A maximum entropy joint demand estimation and capacity control policy

Costis Maglaras
Graduate School of Business
Columbia University
Tel: 212-854-4240; Fax: 212-316-9180
email: c.maglaras@gsb.columbia.edu

Serkan Eren
email: serkanseren@gmail.com

Init. Ver.: 04/11/2009; Revisions: 10/21/2010; 07/10/2011; 04/11/2014 (accepted)

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; this is also known as the repeated newsvendor problem. 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, maximum entropy distributions.

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). This paper proposes 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 many 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 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 an excellent treatment of information theory, entropy, 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)[1].

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:

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

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

[1]See Jaynes (1982) for a lengthy discussion of why ME distributions come up naturally in many estimation problems and natural phenomena. In short, Jaynes argue that by maximizing entropy, one chooses the frequencies that can produce the observed data in the greatest number of ways. In other words, the ME distribution is the most "likely" (in the sense defined by Jaynes (1982)) one that is consistent with known data and constraints. Furthermore, the Entropy Concentration Theorem stated by Jaynes yields a sharp constraint on the entropy of possible outcomes for a random trial under linearly independent constraints.

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.

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 $\mathbf{1}(\cdot)$ is the indicator function which is equal to 1 when the condition is true, and is 0, otherwise. For continuous distributions the summations 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). 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, 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 incorporate the information from censored observations in the form of conditions of fractile probabilities of the cumulative demand distribution – which is precisely what these observations tell us about the demand. 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 repeats. The proposed policy is "passive" or "myopic" in that it does not actively selects its controls to both extract high revenues and learn the unknown demand distribution, but the latter is nevertheless achieved through the fitting of the ME forecast. This is motivated from practical considerations and also leads to a tractable analysis.

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, the 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 response, the ME demand forecast will reallocate some of the censored sales observation to higher demand points, shifting 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, which is known to be irrelevant for the optimal protection level, is ample. We briefly discuss in the concluding section how the results extend to the case of a general low-fare demand distribution. Finally, 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.

Another phenomenon that can lead to the spiral down effect is model misspecification, where in broad terms the mechanism that takes primitive demand as input and generates sales observations is not captured or modeled accurately. A simple example is when low-fare demand could buy-up and purchase high-fare tickets when the low fare capacity is depleted, but this effect is not captured in the process of building demand forecasts. Similar to the effect of demand censoring, other aspects of model misspecification should be incorporated into the forecasting process so as to be plausibly captured in the forecasting procedure. While we will not address such issues in this paper, the ME approach is sufficiently flexible to potentially incorporate and address such issues. Specifically, the buy-up phenomenon mentioned above can be modeled and accurately accounted for in the ME demand estimation procedure, albeit with an increase in complexity.

Literature survey: This paper is directly motivated by the observation of the spiral-down effect in Boyd et al. (2001) and its analysis in Cooper et al. (2006). Another paper by Weatherford and S. (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 (2009). 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 be of more general interest.

The closest paper to our work is the one by Huh et al. (2011) 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 the structure of information that can be incorporated, such as additional information about the moments of the distribution, and their potential applications.

A more recent paper by Besbes and Muharremoglu (2013) study the worse case regret in the repeated newsvendor problem (like ours) due to censored demand observations, and provide asymptotic lower bounds on the performance loss attributed to censoring. They emphasize the difference between continuous and discrete demand distributions, showing that the latter is more sensitive to the issue of censoring and to the tension between revenue maximization and demand learning. The distinction between continuous and discrete demand distributions is one of overall scale and granularity, e.g., is the demand measured in 100s and 1000s or in 10s – in the former case an error of 1 or 2 units is insignificant, whereas in the latter it may have an important impact on performance. Our model is motivated by the high-fare class in an airline revenue management problem and typically ranges in the 10s to 100, so it is best modeled with a discrete demand distribution.

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 (2009), Farias and Van Roy (2010), and the references therein. Assuming a parametric family of distributions for the unknown demand runs the risk of model mis-specification. 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

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

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 and partial sales were allowed, then the optimal protection-level L^* would be given by the following equality

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

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

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

The optimal protection level is given by

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

We will make the simplifying assumption that the last inequality in (5) is strict – this is a mild restriction given that F_1 and γ are real-valued. Specifically,

Assumption 1 *$F_1(L^*) > \gamma$.*

Equation (5) highlights an additional challenge for the discrete version of the problem: 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, in order to identify optimal level L^*, either one needs to sample sufficiently often 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. Instead, when the γ-fractile is estimated to be at position L, we will set the protection level to be $L+1$. This will allow us to draw uncensored observations at level L, which will allow us to verify the optimality of L^* upon convergence; while the estimated fractile position is below the optimal level L^*, the experimentation will help accelerate learning. This is a natural assumption and the experimentation at $L+1$ is occasionally adapted by airlines that open up their protection levels to better sample the underlying demand distribution.

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 affects the sequence of observations, and thus the resulting estimation output, which 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 2 $D_2^k \geq C$ *with probability* 1 *for all* k.

Assumption 3 *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 affect 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 simplifies notation and exposition. We briefly discuss how the proposed approach can handle a setting where this assumption is relaxed in the concluding remarks. 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.

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 $L^{k+1} : \mathcal{I}^k \to [0, S-1]$ 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+1} 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 $i \le j$. The ME distribution that corresponds to the observation vector θ^k is computed as follows:

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

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

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

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

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

The last constrained is redundant since $\sum_j (\kappa_j + \zeta_j) = 1$. The above formulation can be simplified to the following problem (all proofs are given in the Appendix A):

Proposition 1 *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 $\tilde{\kappa}_{S-1}^k = \eta_{S-1}^k$. Formulation (8)-(12) can be reduced to:*

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

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

$$\sum_{i \ge j} p_i \ge \sum_{i \ge j} \eta_i^k \qquad if \;\; \zeta_j^k > 0 \tag{15}$$

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

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. Let $L_{\theta^k} = \min\{L \mid F_{\theta^k}(L) \ge \gamma\}$.
4. Implement $L^{k+1} = L_{\theta^k} + 1$.
5. Observe new sales in period $k+1$ and go to step 1.

The "passive" or "myopic" structure of the policy is reflected in steps 3 and 4 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 γ-fractile estimates L_{θ^k} that converges to the optimal level L^*. Notice that under our assumption that we cannot distinguish whether a sales observations equal to the protection level is censored or not, the actual controls $L^{k+1} = L_{\theta^k} + 1$ will converge to $L^* + 1$. That is, the impact of not being able to classify observations at the protection level as censored or not, implies that the control overshoots the optimal level by 1 unit. Once learning (and convergence) has been achieved, however, this could be corrected. 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 at or 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

$$(K^{k+1}, J^{k+1}) = (K^k, J^k) + (W^{k+1}, Q^{k+1}) \tag{17}$$

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

$$\mathbf{P_k}(W_j^{k+1} = 1) = \begin{cases} \pi_j & \text{for } j < L^{k+1} \\ 0 & \text{otherwise,} \end{cases} \tag{18}$$

and,

$$\mathbf{P_k}(Q_j^{k+1} = 1) = \begin{cases} \sum_{i \geq L^{k+1}} \pi_i & \text{for } j = L^{k+1} \\ 0 & \text{otherwise,} \end{cases} \tag{19}$$

for $j = 0 \dots S-1$, where $\mathbf{P_k}(\cdot)$ denotes conditional probability given the information up to iteration-k. In other words, vectors W^{k+1} and Q^{k+1} track the realization of uncensored and censored observations at step $k+1$. Conditional on the control L^{k+1}, these Bernoulli random vectors are independent over k. Note that $\sum_j (K_j^k + W_j^k) = k$, i.e., the total number of censored and uncensored observations is equal to k. Recall that $\theta^k = (K^k, J^k)/k$ and note that $\sum_j (\kappa_j^k + \zeta_j^k) = 1$; i.e., θ^k can be interpreted as probability mass, where we distinguish between uncensored (i.e., correctly identified mass) in κ^k and censored mass in ζ^k, which we will re-allocate using the ME approach. A result that we will use in the subsequent analysis is the following:

Lemma 1 *For $\theta \in \mathbf{R}_+^{2S}$ and $\|\theta\| = 1$, let p_θ denote the maximum entropy distribution computed through (13)-(16). Then, p_θ, is continuous in θ.*

4.2 Convergence of the protection levels

The proof of our main result will study the sequence $\{\theta^k,\ k = 1, \dots\}$, where $\theta_k = [\kappa^k, \zeta^k] \in \mathbf{R}_+^{2S}$. By construction, $\|\theta^k\|_1 = 1$ for $k = 1, \dots$, and therefore since $\{\theta^k, k = 1, \dots\}$ is bounded and in $\mathbf{R}_+^{2S}$ it follows from the Bolzano-Weierstrass Theorem that it has a converging subsequence. Our main result shows that all converging subsequences have the same limit with probability 1, and that the corresponding limiting γ-fractile converges to L^*.

The gist of the proof is the following: 1) for k large, θ^k will be close to its limit $\bar{\theta}$, which may depend on the subsequence; 2) due to the continuity of the maximum entropy distribution with respect to θ, this implies that the ME distribution F_{θ^k} gets close to $F_{\bar{\theta}}$ and as a result the estimated γ-fractile settles down to its limiting value after some finite point in that subsequence; 3) using the last observation, we can apply the strong law of large numbers to prove almost sure limits for the uncensored and censored probability vectors that have a simple and intuitive structure (Lemma 2); 4) finally, we study the ME optimization problem under that limiting structure and derive its solution in closed form (Lemma 3), from which we can deduce that the limiting γ-fractile has to be equal to L^*. A small technical issue that could complicate 2) above is if the γ-fractile of the limiting $\bar{\theta}$ (which we cannot characterize a priori, but ultimately we can establish its structure through a contradiction argument) is such that $F_{\bar{\theta}}(\bar{L}) = \gamma$, i.e., achieves the target fractile with equality. In that case, the γ-fractile along the converging subsequence may oscillate between $\bar{L}$ and $\bar{L}+1$. This is handled in Lemma 4.

Corollary 1 follows from steps 1 and 3, and establishes that all converging subsequences achieve the same limit $\bar{\theta}$, which was explicitly characterized in Lemma 2. In the sequel we let $\Pi_i^j = \sum_{l=i}^{j} \pi_l$ for any $0 \leq i < j \leq S-1$.

Specifically, let $\{k_1, k_2, \ldots\}$ denote any converging subsequence for which $\theta^{k_i} \rightarrow \bar{\theta}$ for some $\bar{\theta} \in \mathbf{R}_+^{2S}$ and $\|\bar{\theta}\|_1 = 1$ that may depend on the specific subsequence. Denote by $F_{\theta^{k_i}}$ the maximum entropy distribution associated with θ^{k_i} defined through (8)-(12), and denote by $L_{\theta^{k_i}}$ the associated γ-fractile defined through (6). Let $p_j^{k_i}$ denote the probability mass points for the maximum entropy distribution for $j = 0, \ldots, S-1$.

Theorem 1 *Under Assumptions 1, 2 and 3,* $\lim_{k_i \rightarrow \infty} L_{\theta^{k_i}} = L^*$ *almost surely.*

Given that the protection control is $L^{k_i+1} = L_{\theta^{k_i}} + 1$, the controls will converge to $L^* + 1$ almost surely. The proof of the Theorem 1 yields the following corollary that states that the ME algorithm ultimately learns the correct discrete demand distribution up to level L^*.

Corollary 1 *Under Assumptions 1, 2 and 3,* $p_j^{k_i} \rightarrow \pi_j$ *for all* $j \leq L^*$ *almost surely.*

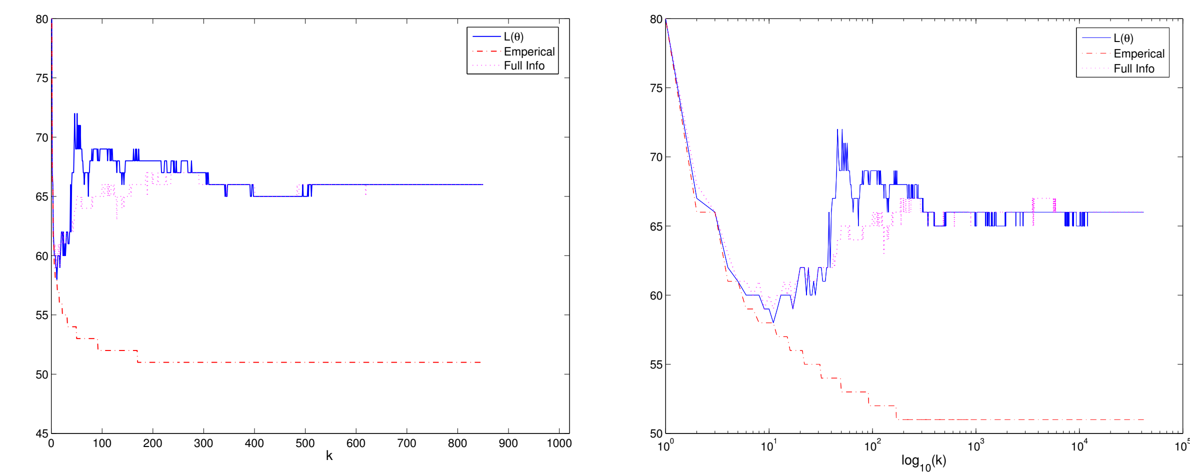

(a) Linear scale (b) Logarithmic scale

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, $p_2/p_1 = 1 - \gamma = 0.5$, $S = 200$, $D_1 \sim U[50, 80]$.

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. In the depicted sample path, it appears that the sample path derived when the true demand (uncensored) is observed under-protects for the high-fare class demand viz the protection levels $\{L^k\}$ under the ME algorithm. This is not true along the entire sample path, however, and indeed the structure of the sample paths can differ depending on the demand realization and its relation to the optimal protection level L^* and the support S assumed by the uncensoring algorithm.

5. 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. The key observation is that censored information corresponds to fractile information on the demand distribution, which, in turn, that can be readily incorporated in the calculation of the ME distribution.

One 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, and the proof provided in the appendix carries through with minor modifications. A somewhat more involved, yet possible, extension would allow low-fare customers to “buy-up” and purchase the high-fare product, say with probability α.

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.[2] 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. [3]

A. Proofs

Proof of Proposition 1: Denote the feasible set for problem (8) defined by constraints (9), (10), (11), and (12) as P_1; and similarly, the feasible region for problem (13) defined by constraints (14), (15), and (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.

[2]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 $\epsilon(\sum_{k \ge j} p_k) = p_j \times j$. 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.

[3]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.

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

$$\sum_{i\geq j} p_i = \sum_{i\geq j} \kappa_i^k + \sum_{i\geq j}\sum_{m\leq i} z_{mi} = \sum_{i\geq j} \kappa_i^k + \left(\sum_{m<j}\sum_{i\geq j} z_{mi} + \sum_{m\geq j}\sum_{i\geq m} z_{mi}\right)$$

$$= \sum_{i\geq j} \kappa_i^k + \left(\sum_{m<j}\sum_{i\geq j} z_{mi} + \sum_{m\geq j} \zeta_m^k\right) \geq \sum_{i\geq j} \eta_i^k, \tag{20}$$

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 $\sum_{i\geq j} p_i \geq \sum_{i\geq j} \eta_i^k$ for all j. Therefore, we have that $\sum_{i<j} p_i \leq \sum_{i<j} \eta_i^k$, and hence, $\sum_{i<j} d_i \leq \sum_{i<j} \zeta_i^k$. 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 $j = 0 \ldots S-1$, and S destination nodes each of which has demand d_j for $j = 0 \ldots S-1$. 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

$$\min_z \left\{ c\,z \;\mid\; \sum_{i\leq j} z_{ij} = d_j \;\forall j, \quad \sum_{j\geq i} z_{ij} = \zeta_i^k \;\forall i, \quad z_{ij} = 0 \;\forall i < j, \quad z_{ij} \geq 0 \right\}. \tag{21}$$

As $\sum_{i<j} d_i \leq \sum_{i<j} \zeta_i^k$, 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. □

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 γ-fractile produced by the ME algorithm is $L_\theta := \min\{L \mid F_\theta(L) \geq \gamma\}$.

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 shows that $\mathcal{D}(\theta)$ is a continuous correspondence.

Now, 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. $\mathcal{D}(\theta)$ 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. $\square$

The proof of the main theorem will study the sequence $\{\theta^k,\ k = 1, \ldots\}$, where $\theta_k = [\kappa^k, \zeta^k] \in \mathbf{R}_+^{2S}$. By construction, $\|\theta^k\|_1 = 1$ for $k = 1, \ldots$, and therefore since $\{\theta^k, k = 1, \ldots\}$ is bounded and in $\mathbf{R}_+^{2S}$ it follows from the Bolzano-Weierstrass Theorem that it has a converging subsequence. Consider any such converging subsequence $\{k_1, k_2, \ldots\}$ for which $\theta^{k_i} \to \bar{\theta}$ for some $\bar{\theta} \in \mathbf{R}_+^{2S}$ and $\|\bar{\theta}\|_1 = 1$ that may depend on the specific subsequence. The proof will first show that along any such converging subsequence and almost all sample paths, the limiting vector $\bar{\theta}$ has a very simple and intuitive structure (Lemma 2). Second, we will exploit that structural result to write in closed form the maximum entropy distribution, $\bar{p}$ associated with the limiting vector $\bar{\theta}$ (Lemma 3). Finally, we will argue by contradiction and show that the protection level $\bar{L}$ associated with $\bar{p}$ is equal to L^* along any such converging subsequence, establishing the desired convergence result. Corollary

1 follows from steps 1 and 3, and establishes that all converging subsequences achieve the same limit $\bar{\theta}$, which was explicitly characterized in Lemma 2. In the sequel we let $\Pi_i^j = \sum_{l=i}^j \pi_l$ for any $0 \le i < j \le S-1$.

Lemma 2 *Consider the bounded sequence $\{\theta^k\}$ in $\mathbf{R}_+^{2S}$, and pick any converging subsequence $\{k_1, k_2, \ldots\}$ for which $\theta^{k_i} \to \bar{\theta}$ for some $\bar{\theta} \in \mathbf{R}_+^{2S}$ and $\|\bar{\theta}\|_1 = 1$ that may depend on the specific subsequence. Let $\bar{p}$ denote the maximum entropy distribution associated with $\bar{\theta}$ defined in (8)-(12), $F_{\bar{\theta}}$ denote the corresponding cumulative distribution, and define $\bar{L}$ through (6). Assume that $F_{\bar{\theta}}(\bar{L}) > \gamma$. Then, with probability 1, $\bar{\theta} = [\bar{\kappa}, \bar{\zeta}]$ is of the following form:*

$$\bar{\kappa}_l = \pi_l \quad \textit{for } l \le \bar{L}, \quad \bar{\kappa}_j = 0 \quad \textit{for } l > \bar{L}; \tag{22}$$

$$\bar{\zeta}_{\bar{L}+1} = \Pi_{\bar{L}+1}^{S-1}, \quad \bar{\zeta}_l = 0 \quad \textit{for } l \le \bar{L} \textit{ or } l > \bar{L}+1. \tag{23}$$

Proof: Let $\nu = \max(\gamma - F(\bar{L}-1), F(\bar{L}) - \gamma) > 0$. By Lemma 1 it follows that there exists a $\nu' > 0$ such that $\|\theta^{k_i} - \bar{\theta}\|_1 < \nu'$ implies that $|F_{\theta^{k_i}}(j) - F_{\bar{\theta}}(j)| < \nu$ for $j = 0, \ldots, S-1$. Given that $\theta^{k_i} \to \bar{\theta}$ it follows that there exists an $M > 0$ such that for $k_i > M$, $\|\theta^{k_i} - \bar{\theta}\|_1 < \nu'$ along that subsequence; M may itself depend on the subsequence. Let $F_{\theta^{k_i}}$ denote the cumulative maximum entropy distribution associated with θ^{k_i} defined through (8)-(12) and L^{k_i} denote the corresponding γ-fractile defined through (6). By the continuity of the maximum entropy distribution on θ and the definition of M, it follows that for all $k_i > M$, $L^{k_i} = \bar{L}$, and, therefore, for all such k_i the protection level was equal to $\bar{L}+1$.

Consider the subsequence for $k_i > M$. Intuitively, the protection level is constant from there onwards, and censored and uncensored observations in each position $j = 0, \ldots, S-1$ become independent Bernoulli random variables with success probabilities as specified in (18)-(19). Applying the strong law of large numbers (SLLN) will give desired almost sure convergence result. In more detail, for $j = 0, \ldots, \bar{L}$, we have that

$$\bar{\kappa}_j = \lim_{i \to \infty} \frac{1}{i} \left(W_j^M + \sum_{l > M}^{k_i} \chi_l(\pi_j) \right),$$

where $\chi_l(\pi_j)$ are Bernoulli random variables with success probability π_j, independent across l (and, in fact, independent of Bernoulli random variables modeling the observations at different values of j). The SLLN for the i.i.d random variables χ_l implies that $\bar{\kappa}_j = \pi_j$ almost surely. For $j > \bar{L}$, the success probabilities for uncensored observations are zero. A similar argument applies to the vector of censored observations that occur only at position $j = \bar{L}+1$, and establishes the limit $\bar{\zeta}$. $\square$

Lemma 3 *Consider any $\bar{\theta} = [\bar{\kappa}, \bar{\zeta}]$ that satisfies (22)-(23). Let $\bar{p}$ denote the maximum entropy distribution associated with $\bar{\theta}$ defined in (8)-(12) and $\bar{L}$ the corresponding γ-fractile defined in (6). Assume that $F_{\bar{\theta}}(\bar{L}) > \gamma$. Then,*

$$\bar{p}_j = \pi_j \quad \text{for } j = 0, \ldots, \bar{L} \quad \text{and} \quad \pi_j = \frac{\bar{\zeta}_{\bar{L}+1}}{S - \bar{L} - 1} = \frac{\Pi_{\bar{L}+1}^{S-1}}{S - \bar{L} - 1} \quad \text{for } j = \bar{L}+1, \ldots, S-1. \tag{24}$$

Proof: Since $\bar{\zeta}_j = 0$ for all $j \le \bar{L}$ it follows that $\bar{p}_j = \bar{\kappa}_j = \pi_j$ for $j \le \bar{L}$; that is, these entries of the distribution are fixed and cannot be changed via the maximum entropy optimization problem that can only select the probability masses at $j = \bar{L}+1, \ldots, S-1$. The resulting problem becomes: choose $x_j \ge 0$ for $j = \bar{L}+1, \ldots, S-1$ to

$$\max_x \left\{ -\sum_{j > \bar{L}} x_j \ln x_j \ : \ x \ge 0, \quad \sum_{j > \bar{L}} x_j = \Pi_{\bar{L}+1}^{S-1} \right\}.$$

As mentioned in the introduction, a standard argument based on the Lagrangian of the above problem, gives that the maximum entropy distribution from $\bar{L}+1$ onwards is uniform, which proves the Lemma. $\square$

The above two Lemmas studied the structure of the limiting vector $\bar{\theta}$ under the assumption that $F_{\bar{\theta}}(\bar{L}) > \gamma$. Given the continuity of the maximum entropy distribution with respect to θ, for large enough k_i such that θ^{k_i} is close to $\bar{\theta}$, the above condition implies that $L^{k_i} = \bar{L}$. If $\bar{\theta}$ is such that $F_{\bar{\theta}}(\bar{L}) = \gamma$, then as k_i grows large, θ^{k_i} gets close to $\bar{\theta}$ but L^{k_i} may oscillate between $\bar{L}$ (when $F_{\theta^{k_i}}(\bar{L}) \ge \gamma$) and $\bar{L}+1$ (when $F_{\theta^{k_i}}(\bar{L}) < \gamma$, but, necessarily, $F_{\theta^{k_i}}(\bar{L}+1) > F_{\theta^{k_i}}(\bar{L}) + \epsilon > \gamma$; the last inequalities follow from Assumption 3). While it seems mild to assume that $F(L^*) > \gamma$, one does not have control over the converging subsequences to enforce an assumption that $F_{\bar{\theta}}(\bar{L}) > \gamma$. As such the case $F_{\bar{\theta}}(\bar{L}) = \gamma$ needs to be addressed, which is done in the next Lemma that mirrors the results of Lemmas 2 and 3.

Lemma 4 *Consider the bounded sequence $\{\theta^k\}$ in $\mathbf{R}_+^{2S}$, and pick any converging subsequence $\{k_1, k_2, \ldots\}$ for which $\theta^{k_i} \to \bar{\theta}$ for some $\bar{\theta} \in \mathbf{R}_+^{2S}$ and $\|\bar{\theta}\|_1 = 1$ that may depend on the specific subsequence. Let $\bar{p}$ denote the maximum entropy distribution associated with $\bar{\theta}$ defined in (8)-(12), $F_{\bar{\theta}}$ denote the corresponding cumulative distribution, and define $\bar{L}$ through (6). Assume* *that $F_{\bar{\theta}}(\bar{L}) = \gamma$. Then, for some $\eta \in [0,1]$, with probability 1, $\bar{\theta} = [\bar{\kappa}, \bar{\zeta}]$ is of the following form:*

$$\bar{\kappa}_l = \pi_l \quad \textit{for } l \leq \bar{L}, \quad \bar{\kappa}_{\bar{L}+1} = \eta\pi_{\bar{L}+1}, \quad \bar{\kappa}_l = 0 \quad \textit{for } l > \bar{L}+1; \tag{25}$$

$$\bar{\zeta}_{\bar{L}+1} = (1-\eta)\Pi_{\bar{L}+1}^{S-1}, \quad \bar{\zeta}_{\bar{L}+2} = \eta\Pi_{\bar{L}+2}^{S-1} \quad \bar{\zeta}_l = 0 \quad \textit{for } l \leq \bar{L} \textit{ or } l > \bar{L}+2. \tag{26}$$

Let $\beta = \frac{\Pi_{\bar{L}+1}^{S-1}}{S-\bar{L}-1}$. Then, the maximum entropy distribution is given by: $\bar{p}_j = \pi_j$ for $j \leq \bar{L}$, and:

i. if $\beta < \eta\pi_{\bar{L}+1}$, then $\bar{p}_{\bar{L}+1} = \eta\pi_{\bar{L}+1}$ and $\bar{p}_j = \frac{\Pi_{\bar{L}+1}^{S-1} - p_{\bar{L}+1}}{S-\bar{L}-2}$ for $j = \bar{L}+2, \ldots, S-1$.

ii. If $\eta\pi_{\bar{L}+1} < \beta \leq \eta\pi_{\bar{L}+1} + \bar{\zeta}_{\bar{L}+1}$, then $\bar{p}_j = \beta$ for $j = \bar{L}+1, \ldots, S-1$.

iii. if $\beta > \eta\pi_{\bar{L}+1} + \bar{\zeta}_{\bar{L}+1}$, then $\bar{p}_{\bar{L}+1} = \eta\pi_{\bar{L}+1} + \bar{\zeta}_{\bar{L}+1}$ and $\bar{p}_j = \frac{\Pi_{\bar{L}+1}^{S-1} - p_{\bar{L}+1}}{S-\bar{L}-2}$.

Proof: For any $\nu > 0$, by Lemma 1 there exists a $\nu' > 0$ such that $\|\theta^{k_i} - \bar{\theta}\|_1 < \nu'$ implies that $|F_{\theta^{k_i}}(j) - F_{\bar{\theta}}(j)| < \nu$ for $j = 0, \ldots, S-1$. Given that $\theta^{k_i} \to \bar{\theta}$, there exists an $M > 0$ such that for $k_i > M$, $\|\theta^{k_i} - \bar{\theta}\|_1 < \nu'$ along that subsequence. Let

$$\eta = \lim_{i\to\infty} \frac{\sum_{l>M}^{k_i} \mathbf{1}\left(L^l = \bar{L}+1\right)}{i},$$

where $\mathbf{1}(\cdot)$ is the indicator function that is equal to 1 if its argument is true, and is equal to 0 otherwise. From the definition of η it follows that for $k_i > M$, the protection level is $\bar{L}+1$ with probability $1-\eta$ and equal to $\bar{L}+2$ with probability η. The argument of Lemma 2 applies unchanged for $\bar{\kappa}_j$ for $j \neq \bar{L}+1$. For $j = \bar{L}+1$,

$$\bar{\kappa}_{\bar{L}+1} = \lim_{i\to\infty} \frac{1}{i}\left(W_{\bar{L}+1}^M + \sum_{l>M}^{k_i} \chi_l(\pi_{\bar{L}+1})\mathbf{1}\left(L^l = \bar{L}+1\right)\right) = \eta\pi_{\bar{L}+1} \quad w.p.\ 1,$$

where the indicator variables on the critical fractile L^l acts as a thinning process over the summation of the i.i.d. Bernoulli random variables, and the almost sure convergence follows from the SLLN; see (Huh et al., 2011, Lemma 1). A similar argument gives the desired result for $\bar{\zeta}_{\bar{L}+1}$ and $\bar{\zeta}_{\bar{L}+1}$ that counts censored observations at the respective positions when the protection level was equal to $\bar{L}+1$ and $\bar{L}+2$, respectively.

Similarly to Lemma 3, since $\bar{\zeta}_j = 0$ for all $j \leq \bar{L}$ it follows that $\bar{p}_j = \bar{\kappa}_j = \pi_j$ for $j \leq \bar{L}$;. The optimization problem for fitting the maximum entropy distribution reduces to choosing the probability mass points x_j at $j = \bar{L}+1, \ldots, S-1$ as follows:

$$\max_x \left\{ -\sum_{j>\bar{L}} x_j \ln x_j \;:\; x \geq 0, \quad \sum_{j>\bar{L}} x_j = \Pi_{\bar{L}+1}^{S-1}, \quad \eta\pi_{\bar{L}+1} \leq x_{\bar{L}+1} \leq \eta\pi_{\bar{L}+1} + \zeta_{\bar{L}+1} \right\}.$$

The only difference with the problem considered in Lemma 3 is the constraint regarding $x_{\bar{L}+1}$. Let $\beta = \frac{\Pi_{\bar{L}+1}^{S-1}}{S-\bar{L}-1}$. In case ii., the value of β is such that the solution obtained in Lemma 3 satisfies the additional constrain on $x_{\bar{L}+1}$ and as such it is optimal. Cases i. and iii. follow by considering the Lagrangian and the effect of the Lagrange multipliers for the lower and upper bound on $x_{\bar{L}+1}$. □

Proof of Theorem 1: The proof will proceed by contradiction. Consider any converging subsequence and the limiting vector $\bar{\theta}$ and the associated protection level $\bar{L}$.

We start by analyzing the setting where the converging subsequence is such that $F_{\bar{\theta}}(\bar{L}) > \gamma$.

a) $\underline{\bar{L} > L^*}$: If $F_{\bar{\theta}}(\bar{L}) > \gamma$, from Lemmas 2 and 3 it follows that $\bar{p}_j = \pi_j$ for all $j \leq \bar{L}$. Since $\bar{L} > L^*$, it follows that $\bar{p}_j = \pi_j$ for all $j \leq L^*$, which by the definition of L^* it follows that $F_{\bar{\theta}}(L^*-1) < \gamma$ and $F_{\bar{\theta}}(L^*) > \gamma$, which contradicts the definition of $\bar{L}$ and the assumption of a). If $F_{\bar{\theta}}(\bar{L}) = \gamma$, Lemma 4 gives that $\bar{p}_j = \pi_j$ for all $j \leq \bar{L}$, and the same contradiction argument applies.

b) $\underline{\bar{L} < L^*}$: Using the structure of $\bar{p}$, and the fact that $\bar{L} < L^*$, we know that $F_{\bar{\theta}}(\bar{L}-1) = \Pi_0^{\bar{L}-1} < \gamma$ and $F_{\bar{\theta}}(\bar{L}) = \Pi_0^{\bar{L}} \leq \Pi_0^{L^*-1} < \gamma$, which contradicts the definition of $\bar{L}$. The same argument applies to both $F_{\bar{\theta}}(\bar{L}) > \gamma$ or $F_{\bar{\theta}}(\bar{L}) = \gamma$.

Together, these imply that $\bar{L} = L^*$ with probability one along any converging subsequence. □

Proof of Corollary 1: It is easy to verify given the structure of $\bar{p}$ that if $\bar{L} = L^*$, $F_{\bar{\theta}}(\bar{L}-1) = \Pi_0^{L^*-1} < \gamma$ and $F_{\bar{\theta}}(\bar{L}) = \Pi_0^{L^*} > \gamma$. This implies that all converging subsequences have limits $\bar{\theta}$ such that $F_{\bar{\theta}}(\bar{L}) > \gamma$. □

References

V. F. Araman and R. A. Caldentey. Dynamic pricing for non-perishable products with demand learning. *Operations Research*, 57(5):1169–1188, 2009.

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.

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. *Management Science*, 59(6):1407–1424, 2013.

O. Besbes and A. Zeevi. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, 57:1407–1420, 2009.

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. *Operations Research*, 58(1):16–29, 2010.

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. *Math. Oper. Res.*, 34:103–123, 2009.

T. Huh, R. Levi, P. Rusmevichientong, and J. B. Orlin. Adaptive data-driven inventory control policies based on Kaplan-Meier estimator. *Operations Research*, 59:646–664, 2011.

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

S. Kunnumkal and H. Topaloglu. A stochastic approximation method for the single-leg revenue management problem with discrete demand distributions. *Math. Oper. Res.*, pages 477–504, 2009.

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. Regret in the newsvendor model with partial information. *Operations Research*, 56(1):188–203, 2006.

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 and Polt S. 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.