L_\theta} \right) \ .
\end{aligned}$$

Hence, we analyze Case 2 by further conditioning for different values of κ_{L_θ}.

Case 2(a): $L_\theta < L^*$ **and** $\kappa_{L_\theta} \leq p^u(\theta)$**.** Note that, we have $p_{\theta,L_\theta} = p^u(\theta)$ in this case, and also, as $\sum_{i\leq L_\theta} p_{\theta,i} \geq \gamma$, we have $\sum_{i<L_\theta} p_{\theta,i} \geq \gamma - p^u(\theta) > \gamma - \pi_{L_\theta} > \sum_{i<L_\theta} \pi_i$. Now, fix some

constant $M > \sum_{i<L_\theta} \pi_i$, and consider the following parameterized optimization problem

$$\max_{\kappa_j,\ j<L_\theta} \left\{ -\sum_{j<L_\theta} ||\kappa_j - \pi_j|| \ : \quad \sum_{j<L_\theta} p_{\theta,j} = M \right\} \tag{37}$$

which has the solution

$$\kappa_j = \pi_j + \frac{M - \sum_{j<L_\theta} \pi_j}{L_\theta - 1} \text{ with optimal objective value } \sum_{j<L_\theta} \pi_j - M = \sum_{j<L_\theta} \pi_j - \sum_{j<L_\theta} p_{\theta,j}.$$

Therefore, we have for $\kappa_{L_\theta} \leq p^u(\theta)$,

$$\begin{aligned}
\Gamma(\theta) &= -\sum_{j<L_\theta} ||\kappa_j - \pi_j|| + \left(\kappa_{L_\theta} - \frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{p^u(\theta)} \pi_{L_\theta}\right) \\
&\leq \sum_{j<L_\theta} \pi_j - \sum_{j<L_\theta} p_{\theta,j} + \left(\kappa_{L_\theta} - \frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{p^u(\theta)} \pi_{L_\theta}\right) \\
&\leq \sum_{j<L_\theta} \pi_j - \sum_{j<L_\theta} p_{\theta,j} + \left(p^u(\theta) - \frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{p^u(\theta)} \pi_{L_\theta}\right) \\
&\leq \sum_{j<L_\theta} \pi_j - \gamma + p^u(\theta) \\
&< \sum_{j\leq L_\theta} \pi_j - \gamma
\end{aligned} \tag{38}$$

The first inequality follows from the optimization problem (37). The second follows from the fact that the term in parenthesis is increasing in κ_{L_θ} which is less than $p^u(\theta)$ by the case assumption. The third inequality follows from the fact that the right hand side of the inequality is increasing in $\sum_{j<L_\theta} p_{\theta,j} < \gamma$ as $p^u(\theta) < \pi_{L_\theta}$. And, the last inequality follows as $p^u(\theta) < \pi_{L_\theta}$ by assumption 2. As a result, we have that $\dot{V}(\theta) < 0$ for all θ with $L_\theta < L^*$ and $\kappa_{L_\theta} \leq p^u(\theta)$ by combining inequalities (35) and (38) above.

Case 2(b): $L_\theta < L^*$ and $p^u(\theta) < \kappa_{L_\theta} < \pi_{L_\theta}$. Here, we have again $\sum_{i<L_\theta} p_{\theta,i} > \gamma - \pi_{L_\theta} > \sum_{i<L_\theta} \pi_i$. And again, fixing some constant $M > \sum_{i<L_\theta} \pi_i$, and considering the optimization problem (37), we have the same optimal objective value of $\sum_{j<L_\theta} \pi_j - \sum_{j<L_\theta} p_{\theta,j}$. Therefore,

in a similar reasoning, we have the following inequalities in this case

$$\begin{aligned}
\Gamma(\theta) =& -\sum_{j<L_\theta} ||\kappa_j - \pi_j|| + \left(\kappa_{L_\theta} - \frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{\kappa_{L_\theta}}\pi_{L_\theta}\right) \\
\leq& \sum_{j<L_\theta} \pi_j - \sum_{j<L_\theta} p_{\theta,j} + \left(\kappa_{L_\theta} - \frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{\kappa_{L_\theta}}\pi_{L_\theta}\right) \\
<& \sum_{j<L_\theta} \pi_j - \sum_{j<L_\theta} p_{\theta,j} + \left(\pi_{L_\theta} - \frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{\pi_{L_\theta}}\pi_{L_\theta}\right) \\
=& \sum_{j\leq L_\theta} \pi_j - \gamma
\end{aligned} \tag{39}$$

where the strict inequality follows from the fact that the term in parenthesis is increasing in $\kappa_{L_\theta} < \pi_{L_\theta}$ for every fixed value of $\sum_{j<L_\theta} p_{\theta,j}$. As a result, we have that $\dot{V}(\theta) < 0$ for all θ with $L_\theta < L^*$ and $p^u(\theta) < \kappa_{L_\theta} < \pi_{L_\theta}$ by combining inequalities (35) and (39) above.

Case 2(c): $L_\theta < L^*$ **and** $\kappa_{L_\theta} \geq \pi_{L_\theta}$**.** In this case, $\sum_{i<L_\theta} p_{\theta,i} \leq \sum_{i<L_\theta} \pi_i$ is possible. However, the optimal solution to problem (37) remains $\kappa_j = \pi_j + \frac{M - \sum_{j<L_\theta} \pi_j}{L_\theta - 1}, \;\; j < L_\theta$, but the optimal objective value is now $-|| \sum_{j<L_\theta} \pi_j - \sum_{j<L_\theta} p_{\theta,j}||$. Therefore, we have that

$$\begin{aligned}
\Gamma(\theta) =& -\sum_{j<L_\theta} ||\kappa_j - \pi_j|| + \left(\frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{\kappa_{L_\theta}}\pi_{L_\theta} - \kappa_{L_\theta}\right) \\
\leq& -|| \sum_{j<L_\theta} \pi_j - \sum_{j<L_\theta} p_{\theta,j}|| + \left(\frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{\kappa_{L_\theta}}\pi_{L_\theta} - \kappa_{L_\theta}\right)
\end{aligned} \tag{40}$$

Consequently, if $\kappa_{L_\theta} \geq \pi_{L_\theta}$ and $\sum_{i<L_\theta} p_{\theta,i} < \sum_{i<L_\theta} \pi_i$ above, we have that

$$\begin{aligned}
\Gamma(\theta) \leq& \; -|| \sum_{j<L_\theta} \pi_j - \sum_{j<L_\theta} p_{\theta,j}|| + \left(\frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{\kappa_{L_\theta}}\pi_{L_\theta} - \kappa_{L_\theta}\right) \\
=& \sum_{j<L_\theta} p_{\theta,j} - \sum_{j<L_\theta} \pi_j + \left(\frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{\kappa_{L_\theta}}\pi_{L_\theta} - \kappa_{L_\theta}\right) \\
\leq& \sum_{j<L_\theta} p_{\theta,j} - \sum_{j<L_\theta} \pi_j + \left(\frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{\gamma - \sum_{j<L_\theta} p_{\theta,j}}\pi_{L_\theta} - (\gamma - \sum_{j<L_\theta} p_{\theta,j})\right) \\
=& 2\sum_{j<L_\theta} p_{\theta,j} - \sum_{j<L_\theta} \pi_j + \pi_{L_\theta} - \gamma \\
<& \sum_{j\leq L_\theta} \pi_j - \gamma
\end{aligned} \tag{41}$$

where the second inequality follows from the fact that the term in big parenthesis is strictly decreasing in κ_{L_θ} and $\kappa_{L_\theta} = p_{\theta,L_\theta} \geq \gamma - \sum_{j<L_\theta} p_{\theta,j}$ must hold by definition of L_θ in this region. The strict inequality follows from the case assumption $\sum_{i<L_\theta} p_{\theta,i} < \sum_{i<L_\theta} \pi_i$ and assumptions 1 and 2. As a result, we have that $\dot{V}(\theta) < 0$ for all θ with $L_\theta < L^*$, $\kappa_{L_\theta} \geq \pi_{L_\theta}$ and $\sum_{i<L_\theta} p_{\theta,i} < \sum_{i<L_\theta} \pi_i$ by combining inequalities (35) and (41) above.

On the other hand, if $\kappa_{L_\theta} \geq \pi_{L_\theta}$ and $\sum_{i<L_\theta} \pi_i \leq \sum_{i<L_\theta} p_{\theta,i} < \gamma - \pi_{L_\theta}$, we have that

$$
\begin{aligned}
\Gamma(\theta) \leq & -|| \sum_{j<L_\theta} \pi_j - \sum_{j<L_\theta} p_{\theta,j}|| + \left(\frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{\kappa_{L_\theta}} \pi_{L_\theta} - \kappa_{L_\theta}\right) \\
&= \sum_{j<L_\theta} \pi_j - \sum_{j<L_\theta} p_{\theta,j} + \left(\frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{\kappa_{L_\theta}} \pi_{L_\theta} - \kappa_{L_\theta}\right) \\
&\leq \sum_{j<L_\theta} \pi_j - \sum_{j<L_\theta} p_{\theta,j} + \left(\frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{\gamma - \sum_{j<L_\theta} p_{\theta,j}} \pi_{L_\theta} - (\gamma - \sum_{j<L_\theta} p_{\theta,j})\right) \\
&= \sum_{j\leq L_\theta} \pi_j - \gamma
\end{aligned} \tag{42}
$$

where the second inequality follows from the fact that the third term in line two is strictly decreasing in κ_{L_θ}, and that $\kappa_{L_\theta} = p_{\theta,L_\theta} \geq \gamma - \sum_{j<L_\theta} p_{\theta,j} > \pi_{L_\theta}$ must hold by definition of L_θ in this region. As a result, we have that $\dot{V}(\theta) < 0$ for all θ with $L_\theta < L^*$, $\kappa_{L_\theta} \geq \pi_{L_\theta}$ and $\sum_{i<L_\theta} \pi_i \leq \sum_{i<L_\theta} p_{\theta,i} < \gamma - \pi_{L_\theta}$ by combining inequalities (35) and (42) above.

Lastly, if $\kappa_{L_\theta} \geq \pi_{L_\theta}$ and $\sum_{i<L_\theta} p_{\theta,i} \geq \gamma - \pi_{L_\theta}$, we have that

$$
\begin{aligned}
\Gamma(\theta) \leq & -|| \sum_{j<L_\theta} \pi_j - \sum_{j<L_\theta} p_{\theta,j}|| + \left(\frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{\kappa_{L_\theta}} \pi_{L_\theta} - \kappa_{L_\theta}\right) \\
&= \sum_{j<L_\theta} \pi_j - \sum_{j<L_\theta} p_{\theta,j} + \left(\frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{\kappa_{L_\theta}} \pi_{L_\theta} - \kappa_{L_\theta}\right) \\
&\leq \sum_{j<L_\theta} \pi_j - \sum_{j<L_\theta} p_{\theta,j} + \left(\frac{\gamma - \sum_{j<L_\theta} p_{\theta,j}}{\pi_{L_\theta}} \pi_{L_\theta} - \pi_{L_\theta}\right) \\
&\leq \gamma - \sum_{j<L_\theta} p_{\theta,j} \\
&\leq \sum_{j\leq L_\theta} \pi_j - \gamma
\end{aligned} \tag{43}
$$

where the second inequality follows again from the fact that the term in big parenthesis is strictly decreasing in κ_{L_θ} and $\kappa_{L_\theta} = p_{\theta,L_\theta} \geq \gamma - \sum_{j<L_\theta} p_{\theta,j}$ must hold by definition of L_θ in this region. As a result, we have that $\dot{V}(\theta) < 0$ for all θ with $L_\theta < L^*$, $\kappa_{L_\theta} \geq \pi_{L_\theta}$ and

$\sum_{i<L_\theta} \pi_i \leq \sum_{i<L_\theta} p_{\theta,i} < \gamma - \pi_{L_\theta}$ by combining inequalities (35) and (43) above.

We have covered all three cases, 2(a), 2(b), and 2(c) for θ with $L_\theta < L^*$, and showed that $\dot{V}(\theta) < 0$ for such θ. Combining with Case 1 for $L_\theta > L^*$, so far we have shown that $\dot{V}(\theta) < 0$ for θ with $L_\theta \neq L^*$.

Case 3: $L_\theta = L^*$ and $\theta \notin \Theta$. The derivative $\dot{V}(\theta)$ in this case satisfies

$$\dot{V}(\theta) = -\sum_{j<L^*} ||\kappa_j - \pi_j|| - \mathbb{I}(\gamma - \sum_{j<L^*} \pi_j > \kappa_{L^*})(q_\theta\, \pi_{L^*} - \kappa_{L^*}) \ .$$

We study two sub-cases depending on the value of the indicator function.

Case 3(a): $L_\theta = L^*$, $\theta \notin \Theta$, and $\kappa_{L^*} \geq \gamma - \sum_{j<L^*} \pi_j$. Remember that by definition of Θ and $V(\cdot)$, we have that $\Theta = \{\theta \mid V(\theta) = 0\}$. Hence, for $\theta \notin \Theta$ satisfying $L_\theta = L^*$ and $\kappa_{L^*} \geq \gamma - \sum_{j<L^*} \pi_j$, we have that $V(\theta) = \sum_{j<L^*} ||\kappa_j - \pi_j|| > 0$ by construction of the Lyapunov function. Consequently, in this case, we have $\dot{V}(\theta) = -\sum_{j<L^*} ||\kappa_j - \pi_j|| = -V(\theta) < 0$.

Case 3(b): $L_\theta = L^*$, $\theta \notin \Theta$, and $\kappa_{L^*} < \gamma - \sum_{j<L^*} \pi_j$. When $\kappa_{L^*} < \gamma - \sum_{j<L^*} \pi_j$, for $L_\theta = L^*$ to hold, $\theta \notin \Theta$ needs to satisfy $\sum_{j<L_\theta} \kappa_j = \sum_{j<L^*} \kappa_j > \sum_{j<L^*} \pi_j$ by definition of L_θ and L^*. Hence, in this case, we have

$$\begin{aligned}
\dot{V}(\theta) &= -\sum_{j<L^*} ||\kappa_j - \pi_j|| + \kappa_{L^*} - q_\theta\, \pi_{L^*} && (44)\\
&\leq -||\sum_{j<L^*} \kappa_j - \sum_{j<L^*} \pi_j|| + \kappa_{L^*} - q_\theta\, \pi_{L^*} \\
&= -\sum_{j<L^*} \kappa_j + \sum_{j<L^*} \pi_j + \kappa_{L^*} - \frac{\gamma - \sum_{j<L^*} \kappa_j}{p_{L^*}} \pi_{L^*} \\
&\leq -\gamma - \sum_{j<L^*} \pi_j + \kappa_{L^*} \\
&< 0, && (45)
\end{aligned}$$

where the first inequality follows due to the triangle inequality. Under Assumptions 1, 2, and 3, $p_{L^*} < \pi_{L^*}$ holds, as previously discussed above, and consequently, the right hand side of the second equality is increasing in $\sum_{j<L^*} \kappa_j < \gamma$, which yields the second inequality by replacing $\sum_{j<L^*} \kappa_j$ with γ. The strict inequality follows from the case assumption, yielding $\dot{V}(\theta) < 0$ for this case. As a result, we have that $\dot{V}(\theta) < 0$ for $\theta \notin \Theta$ satisfying $L_\theta = L^*$ and $\kappa_{L^*} < \gamma - \sum_{j<L^*} \pi_j$.

We have analyzed all cases for $\theta \notin \Theta$ and shown that $\dot{V}(\theta) < 0$. Also observing that $\dot{V}(\theta) = V(\theta) = 0$ if and only if $\theta \in \Theta$, we conclude that Θ is globally asymptotically stable. $\square$

References

V. F. Araman and R. A. Caldentey. Dynamic pricing for non-perishable products with demand learning. *Working Paper, New York University, New York, NY 10012*, 2005.

Y. Aviv and A. Pazgal. Pricing of short life-cycle products through active learning. *Working Paper, Washington University, St. Louis, MO 63130*, 2005.

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

A. Benveniste, A. M. Mtivier, and P. Priouret. Adaptive algorithms and stochastic approximations. *Springer Verlag, Applications of Mathematics, vol. 22, Berlin, Heidelberg, New York,*, 1990.

D. Bertsimas and A. Thiele. A robust optimization approach to supply chain management. *Proceedings of 14th IPCO*, pages 86–100, 2004.

O. Besbes and A. Muharremoglu. On the price of demand censoring in the newsvendor problem. *Working Paper, Columbia University, New York, NY 10027*, 2010.

O. Besbes and A. Zeevi. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, 2009. Forthcoming.

I. C. Bilegan, C. A. N. Cosenza, S. Gonzlez-Rojo, and F. Mora-Camino. A maximum entropy approach to update airlines demand distributions. *Inverse Problems, Design and Optimization Symposium. Rio de Janerio, Brazil.*, 2004.

E. A. Boyd, E. Kambour, and J. Tama. The impact of buy down on sell up, unconstraining, and spiral down. *Presented at 1st Annual INFORMS Revenue Management Section Conference, New York.*, 2001.

A. N. Burnetas and C. E. Smith. Adaptive ordering and pricing for perishable products. *Operations Research*, 48(3):436–443, 2000.

W. L. Cooper, T. Homem-de Mello, and A. J. Kleywegt. Models of the spiral-down effect in revenue management. *Operations Research*, 54(5):968–987, 2006.

T. Cover and J. Thomas. *Elements of Information Theory.* John Wiley and Sons, 1991.

V. F. Farias and B. Van Roy. Dynamic pricing with a prior on market response. *Working Paper, Stanford University, Stanford, CA 94305*, 2006.

G. Gallego and I. Moon. A min-max distribution newsboy problem: Review and extensions. *Journal of Operational Research Society*, 44:825–834, 1993.

G. A. Godfrey and W. B. Powell. An adaptive, distribution-free algorithm for the newsvendor problem with censored demands, with applications to invntory and distributions. *Management Science*, 47:1101–1112, 2001.

T. Huh and P. Rusmevichientong. A non-parametric asymptotic analysis of inventory planning with censored demand. 2008. Technical Report, Columbia University.

T. Huh, R. Levi, P. Rusmevichientong, and J. B. Orlin. Adaptive data-driven inventory control policies based on Kaplan-Meier estimator. *Operations Research*, 2009. Forthcoming.

E. T. Jaynes. On the rationale of maximum-entropy methods. *Proceedings of the IEEE*, 70(9): 939–952, 1982.

H. K. Khalil. *Nonlinear Systems.* Prentice-Hall, Englewood Cliffs, NJ, 2nd edition, 1996.

S. Kunnumkal and H. Topaloglu. A stochastic approximation method for the single-leg revenue management problem with discrete demand distributions. *Math. Oper. Res.*, 2009. Forthcoming.

H. J. Kushner and G. Yin. Stochastic approximation and recursive algorithms and applications. 2nd ed. *Springer New York*, 2003.

M. S. Lobo and S. Boyd. Pricing and learning with uncertain demand. *Working Paper, Duke University, Durham, NC 27708*, 2003.

G. Perakis and G. Roels. The distribution-free newsvendor: inventory management with limited demand information. *Operations Research*, 2006. Forthcoming.

Rangarajan K. Sundaram. *A First Course in Optimization Theory.* Cambridge University Press, 1996.

G. van Ryzin and J. McGill. Revenue management without forecasting or optimization: An adaptive algorithm for determining airline seat protection levels. *Management Science*, 46(6): 760–775, 2000.

L.R. Weatherford. Better unconstraining of airline demand data in revenue management systems for improved forecast accuracy and greater revenues. *Journal of Revenue and Pricing Management*, 1:234–254(21), 2002.