INSTITUTT FOR FORETAKSØKONOMI
DEPARTMENT OF FINANCE AND MANAGEMENT SCIENCE

FOR 14 2011
ISSN: 1500-4066
August 2011

Discussion paper

A maximum entropy approach to the newsvendor problem with partial information

BY
Jonas Andersson, Kurt Jörnsten, Sigrid Lise Nonås, Leif K. Sandal, AND **Jan Ubøe**

Norges
Handelshøyskole
NORWEGIAN SCHOOL OF ECONOMICS

A maximum entropy approach to the newsvendor problem with partial information

Jonas Andersson, Kurt Jörnsten, Sigrid Lise Nonås, Leif Sandal and Jan Ubøe
Norwegian School of Economics and Business Administration,
Helleveien 30, N-5045 Bergen, Norway.

August 26, 2011

Abstract

In this paper, we consider the newsvendor model under partial information, i.e., where the demand distribution D is partly unknown. We focus on the classical case where the retailer only knows the expectation and variance of D. The standard approach is then to determine the order quantity using conservative rules such as minimax regret or Scarf's rule. We compute instead the most likely demand distribution in the sense of maximum entropy. We then compare the performance of the maximum entropy approach with minimax regret and Scarf's rule on large samples of randomly drawn demand distributions. We show that the average performance of the maximum entropy approach is considerably better than either alternative, and more surprisingly, that it is in most cases a better hedge against bad results.

Keywords: newsvendor model, entropy, partial information

1 Introduction

The single-period newsvendor model is a heavily studied tool that has attracted increasing interest in the last two decades. The basic setting is that a retailer wants to order a quantity q from a manufacturer, demand D is a random variable, and the retailer wishes to select an order quantity maximizing the expected profit $\mathrm{E}[\Pi_r(q, D)]$. When the distribution of D is known, this problem is easily solved; see Section 3 below.

In many cases, the retailer only has partial information on D. The typical case we have in mind is when the retailer has, e.g., 10–20 observations of D. While this information can be used to form a qualified opinion on the expectation and possibly the variance of D, it offers little clue as to the shape of the distribution. In such situations, the retailer wishes to hedge against bad results, particularly results with negative profits. Two commonly used rules are Scarf's rule and minimax regret, both discussed in detail later. We study the performance of these rules in comparison with an approach where we calculate the order size using the most likely distribution in the sense of maximum entropy. Initially, we very briefly discuss Scarf's rule to focus on a central and seemingly counterintuitive point in our paper: that is, Scarf's rule is not necessarily the best hedge.

Scarf's rule is the solution of the following maximin problem, i.e., given μ and σ, find an order quantity q such that q is the solution to the problem:

$$\max_{q\geq 0}\ \min_{D\in\mathcal{D}_{\mu,\sigma}}\ \mathrm{E}[\Pi_r(q,D)]. \tag{1}$$

The minimum is taken over all distributions D in the set $\mathcal{D}_{\mu,\sigma}$ of all distributions with $\mathrm{E}[D]=\mu$ and $\mathrm{Var}[D]=\sigma^2$. Scarf's construction is certainly relevant and interesting, but there is an obvious objection. That is, Scarf's rule hedges against the worst cases in $\mathcal{D}_{\mu,\sigma}$, but what if these worst cases are very unlikely? From a practical viewpoint, we consider that a manager should seek strategies with small probabilities for bad results, and from this perspective, Scarf's rule may not be the optimal choice. To be more precise, the performance of Scarf's rule may be challenged where there is a small probability of being close to the worst-case scenarios hedged by Scarf's rule.

For example, if a retailer wishes to sell goods in California, they could potentially be concerned about the worst-case scenario where there is a massive slide in the San Andreas Fault. If this expected event materializes, the retailer will incur huge losses, and to hedge against such losses, the retailer should make an order of zero. In practice, most retailers would be willing to accept the risk, i.e., they have no wish to hedge against extremely unlikely events. On the other hand, they may wish to hedge against events that are rare, but not extremely rare. In such cases, it could be appropriate to compare performance at, say, the 95% or 99% percentiles. We demonstrate that ordering schemes based on maximum entropy outperform Scarf's rule at these percentiles, suggesting that Scarf's rule is not an optimal hedge at these levels. As maximum entropy rules also produce greater expected profits, we suggest that such rules are often superior to Scarf's rule.

While the point made above may appear obvious at first, on closer inspection it is not, as it relies heavily on the assumption that we can discuss the probability of a demand distribution. As there is no universal choice for a probability measure on $\mathcal{D}_{\mu,\sigma}$, any particular choice will have to be subjective, and the results will then depend critically on the particular choice made. Trivially, we could equip $\mathcal{D}_{\mu,\sigma}$ with a probability measure where all mass is concentrated at the distributions where Scarf's rule is optimal, and so also trivially, Scarf's rule would be the best choice under this particular scenario. However, in sampling distributions, we should avoid as a rule sampling schemes that have an obvious bias in favor of any of the selection rules. We believe that the sampling schemes we apply in this paper represent a fair way of assessing the performance of different ordering rules.

The newsvendor model under partial information has previously been studied in many different settings, with regard to both the available information on the demand distribution and different extensions to the classical newsvendor model. Two ordering rules are generally discussed when only limited information on the demand distribution is known: the conservative maximin rule (Scarf's rule), which minimizes the worst-case performance (Scarf (1958), Gallego and Moon (1993), Moon and Choi (1995), Gallego et al. (2001)); and the less conservative minimax regret ordering rule (Savage (1951), Yue et al. (2006), Perakis and Roels (2008)). Ordering schemes based on maximum entropy are also mentioned in the literature (Jaynes (1957, 2003), Eren and Maglaras (2006), Perakis and Roels (2008)).

Perakis and Roels (2008) examine the performance of Scarf's rule and the minimax regret rule where both the mean and the variance are known (as is common in the fashion and sporting goods industries (Gallego and Moon (1993)). In particular, Perakis and Roels (2008) focus on how this approach performs in relation to maximum regret (as compared with Scarf's rule and the newsvendor solution given a normal- or gamma-demand distribution) and mean expected profit loss (compared with the newsvendor solution when the demand distribution is truncated normal, gamma, lognormal, or negative binomial). In their findings, Perakis and Roels (2008) recommend the use of the truncated normal distribution for situations with large coefficients of variation and the gamma distribution for large profit margins.

Perakis and Roels (2008) also briefly discuss the possibility of selecting an order scheme based on the maximum entropy distribution, listing the classical entropy maximizing distributions and showing how the ordering decision using these distributions performs, in terms of maximum regret, compared with Scarf's rule and the minimax regret decision. Eren and Maglaras (2006) suggest a maximum entropy approach similar to that in the current paper. However, while Eren and Maglaras (2006) briefly mention a few relevant formulas, they never enter into a discussion of much depth. Lastly, Lim and Shanthikumar (2007) discuss how we can use relative entropy to measure uncertainty in the demand rate for dynamic revenue management problems.

In this paper, we assume knowledge of the mean and variance of the demand distribution, and use this information to find the particular distribution that maximizes entropy. We then calculate the expected profit or loss using the maximum entropy distribution instead of the real demand distribution. We use this to compare the results with Scarf's rule and the minimax regret ordering rule in terms of average expected loss from the real distribution (mean loss and loss at the 95% and 99% percentiles) using a large sample of randomly drawn demand distributions. The three ordering rules are also compared (in terms of mean expected loss) when the real demand distribution is a mixture of normal and exponential distributions. We find that the average performance of the entropy approach in terms of mean expected loss is considerably better. More surprisingly, the maximum entropy approach also outperforms Scarf's rule (as well as the minimax regret approach) when we consider the 95% and 99% percentiles of the expected loss.

The most recent literature considering the newsvendor model with limited demand information discusses both Scarf's rule (Perakis and Roels (2010), Tajbakhsh et al. (2010), Berman et al. (2011)) and the minimax regret rule in different settings (Jiang et al. (2010), Perakis and Roels (2010), Lan et al. (2011), Lin and Ng (2011)). For instance, Lan et al. (2011) derive optimal overbooking levels and booking limits, minimizing maximum relative regret for an overbooking and fare-class allocation model in revenue management when demand and no-shows are characterized using interval uncertainty. In other work, Berman et al. (2011) compare a centralized (pooled) system with a decentralized (nonpooled) system using a number of common demand distributions, as well as the distribution-free approximation when only the first two moments of the demand distribution are known (Gallego and Moon (1993)). Likewise, Lin and Ng (2011) consider the minimax regret multimarket newsvendor model when demand is only known to be bounded within some given interval. Similarly, Jiang et al. (2010) study minimax regret where we have multiple newsvendors that compete in a setting with asymmetric information, such that each newsvendor knows the support of their own demand distribution but only have estimates of the demand distribution support for their competitors. Finally, Perakis and Roels (2010) consider the maximin and the minimax regret decision criteria for the capacity allocation problem in revenue management under general polyhedral uncertainty sets.

Benzion et al. (2010) and Lee and Hsu (2011) present other related literature. To start with, Lee et al. (2011) study the effect of advertising in the newsvendor problem when only the mean and variance of the demand distribution are known. Benzion et al. (2010) present an experimental study where all participants should assume the newsvendor role but only half have knowledge of the underlying demand distribution. This study assumes that the newsvendor should decide the order quantity each day for 100 days. The findings in Benzion et al. (2010) indicate that the two groups behave differently, but that knowing the demand distribution does not necessarily lead to better decisions. The reason is that the supply surplus in a given period strongly affects the order quantity in the following period for both groups. In the earlier literature, Perakis and Roels (2008) provide a good overview of the different distribution-free approaches as well as the concept of entropy.

The remainder of the paper is organized as follows. In Section 2, we set up the framework, formally state the basic definitions, and specify the central results used later in the paper. In Section 3, we show how to construct maximum entropy distributions, and illustrate how we can use these distributions to specify alternative ordering rules. In Section 4, we compare the maximum entropy approach with Scarf's rule and minimax regret. We compute the average expected loss (compared with the situation where the distribution is fully known) over large samples of randomly drawn discrete distributions. In Section 5, we refine the simplistic framework in Section 4 using mixture distributions. In this, we equip the set of mixture coefficients with a probability measure, and compute the expected loss (expectation taken over the random coefficients and the random values of demand). In Section 6, we compute the optimal ordering rule for a known mixture distribution. This rule is useful as it provides a lower bound on the expected loss and we can then compute how close the other ordering rules come to this lower bound. We also consider the performance of the maximum entropy approach under the worst-case scenario, i.e., the distributions where Scarf's rule is (by definition) the best hedge. In Section 7, we briefly describe how to construct the maximum entropy distribution in cases where the higher-order moments and/or percentiles are known. While it is straightforward to solve maximum entropy problems under such additional information, maximin and minimax problems can become very difficult to solve. This shows that the maximum entropy approach is more flexible. Finally, in Section 8, we offer some concluding remarks.

2 The newsvendor model under partial information

In this section, we offer a brief review of the theory underlying this paper. However, as it is mostly a summary of classical results, we provide only a minimum of detail. For complete proofs and detailed discussion, we refer to Scarf (1958), Gallego and Moon (1993), Perakis and Roels (2008), and Jörnsten et al. (2011).

2.1 The classical newsvendor model

In the classical newsvendor model, a retailer plans to sell a commodity in a market with uncertain demand. The retailer orders a number of units of the commodity from a manufacturer, and expects to sell a sufficient number of these units to make a profit. The manufacturer decides the wholesale price W (which is fixed in this analysis), while the retailer has an exogenously given selling price (revenue) R and decides the order quantity q. Any unsold items can be salvaged at the price S.

2.2 Retailer's profit

The retailer's profit is denoted by $\Pi_r(q, D)$. Profits in the newsvendor model can be rewritten in several different ways. For the analysis in this paper, it is convenient to express everything in terms of the random variable $\min[D, q]$. Using the relation $(q - D)^+ = q - \min[D, q]$, we obtain:

$$\Pi_r(q, D) = R \min[D, q] + S(q - D)^+ - Wq \tag{2}$$

$$= R \min[D, q] + S(q - \min[D, q]) - Wq \tag{3}$$

$$= (R - S) \min[D, q] - (W - S)q. \tag{4}$$

From the expression above, we obtain:

$$\mathrm{E}[\Pi_r] = (R - S)\mathrm{E}[\min[D, q]] - (W - S)q = \frac{1}{R - S}\left(\mathrm{E}[\min[D, q]] - \frac{W - S}{R - S}\right) q. \tag{5}$$

Without loss of generality, we can assume that units are chosen such that $R - S = 1$. If we define:

$$\beta = \frac{W - S}{R - S} = \frac{\text{overage cost}}{\text{underage cost}}, \tag{6}$$

(5) can be simplified to:

$$\mathrm{E}[\Pi_r] = \mathrm{E}[\min[D, q]] - \beta \cdot q. \tag{7}$$

If D is a continuous distribution with cumulative distribution F_D, the order quantity maximizing $\mathrm{E}[\Pi_r]$ is $q = F_D^{-1}(1-\beta)$. If D is a discrete distribution with values $d_1, d_2, \ldots, d_n$ and probabilities $p_1, p_2, \ldots, p_n$, the optimal order quantity can be found as follows: let $1 \leq k \leq n$ be the smallest integer s.t. $\sum_{i=1}^{k} p_i \geq 1 - \beta$. Maximum expected profit in (7) is achieved using $q = d_k$. In the degenerate case where $\sum_{i=1}^{k} p_i = 1 - \beta$, the expected value is constant if $q \in [d_k, d_{k+1}]$. In all other cases, the optimal order quantity is unique.

Note that the distributions materializing the worst-case scenario under Scarf's rule are, in fact, discrete and degenerate; see Gallego and Moon (1993). Hence, all ordering rules suggesting order quantities reasonably close to the ordering in Scarf's rule are also optimal when applied to the worst case distributions. Hence, the hedge in Scarf's rule is more subtle than it at first appears. We discuss these issues in more detail in Section 6.

2.3 Scarf's rule

This rule assumes that the retailer only knows the expected demand $\mu = \mathrm{E}[D]$ and the standard deviation $\sigma = \mathrm{sd}[D]$. Let μ and σ be given and let $\mathcal{D}_{\mu,\sigma}$ denote the collection of all distributions with $\mathrm{E}[D] = \mu, \mathrm{Var}[D] = \sigma^2$. Scarf (1958) considered the following maximin problem:

$$\max_{q \geq 0} \min_{D \in \mathcal{D}_{\mu,\sigma}} \mathrm{E}[\Pi_r(D, q)], \tag{8}$$

and showed that the optimal maximin order quantity is given by:

$$q = \mu + \frac{\sigma}{2} \left(\frac{1 - 2\beta}{\sqrt{\beta(1-\beta)}} \right). \tag{9}$$

Scarf's rule is often modified to the form:

$$q = \begin{cases} \mu + \frac{\sigma}{2} \left(\frac{1-2\beta}{\sqrt{\beta(1-\beta)}} \right) & 0 \leq \beta \leq \frac{\mu^2}{\mu^2+\sigma^2} \\ 0 & \text{otherwise,} \end{cases} \tag{10}$$

in which case the retailer obtains a nonnegative expected profit for all $D \in \mathcal{D}_{\mu,\sigma}$.

2.4 Minimum regret

In their seminal paper on minimax regret, Perakis and Roels (2008) consider the minimum regret order, i.e., the solution to:

$$\min_{q \geq 0} \max_{D \in \mathcal{D}} \max_{z \geq 0} \mathrm{E}[\Pi_r(D, z)] - \mathrm{E}[\Pi_r(D, q)], \tag{11}$$

where the second maximum is taken over some set $\mathcal{D}$ of distributions, typically characterized by moment conditions. Perakis and Roels (2008) consider eight different cases in their paper. Of these, the case most relevant to our analysis is where μ and σ are known, together with a positivity constraint $D \geq 0$. The optimal order quantity q can then be obtained (Theorem 8 in

Perakis and Roels (2008)) by solving the equation:

$$\max \left\{ \max_{\max\{\mu,q\}\le x\le(\mu^2+\sigma^2)/\mu} \left(\frac{\mu}{x} - \beta\right)(x-q), \right. \tag{12}$$

$$\left. \max_{q\le x\le\mu \cup \max\{q,(\mu^2+\sigma^2)/\mu\}\le x\le q+\sqrt{(q-\mu)^2+\sigma^2}} \left(\frac{\sigma^2}{(x-\mu)^2+\sigma^2} - \beta\right)(x-q) \right\} \tag{13}$$

$$= \max_{\max\{0,y-\sqrt{(q-\mu)^2+\sigma^2}\}\le x\le\min\{q,\mu\}} \left(\frac{(x-\mu)^2}{(x-\mu)^2+\sigma^2} - \beta\right)(x-q), \tag{14}$$

w.r.t. q. When μ, σ, β are given, the left-hand side of this equation is a decreasing function of q alone, while the right-hand side is an increasing function of q. The minimum for the maximum of both sides is obtained at the unique value q where the two sides are equal, and q is easily obtained by a numerical method.

3 Maximum entropy

In this section, we discuss how to solve a maximum entropy problem and how we can use the solution of this problem to select an order quantity for the newsvendor problem.

The maximum entropy principle is often used to derive the prior probability distribution. Jaynes (1957, 2003), for example, argues that the maximum entropy distribution, as it is the least informative given the available information, is a good choice for a prior distribution. Classical examples of entropy-maximizing distributions, among others listed in Perakis and Roels (2008), are the uniform distribution, when only the range of the distribution is known; the exponential distribution, when the distribution is known to be nonnegative and has a certain mean; and the most relevant for our study, the normal distribution, when the distribution has known mean and variance.

If the mean and variance are known, Perakis and Roels (2008, 2010) claim that the normal distribution (if not dramatically skewed) performs quite well for a variety of demand distributions. More specifically, Perakis and Roels (2010) suggest that both the newsvendor policy based on the normal distribution and Scarf's rule (a minimax cost policy) perform quite well when the true distribution is bell shaped (gamma) or skewed (gamma and exponential distribution). However, Gallego et al. (2007) caution against the use of the normal distribution when the coefficient of variation is large as excessive orders and large financial losses may occur. To avoid this, Gallego et al. (2007) recommend the use of nonnegative distributions, such as the gamma, negative binomial, or the lognormal, for products with large coefficients of variation. If managers insist on using the normal distribution, Gallego et al. (2007) provide a tight, distribution-free upper bound to limit the order size. Finally, in related work, Silver et al. (1998) argued that the normal distribution should not be used when the coefficient of variation of demand was greater than 0.5.

3.1 Equations for maximum entropy

The entropy of a discrete distribution is defined by:

$$\text{entropy} = \sum_{i=1}^{n} p_i \ln[p_i]. \tag{15}$$

In the continuous case, we have:

$$\text{entropy} = \int_{-\infty}^{\infty} f(x) \ln[f(x)] dx. \tag{16}$$

In a maximum entropy problem, we want to maximize the entropy over all distributions with given moments. In this paper, we discuss the case where μ and σ are given together with the support of the distribution. In the newsvendor problem, it is natural to restrict the support to the positive real axis. If the maximum value N is known, we can consider an $N+1$-variable problem of the form:

$$\max_{p_0,\ldots,p_N} \sum_{i=0}^{N} p_i \ln[p_i], \tag{17}$$

subject to the constraints:

$$\sum_{i=0}^{N} p_i = 1 \qquad \sum_{i=0}^{N} i \cdot p_i = \mu \qquad \sum_{i=0}^{N} i^2 \cdot p_i = \sigma^2 + \mu^2. \tag{18}$$

The constraints are all linear in p, and the first-order condition reads:

$$\ln[p_i] + 1 - \lambda_1 \cdot 1 - \lambda_2 \cdot i - \lambda_3 \cdot i^2 = 0. \tag{19}$$

Hence, it is possible to find constants a, b, c s.t.:

$$p_i = \mathrm{e}^{a+b\cdot i+c\cdot i^2}. \tag{20}$$

We find numerical values for these constants using the three constraints. The simple method is also easily modified to include the case where we have specified the range, i.e., find constants a, b, c such that:

$$\sum_{i=d}^{e} \mathrm{e}^{a+b\cdot i+c\cdot i^2} = 1 \qquad \sum_{i=d}^{e} i \cdot \mathrm{e}^{a+b\cdot i+c\cdot i^2} = \mu \qquad \sum_{i=d}^{e} i^2 \cdot \mathrm{e}^{a+b\cdot i+c\cdot i^2} = \sigma^2 + \mu^2, \tag{21}$$

where d is the minimum and e is the maximum of D. Passing to the limit, we obtain:

$$\int_d^e \mathrm{e}^{a+b\cdot x+c\cdot x^2} dx = 1 \qquad \int_d^e x \cdot \mathrm{e}^{a+b\cdot x+c\cdot x^2} dx = \mu \qquad \int_d^e x^2 \cdot \mathrm{e}^{a+b\cdot x+c\cdot x^2} dx = \sigma^2 + \mu^2, \tag{22}$$

for the continuous case. If D is supported on the whole real axis, it is straightforward to verify that:

$$\mathrm{e}^{a+b\cdot x+c\cdot x^2} = \frac{1}{\sqrt{2\pi}\sigma} \mathrm{e}^{-(x-\mu)^2/(2\sigma^2)}, \tag{23}$$

i.e., the normal distribution is the solution to the maximum entropy problem. If D is supported on the positive real axis, however, this is no longer true. In particular, if $\sigma \geq \mu$, the difference is considerable as no such distribution can be approximately normal. Figure 1 depicts two particular cases. In the first case, $\mu = 100, \sigma = 100$, we can see that the distribution is far from normal. In the second case, $\mu = 100, \sigma = 50$, the solution resembles the normal distribution, but is not exactly equal.

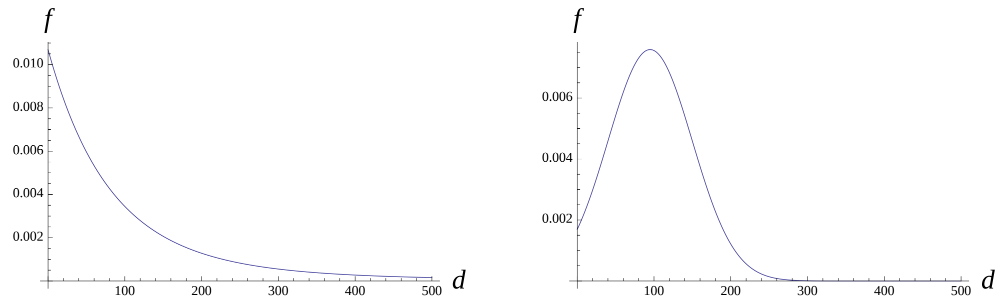

Figure 1: Maximum entropy distributions for the cases

$\mu = 100, \sigma = 100$ (left), $\mu = 100, \sigma = 50$ (right)

The qualitative behavior is governed by the fraction $\frac{\sigma}{\mu}$, and the main cases are described in Table 1.

Table 1: Maximum entropy distributions with positive support

$\frac{\sigma}{\mu} = 0.25$	$\frac{\sigma}{\mu} = 0.50$	$\frac{\sigma}{\mu} = 1$
Small	Medium	Large
Close to normal distribution	Resembling normal distribution	Far from normal distribution

The newsvendor problem is very well understood where D has a normal distribution. Henceforth, we focus on cases where $\frac{\sigma}{\mu}$ is medium/large.

3.2 Predicting order quantities by maximum entropy

The construction above offers an alternative order quantity for the newsvendor model. Subject to the partial information we have on D, we solve a maximum entropy problem to find a distribution $D^{\text{maxentropy}}$ and select the order quantity *we would have used if* $D^{\text{maxentropy}}$ *was the true distribution*, i.e.:

$$q = F^{-1}_{D^{\text{maxentropy}}}(1 - \beta). \tag{24}$$

To examine how this works in a particular case, and to explain how we will evaluate performance later, we assume that the true demand distribution is the discrete distribution to the left in Figure 2.

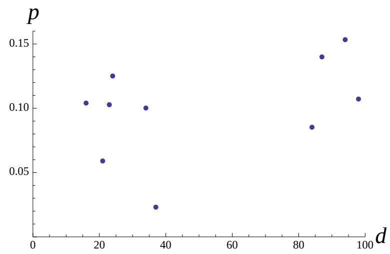

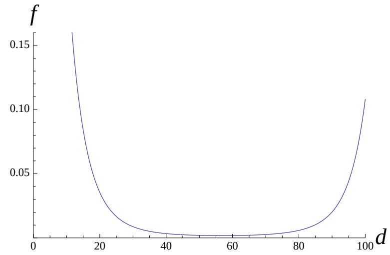

Figure 2: True distribution (left) and maximum entropy distribution (right)

Assume that we have available the following (partial) information about the true D:

$$\mu(D) = 56.8 \qquad \sigma(D) = 33.9 \qquad \min(D) = 16 \qquad \max(D) = 98 \tag{25}$$

If we use this information and solve the nonlinear system in (22), we obtain a solution of the maximum entropy problem that has a density $f_{D^{\text{maxentropy}}}$ given by:

$$f_{D^{\text{maxentropy}}}(x) = \begin{cases} \exp[a + bx + cx^2] & 16 \leq x \leq 98 \\ 0 & \text{otherwise,} \end{cases} \tag{26}$$

where:

$$a = 0.770388 \qquad b = -0.249481 \qquad c = 0.00219509. \tag{27}$$

Figure 2 illustrates a plot of this density to the right. Assuming that $\beta = 0.6$, we can solve (24) to obtain an order quantity of $q = 31.79$. Table 2 shows the corresponding values we obtain using Scarf's rule and minimax regret using the same values.

Table 2: Expected profit using different ordering rules; $\beta = 0.6$

Order	Optimal	Maxentropy	Scarf's rule	Minimax regret
Quantity	34.00	31.79	49.87	49.27
Expected profit	8.58	8.56	6.83	6.90

With full information, the optimal order quantity is 34 with an expected profit of 8.58. From Table 2, we can see that the expected profit using the order quantity from the maximum entropy distribution is very close to the global optimum, while the results from either Scarf's rule or minimax regret are considerably weaker.

We hasten to remark that this is not really a fair comparison. The maximum entropy distribution uses more information (min and max) in this case; hence, the results are not directly comparable. Nevertheless, we consider that the example demonstrates the strength of the maximum entropy approach. The maximum entropy is, to some extent, able to “sense” the U-shape of the true distribution, and that gives some valuable intuition as to why this strategy works so well.

When we evaluate performance in the following section, we compute the maximum entropy distributions for the case with positive support, i.e., where the values of min and max are unknown and hence not used in the calculations. In that case, we are then able to compare performance on an equal basis.

4 Performance evaluation by random sampling

In this section, we evaluate the performance of different ordering rules by sampling from a reasonably wide class of distributions. From each sampled D, we compute the optimal order quantity we would have used under full information, and find the expected profit using that particular order quantity. We then compute $\mu(D)$ and $\sigma(D)$, and use these values to find the order quantities suggested by Scarf's rule, minimax regret, and maximum entropy (assuming D is supported on $[0, \infty)$). Finally, we compute the expected values we would have obtained if these order quantities were used with the true distribution, i.e., the expected values are computed w.r.t. full information on D.

4.1 A simple sampling procedure

As a first step, we sample distributions as follows. We sample $n = 10$ different values uniformly on the interval $[0, 300]$; these values are then sorted to provide 10 values $d_1, d_2, \ldots, d_{10}$. We then sample 10 new values $s_1, s_2, \ldots, s_{10}$ from, e.g., a uniform distribution on the unit interval $[0, 1]$. We do not sort these values, instead normalizing them into probabilities using:

$$p_i = \frac{s_i}{\sum_{j=1}^{10} s_j}. \tag{28}$$

The two sets of numbers define a sampled demand distribution where D has the values $d_1, d_2, \ldots, d_{10}$

with probabilities $p_1, p_2, \dots, p_{10}$. The particular choice of $n = 10$ here is, of course, open for discussion, but if n is too small, we believe that the distributions are too special, whereas if n is very large, we essentially end up with a small perturbation of a uniform distribution. In our opinion, $n = 10$ is a fair trade-off between these two effects, leading to a procedure that is both transparent and easy to implement. Nonetheless, we discuss some more refined constructions in Section 5.

To evaluate the performance of the different ordering rules, we sampled 100 000 distributions for each of the cases $\beta = 0.2, \beta = 0.5, \beta = 0.8$ using the above procedure. For each sample, we compared the expected profit with the profit we would have obtained with full information. Table 3 provides the mean losses in expected value. In Table 3, we have computed the mean expected loss. Alternatively, we could have computed the mean values of the relative losses. We can average relative effects in several different ways, however, and any particular choice will have its own strengths and weaknesses. For example, a very large relative error on a small profit may be misleading in terms of economic importance. For our part, we believe that the mean expected loss (measured directly in terms of monetary units) best reflects the economic impact.

Table 3: Expected profit loss using different ordering rules

$1-\beta = 0.8$	Mean expected profit 66.17		
	Scarf's rule	Minimax regret	Maxentropy
Mean expected loss	1.14	2.13	0.49
(standard deviation)	(1.05)	(1.62)	(0.47)
$1-\beta = 0.5$	Mean expected profit 28.23		
	Scarf's rule	Minimax regret	Maxentropy
Mean expected loss	1.01	0.93	0.72
(standard deviation)	(1.09)	(0.95)	(0.73)
$1-\beta = 0.2$	Mean expected profit 6.13		
	Scarf's rule	Minimax regret	Maxentropy
Mean expected loss	2.55	1.90	0.51
(standard deviation)	(2.20)	(1.50)	(0.50)

The mean expected profit in Table 3 is the mean of the expected profits we would have obtained under full information. This value is relevant here as it makes it possible to assess the relative loss. The results in Table 3 very clearly show that the maximum entropy approach is superior in terms of average loss. The increase in performance is considerable, particularly when β is large. If $\beta = 0.8$, the fraction

$$\frac{\text{Mean expected loss}}{\text{Mean expected profit}} \tag{29}$$

is reduced from 41.6% using Scarf's rule to 8.3% using the maximum entropy approach.

To assess the way these rules hedge against bad profits, we computed the expected loss at the 95% and 99% percentiles. Table 4 details the results.

Table 4: Expected profit loss at 95% and 99% percentiles

$1-\beta = 0.8$	Scarf's rule	Minimax regret	Maxentropy
Expected loss at 95% percentile	3.27	5.21	1.45
Expected loss at 99% percentile	4.37	6.65	2.03
$1-\beta = 0.5$	Scarf's rule	Minimax regret	Maxentropy
Expected loss at 95% percentile	3.14	2.87	2.21
Expected loss at 99% percentile	4.93	4.23	3.23
$1-\beta = 0.2$	Scarf's rule	Minimax regret	Maxentropy
Expected loss at 95% percentile	6.82	4.76	1.53
Expected loss at 99% percentile	8.44	6.20	2.19

From the results in Table 4, we can conclude that the maximum entropy approach is also superior in performance for results far out in the tail. How is this possible? The answer is very simple. When we sample distributions using the above procedure, we almost never sample a distribution that is close to the cases where either Scarf's rule or minimax regret is superior.

4.2 Large coefficients of variation

To examine how the size of $\frac{\sigma}{\mu}$ influences performance, we repeated the simulation above discarding all cases where $\frac{\sigma}{\mu} < 0.5$. Table 5 and Table 6 detail the results.

Table 5: Expected profit loss using different ordering rules, $\frac{\sigma}{\mu} \geq \frac{1}{2}$			
$1-\beta=0.8$ Mean expected profit 85.02			
	Scarf's rule	Minimax regret	Maxentropy
Mean expected loss	2.06	3.71	0.78
(standard deviation)	(1.73)	(2.62)	(0.73)
$1-\beta=0.5$ Mean expected profit 31.72			
	Scarf's rule	Minimax regret	Maxentropy
Mean expected loss	1.66	1.46	1.15
(standard deviation)	(1.79)	(1.48)	(1.15)
$1-\beta=0.2$ Mean expected profit 5.54			
	Scarf's rule	Minimax regret	Maxentropy
Mean expected loss	5.54	3.51	0.73
(standard deviation)	(3.07)	(2.35)	(0.70)

Table 6: Expected profit loss at 95% and 99% percentiles, $\frac{\sigma}{\mu} \geq \frac{1}{2}$			
$1-\beta=0.8$	Scarf's rule	Minimax regret	Maxentropy
Expected loss at 95% percentile	5.43	8.47	2.26
Expected loss at 99% percentile	7.07	10.58	3.11
$1-\beta=0.5$	Scarf's rule	Minimax regret	Maxentropy
Expected loss at 95% percentile	5.16	4.48	3.50
Expected loss at 99% percentile	8.15	6.49	5.13
$1-\beta=0.2$	Scarf's rule	Minimax regret	Maxentropy
Expected loss at 95% percentile	11.11	7.83	2.16
Expected loss at 99% percentile	13.37	9.73	3.03

Comparing Table 3 and 4 with Table 5 and 6, we see that when cases with a small coefficient of variation are excluded, all expected losses increase by roughly 60%, indicating that the choice of method is more important when the coefficient of variation $\frac{\sigma}{\mu}$ is large.

5 Mixture distributions

In this section, we examine the performance of different ordering rules computing expected losses over mixtures of distributions. This may appear to be different from the sampling procedure carried out in Section 4, but it is not really. Alternatively, we could have obtained the results in this section by sampling a large number of mixture coefficients. The advantage of the mixture approach is that with it we can compute expected values (expectation taken over a set of distributions) instead of mean values, and hence avoid problems related to finite samples.

5.1 General mixtures

The normal distribution is very often used to model demand distributions. As total demand can be viewed as the sum of individual, independent demands, invoking the central limit theorem is an obvious first step. On closer examination, however, this argumentation fails because individual demands are not generally identically distributed.

Consider the following special case: we want to sell newspapers. The buyers can be divided into two groups, regular buyers and people that only buy newspapers on special events. Conditional on scenario S_1, for regular events, the demand is $N(\mu_1, \sigma_1^2)$. In the case of scenario S_2, special events, the other group also wants to buy newspapers. If a special event occurs, the total demand is then $N(\mu_2, \sigma_2^2)$ with $\mu_2 > \mu_1$. If we apply the law of total probability, we obtain:

$$P(D \leq d) = P(D \leq d|S_1)P(S_1) + P(D \leq d|S_2)P(S_2), \tag{30}$$

and the resulting demand is hence a mixture distribution of the form:

$$f_D(x) = \alpha f_1(x) + (1-\alpha)f_2(x) \qquad\qquad 0 \leq \alpha \leq 1, \tag{31}$$

where $\alpha = P(S_1)$ and $(1-\alpha) = P(S_2)$. This can clearly be generalized to situations with n different scenarios, resulting in a mixture distribution:

$$f_D(x) = \sum_{i=1}^{n} \alpha_1 f_i(x) \qquad\qquad \sum_{i=1}^{n} \alpha_i = 1. \tag{32}$$

Note that the distributions we sampled in Section 4 have an obvious interpretation within this framework; d_i is the demand under scenario S_i and p_i is the probability that scenario S_i occurs.

Gallego et al. (2001) consider a related setting. They propose expressions for the expected cost for two similar single-period models when the cost depends on the joint distribution of two random variables and only limited information on the marginal distributions is available. Further, Gallego et al. (2001) consider the newsvendor problem when the demand at the salvage price is a random variable and not indefinite as assumed in the classical model. Lastly, they examine the one-period seat allocation problem when the demand for super fare seats is a random variable and not indefinite as generally assumed in the standard model.

5.2 Mixtures of normal distributions

If D follows a normal distribution, it is well known that the error using Scarf's rule is always small. Gallego and Moon (1993) report that the error seldom exceeds 0.36% of the expected profit. If we know that demand is (approximately) normally distributed, however, there is little reason to apply Scarf's rule in the first place, so that argument is not as good as it first might appear.

To pursue the performance under normal distributions in more detail, we considered a mix of two normal distributions. Given $\mu_1 = 70, \sigma_1 = 20$ and $\mu_2 = 140, \sigma_2 = 30$ define a one-parameter family of distributions D_α with densities:

$$f_\alpha(x) = \alpha f_1(x) + (1 - \alpha) f_2(x). \tag{33}$$

To work with distributions with positive support, the f_1 and f_2 were truncated at zero and renormalized. With expectations and variances as above, the effect of the truncation is very small. Truncation effects are hence unimportant in this case. To evaluate the performance of the different ordering rules, we assume that α is uniformly distributed on $[0, 1]$.

For each fixed choice of α, we can compute $\mu = \mu(D_\alpha), \sigma = \sigma(D_\alpha)$ and use these values to find the corresponding order quantities using Scarf's rule, minimax regret, and maximum entropy. Given a particular value on β, we can then compute the expected loss EL_α for each ordering rule.

Given we have assumed that α is uniformly distributed on $[0,1]$, we can then take expectation over α to compute expected performance, i.e.:

$$\text{Expected performance} = \int_0^1 \text{EL}_\alpha d\alpha \tag{34}$$

for each of the different ordering rules. The results are shown as functions of β in Figure 3.

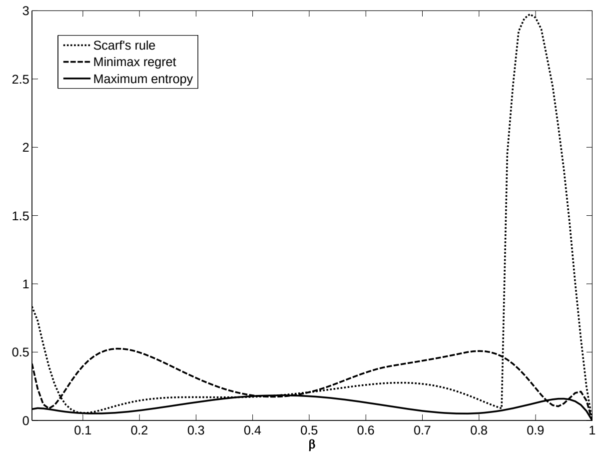

Figure 3: Expected loss using a mixture of normal distributions

The results in Figure 3 support the conclusions in Section 4. With the exception of a small interval near $\beta = 0.45$, we can see that the performance of the maximum entropy approach is clearly better than the alternatives. It is particularly interesting to note the behavior of Scarf's rule when $\beta \in [0.84, 0.96]$. Here the interval $[0.84, 0.96]$ is simply the range of the function:

$$\alpha \mapsto \frac{\mu_\alpha^2}{\mu_\alpha^2 + \sigma_\alpha^2}, \tag{35}$$

where μ_α and σ_α^2 are the expectation and variance for the mixture D_α. At $\beta = 0.84$, truncation is activated, and all orders are truncated at $\beta = 0.96$. We note that the effect of the truncation is quite devastating, as while the truncation hedges against negative expected profit, the hedge is clearly very costly.

The expected profit under full information is roughly 100 at $\beta = 0$, and decreases monotonically to 0 at $\beta = 1$. For small β, the losses are relatively minor, but increase considerably with β. In the interval $\beta \in [0.7, 0.8]$ (i.e., before truncation comes into play), the value of the fraction

$$\frac{\text{Expected loss}}{\text{Expected profit}} \tag{36}$$

is roughly 1.5% with Scarf's rule and reduces to about 0.5% using the maximum entropy approach. This supports the effect of relative losses in Table 3.

5.3 Mixtures of exponential distributions

To examine the performance in another special case, we repeated the construction in Section 5.2 with a mixture of exponential distributions. The exponential distribution is a special case here related to maximum entropy. It is straightforward to prove that if $\mu = \sigma$, the solution to the maximum entropy problem is an exponential distribution (see also the graph to the left in Figure 1).

In a mixture of exponential distributions, the maximum entropy approach provides a perfect fit at the extremes $\alpha = 0$ and $\alpha = 1$. If $0 < \alpha < 1$, the fit is not perfect, but still extremely good. In the calculations, we used the particular values $\mu_1 = 40$ and $\mu_2 = 70$, and the resulting expected performance (expectation taken over a uniform α mixture) is shown in Figure 4.

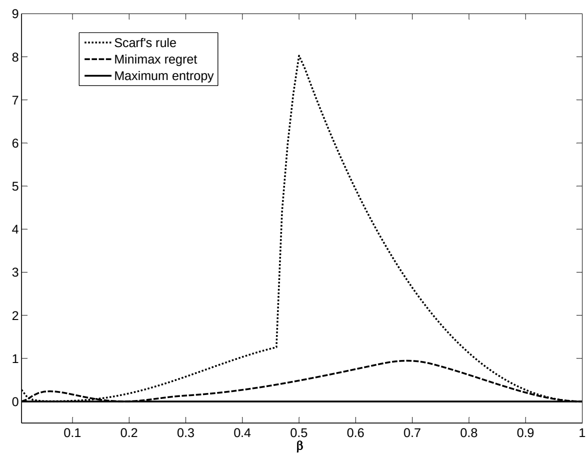

Figure 4: Expected loss using a mixture of exponential distributions

We can see that the maximum entropy approach is extremely efficient in this case (the expected loss is of an order of magnitude of 10^{-3}). Another effect worth noting is the truncation effect in Scarf's rule. In this particular case, truncation takes effect when $\beta \in [0.46, 0.5]$. We note again that the truncation effect is quite devastating for performance, supporting the point raised at the end of Section 5.2.

6 Alternative comparisons

In this section, we show how to construct a lower bound for the expected loss for a known mixture distribution, i.e., the case where the distribution of the mixture parameters is known, but the actual values of these parameters are unknown. We will also compare the performance of the ordering rules under the worst-case scenario, i.e., the cases we hedge using Scarf's rule.

6.1 Maximizing double expectation

Performance in this paper is measured in terms of double expectation, i.e., expectation taken over both $D \in \mathcal{D}$ and $\omega \in \Omega$. If we consider a mixture distribution D_α with density:

$$f_{D_\alpha}(x) = \sum_{i=1}^{n} \alpha_i f_i(x) \qquad\qquad \sum_{i=1}^{n} \alpha_i = 1, \tag{37}$$

where $\alpha = (\alpha_1, \ldots, \alpha_n)$ is a random variable on a probability space (A, Q), we can easily compute the fixed order q that maximizes double expectation. The reason is that $D(\alpha, \omega) = D_\alpha(\omega)$ defines a random variable on the product space $A \times \Omega$, and the problem of finding the optimal q reduces to a standard newsvendor problem. Note that the value of q must be determined before the draw, i.e., at a point in time where the value of α is still unknown.

If α has a density g, the density of $D = D(\alpha, \omega)$ can be computed via:

$$f_D(x) = \int_A f_{D_\alpha}(x) g(\alpha) dQ(\alpha). \tag{38}$$

In particular, if α is uniformly distributed, it is straightforward to verify that $f_D(x) = \frac{1}{n}\sum_{i=1}^{n} f_i(x)$, and the optimal fixed order quantity $q^* = F_D^{-1}(1-\beta)$ is found by solving the equation:

$$\int_0^{q^*} \frac{1}{n} \sum_{i=1}^{n} f_i(x) dx = 1 - \beta. \tag{39}$$

The scenario we focus on in this paper is the case where $\mu = \mu(D)$ and $\sigma = \sigma(D)$ are known, and infinitely many distributions have the same μ and σ. In the above context, it is hence natural to focus on the particular case where all the distributions making up the mixture have the same μ and σ. When this happens, for any α, we have $\mu(D_\alpha) = \mu$ and $\sigma(D_\alpha) = \sigma$, i.e., any mixture leads to the same μ and σ. Moreover, in the double expectation:

$$\mathrm{E}[D(\alpha,\omega)] = \int_0^\infty \int_A x f_{D_\alpha}(x) g(\alpha) dQ(\alpha) dx = \int_A \left(\int_0^\infty x f_{D_\alpha}(x) dx \right) g(\alpha) dQ(\alpha) = \mu, \tag{40}$$

$$\begin{aligned} \mathrm{Var}[D(\alpha,\omega)] &= \int_0^\infty \int_A (x-\mu)^2 f_{D_\alpha}(x) g(\alpha) dQ(\alpha) dx \\ &= \int_A \left(\int_0^\infty (x-\mu)^2 f_{D_\alpha}(x) dx \right) g(\alpha) dQ(\alpha) = \sigma^2. \end{aligned} \tag{41}$$

In that case, the order quantity $q^* = F_D^{-1}(1-\beta)$ provides a lower bound on the loss, in the sense that no strategy depending only on μ and σ can achieve a higher double expectation.

Example

Assume that we have a uniform mixture of two distributions, where:

$$f_1(x) = \frac{1}{\sqrt{2\pi}\sigma_1} e^{-(x-\mu_1)^2/(2\sigma_1^2)} \qquad f_2(x) = \frac{1}{2} \cdot \frac{1}{\sqrt{2\pi}\sigma_2} e^{-(x-\mu_2)^2/(2\sigma_2^2)} + \frac{1}{2} \cdot \frac{1}{\sqrt{2\pi}\sigma_3} e^{-(x-\mu_3)^2/(2\sigma_3^2)}. \tag{42}$$

If $\mu_1 = 100, \sigma_1 = 25$ and $\mu_2 = 77.088, \mu_3 = 122.912, \sigma_2 = 10, \sigma_3 = 10$, it is easy to verify that both distributions have $\mu = 100, \sigma = 25$. For this mixture, we can carry out the same exercise as shown in Figures 3 and 4. This time, however, we can include an additional strategy based on q^* from (39). The result is shown in Figure 5.

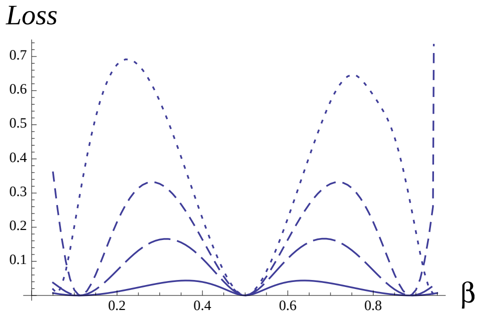

Figure 5: Expected loss using minimax regret (short dashed line), Scarf's rule (medium dashed line), maximum entropy (long dashed line), and q^* from (39) (undashed line)

Figure 5 supports a strategy based on q^* from (39) that provides a lower bound on the expected loss.

6.2 Worst-case scenarios

To put the above results into perspective, we wish to test performance under the worst-case scenario. A worst case under Scarf's rule, see Gallego and Moon (1993), occurs using a two-point distribution where:

$$d_1 = \mu - \sigma\sqrt{\frac{\beta}{1-\beta}} \qquad d_2 = \mu + \sigma\sqrt{\frac{1-\beta}{\beta}}, \tag{43}$$

$$p_1 = 1 - \beta \qquad p_2 = \beta. \tag{44}$$

For simplicity, we assume that $\beta \leq \frac{\mu^2}{\mu^2+\sigma^2}$, in which case d_1 and d_2 are both nonnegative and no truncation occurs. A close examination of this case reveals that the hedging is more subtle than it at first appears. As $p_1 = 1 - \beta$, this case is degenerate (see Section 2.2) and the expected value is in fact constant when $q \in [d_1, d_2]$. Hence, all the ordering rules offer a perfect fit for this particular distribution. The hedge in Scarf's rule, however, is against all the worst cases that might occur *if* we make an order different from that suggested by Scarf's rule. We should then not only consider the case given by (43) and (44), but all those other cases as well.

Given values for μ, σ, β, we can compute the order quantities q_1 for Scarf's rule and q_2 for the maximum entropy approach. The particular cases where $\mu = 100, \sigma = 50$ and $\mu = 100, \sigma = 100$ are shown as functions of β in Figure 6. In the left plot, we have used $\mu = 100, \sigma = 50$, while the values $\mu = 100, \sigma = 100$ were used in the plot to the right.

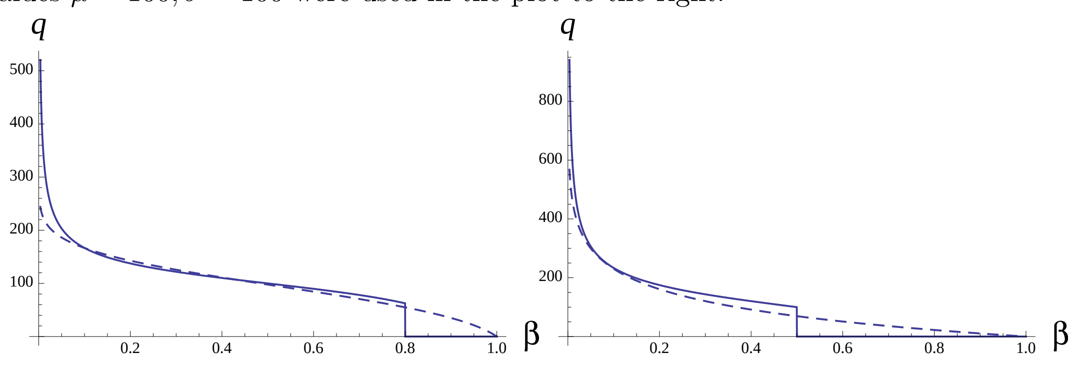

Figure 6: Ordering quantities using Scarf's rule and maximum entropy

To assess the performance of Scarf's rule, we examine the expected loss for the worst case that can occur when the "wrong order" q_2 is made. A distribution leading to the worst expected profit given q_2, see Gallego and Moon (1993), occurs on a two-point distribution where:

$$d_1 = \mu - \sigma\sqrt{\frac{\gamma}{1-\gamma}} \qquad d_2 = \mu + \sigma\sqrt{\frac{1-\gamma}{\gamma}}, \tag{45}$$

$$p_1 = 1-\gamma \qquad p_2 = \gamma, \tag{46}$$

$$\gamma = 1 - \frac{\sqrt{\sigma^2 + (q_2-\mu)^2} + (q_2-\mu)}{2\sqrt{\sigma^2 + (q_2-\mu)^2}}. \tag{47}$$

It is straightforward to verify that if q_2 equals q_1 (given by (9)), then $\gamma = \beta$. In Figure 6 (left part), the two curves intersect at $\beta = 0.1$ and $\beta = 0.43$. At points $q_2 = q_1$, there is a perfect fit for both strategies. At any other point $q_2 \neq q_1$, Scarf's rule performs better. Figure 7 shows the expected loss for both strategies.

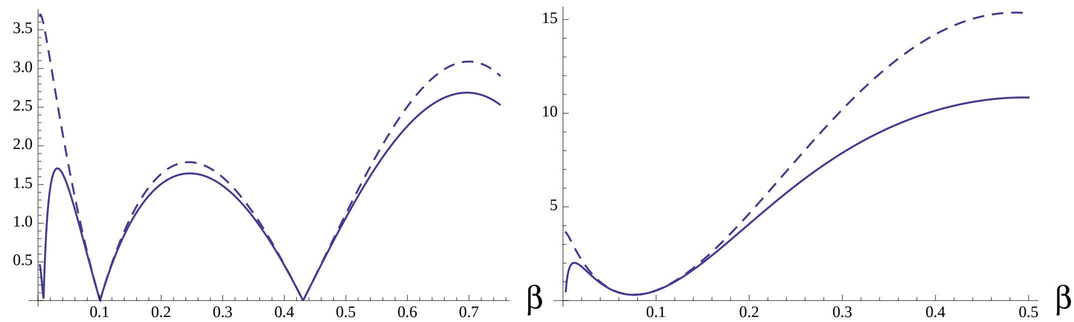

Figure 7: Expected loss using Scarf's rule and maximum entropy (dashed line)

In the left plot, we have used $\mu = 100, \sigma = 50$, while the values $\mu = 100, \sigma = 100$ were used in the plot to the right. For each β in Figure 7, we have used the distribution given by (45)–(47), and computed the expected loss (compared with full information) for this particular distribution. Note, we are using different distributions for different values of β.

We can see that the performance of Scarf's rule is better in this case. That is hardly surprising as it is true by definition of Scarf's rule. From the figures, however, we can also conclude that the effect is moderate for most values of β.

Figure 8 shows the result when the same exercise is used to compare Scarf's rule with minimax regret. For each β, we use the distribution with the worst expected value at the order quantity suggested by minimax regret.

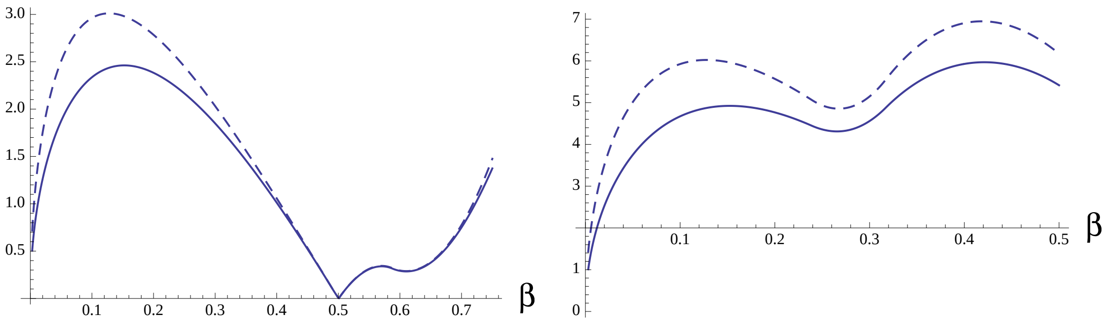

Figure 8: Expected loss using Scarf's rule and minimax regret (dashed line)

In the left plot, we have used $\mu = 100, \sigma = 50$, while the values $\mu = 100, \sigma = 100$ were used in the plot to the right. Note again that we are using different distributions for different values of β, and that these distributions are different from those we used to compare Scarf's rule and maximum entropy. Importantly, the distributions depend on the alternative to Scarf's rule and hence a comparison between maximum entropy and minimax regret is inappropriate.

7 More general cases

In this paper, we have focused on the case where μ, σ are known and $D \geq 0$. The constructions can easily be modified to cover more general moments/percentiles. For example, if the median m is known, it amounts to a linear restriction:

$$\int_0^m f_D(x)dx = \frac{1}{2}. \tag{48}$$

The same derivation as before leads to a density on the form:

$$f_D(x) = e^{a+bx+cx^2+d\mathcal{X}_{[0,m]}(x)} = \begin{cases} e^{a+bx+cx^2+d} & \text{if } 0 \leq x \leq m \\ e^{a+bx+cx^2} & \text{if } x > m. \end{cases} \tag{49}$$

In general, we can handle any additional restriction on the form:

$$\int_r^s x^u f_D(x)dx = \alpha, \tag{50}$$

when r, s, u, α are known; see the introductory example in Section 3.2 where we specified $r = \min[D], s = \max[D]$ and used this to construct the maximum entropy distribution. It thus appears to be straightforward to cover cases where we have such extra information. By comparison, maximin and minimax problems become very difficult to solve in these cases. The case covered in Section 2.4 is already quite difficult to solve, and to our best knowledge no explicit formulas exist for more general cases.

Even if systems such as (22) are nonlinear, they appear to contract quickly to a solution. Fixing all but one parameter, we could update that parameter such that it gives perfect fit in the restriction creating the parameter. This is repeated until all the parameters are updated. This idea is similar to the Bregman-balancing algorithm in Bregman (1967). Each iteration is linear in computation time with respect to the number of parameters, hinting that systems with several extra parameters can be handled numerically. We leave discussion of such extended cases to future research.

8 Concluding remarks

While several authors briefly mention maximum entropy as a possible line of approach, we believe that the current paper is the first that provides an in-depth analysis of maximum entropy as a viable alternative. The intention of this work has not been to disparage Scarf's rule or minimax regret as both are interesting and relevant theories. Instead, we wished to focus on the positive performance of the maximum entropy approach. In the maximum entropy approach, we seek the most probable distribution given the information we have at hand. It then follows intuitively that this approach leads to a better average simply because many distributions are fairly close to the most likely one.

The interesting question then is not that the maximum entropy approach does better on average, but rather to form a qualified opinion about how much better it performs in a wide variety of cases. From the experiments reported in this paper, it came as a complete surprise that this approach also appears to perform better at the tails, e.g., the worst case remaining when the 5% or 1% worst cases are excluded. In fact, we initially examined the tails merely to verify that

Scarf's rule was performing better and were quite puzzled to see the opposite.

In summary, the experiments carried out in Sections 4 and 5 yield a very clear vote in favor of the maximum entropy approach. Overall, it provides better expected profits, while at the same time offering better profits in cases that are rare, i.e., at the 95% or 99% percentiles. To balance the discussion, we included Section 6.2 on the worst-case scenarios. This was to ensure that there are cases where Scarf's rule performs better than any other approach. However, our general impression is that it appears overly conservative to base decisions on such extreme cases.

When compared with Scarf's rule and minimax regret, the maximum entropy approach can be more easily adjusted to take into account additional information on range, moments and percentiles. If the reader agrees that the performance of the maximum entropy approach is at least as good as the alternatives, the extra versatility provides an additional vote in favor of this approach.

Given μ and σ, it is fairly straightforward to solve the system in (22) (with $d = 0$ and $e = \infty$). Once this is done, one can easily find the optimal order quantity using (24). Of course, we cannot do this using a calculator, but can do so using very few lines of programming code. Why not give it a go?

References

[1] Benzion U, Y Cohen, T Shavit, 2010. The newsvendor problem with unknown distribution, *Journal of the Operational Research Society*, 61, 6, 1022–1031.

[2] Berman O, D Krass, MM Tajbakhsh, 2011. On the benefits of risk pooling in inventory management, *Production and Operations Management*, 20, 1, 57–71.

[3] Bregman LM, 1967. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming, *USSR Computational Mathematics and Mathematical Physics*, 7, 200–217.

[4] Eren S, C Maglaras, 2006. Revenue management heuristics under limited market information: a maximum entropy approach. Paper presented at the 6th Annual INFORMS Revenue Management Conference, June 5–6, Columbia University.

[5] Gallego G, I Moon, 1993. The distribution free newsboy problem: review and extensions, *Journal of the Operational Research Society*, 44, 8, 825–834.

[6] Gallego G, K Katiricioglu, B Ramachandran, 2007. Inventory management under highly uncertain demand, *Operations Research Letters*, 35, 281–289.

[7] Gallego G, JK Ryan, D Simchi-Levi, 2001. Minimax analysis for finite-horizon inventory models, *IIE Transaction*, 33, 861–874.

[8] Jaynes ET, 1957. Information theory and statistical mechanics, *Physical Review*, 106, 620–630.

[9] Jaynes ET, 2003. *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge, UK.

[10] Jiang HY, S Netessine, S Savin, 2010. Robust newsvendor competition under asymmetric information, *Operations Research*, 59, 1, 254–261.

[11] Jörnsten K, SL Nonås, L Sandal, J Ubøe, 2011. Transfer of risk in the newsvendor model with discrete demand, *OMEGA*, doi:10.1016/j.omega.2011.07.001 (forthcoming).

[12] Lan YJ, MO Ball, IZ Karaesmen, 2011. Regret in overbooking and fare-class allocation for single leg, *Manufacturing & Service Operations Management*, 13, 2, 194–208.

[13] Lee CM, SL Hsu, 2011. The effect of advertising on the distribution-free newsboy problem, *International Journal of Production Economics*, 129, 1, 217–224.

[14] Lim AEB, JG Shanthikumar, 2007. Relative entropy, exponential utility, and robust dynamic pricing, *Operations Research*, 55, 2, 198–214.

[15] Lin J, TS Ng, 2011. Robust multi-market newsvendor models with interval demand data, *European Journal of Operational Research*, 212, 2, 361–373.

[16] Moon I, S Choi, 1995. The distribution free newsboy problem with balking, *Journal of Operational Research Society*, 45, 6, 651–658.

[17] Perakis G, G Roels, 2008. Regret in the newsvendor model with partial information, *Operations Research*, 56, 1, 188–203.

[18] Perakis G, G Roels, 2010. Robust controls for network revenue management, *Manufacturing & Service Operations Management*, 12, 1, 56–76.

[19] Savage, LJ, 1951. The theory of statistical decisions. *J. Amer. Statist. Assoc.*, 46(253), 55–67.

[20] Scarf H, 1958. A min-max solution of an inventory problem. In *Studies in The Mathematical Theory of Inventory and Production*, (K Arrow, S Karlin and H Scarf, Eds), 201–209.

[21] Silver EA, DF Pyke, R Peterson, 1998. *Inventory Management and Production Planning and Scheduling*. John Wiley & Sons, New York, Third edition.

[22] Tajbakhsh MM, 2010. On the distribution free continuous-review inventory model with a service level constraint, *Computers & Industrial Engineering*, 59, 4, 1022–1024.

[23] Yue J, B Chen, M-C Wang, 2006. Expected value of distribution information for the newsvendor problem, *Operations Research*, 54, 6, 1128–1136.