

Confidence-based Optimization for the Newsvendor Problem

Roberto Rossi

University of Edinburgh Business School, UK,
roberto.rossi@ed.ac.uk

Steven Prestwich

Dept. of Computer Science, University College Cork, Ireland,
s.prestwich@4c.ucc.ie

S. Armagan Tarim

Dept. of Management, Hacettepe University, Turkey,
armagan.tarim@hacettepe.edu.tr

Brahim Hnich

Dept. of Computer Engineering, Izmir University of Economics, Turkey,
brahim.hnich@ieu.edu.tr

Abstract

We introduce a novel strategy to address the issue of demand estimation in single-item single-period stochastic inventory optimisation problems. Our strategy analytically combines confidence interval analysis and inventory optimisation. We assume that the decision maker is given a set of past demand samples and we employ confidence interval analysis in order to identify a range of candidate order quantities that, with prescribed confidence probability, includes the real optimal order quantity for the underlying stochastic demand process with unknown stationary parameter(s). In addition, for each candidate order quantity that is identified, our approach can produce an upper and a lower bound for the associated cost. We apply our novel approach to three demand distribution in the exponential family: binomial, Poisson, and exponential. For two of these distributions we also discuss the extension to the case of unobserved lost sales. Numerical examples are presented in which we show how our approach complements existing frequentist — e.g. based on maximum likelihood estimators — or Bayesian strategies.

keywords: inventory control, newsvendor problem, confidence interval analysis, demand estimation, sampling.

1 Introduction

We consider the problem of controlling the inventory of a single item with stochastic demand over a single period. This problem is known as the “newsvendor” problem [Silver et al., 1998]. Most of the literature on the newsboy problem has focused on the case in which the demand distribution and its parameters are known. However, what happens in practice is that the decision maker must estimate the order quantity from a — possibly very limited — set of past demand realisations. This task is often complicated by the fact that unobserved lost sales must be taken into account.

Existing approaches to the newsvendor under limited historical demand data can be classified in non-parametric and parametric approaches. Non-parametric approaches operate without any access to and assumptions on the true demand distributions. Parametric approaches assume that demand realisations come from a given probability distribution — or from a family of probability distributions — and make inferences about the parameters of the distribution. When the class of the distribution is known, but its parameters must be estimated from a set of samples, non-parametric approaches may produce conservative results. For this reason several works in the literature investigated parametric approaches to the newsvendor under limited historical demand data; a complete overview on these works will be provided in Section 2. Two classical parametric approaches for dealing with the newsvendor under limited historical demand data are the maximum likelihood [see e.g. Scarf, 1959, Fukuda, 1970, Gupta, 1960] and the Bayesian approach [see e.g. Hill, 1997, 1999]. Both these strategies feature a number of asymptotical properties that guarantee their convergence towards the optimal control strategy. However, a decision maker finds herself rarely in an asymptotic situation, since only few samples are generally available to estimate an order quantity. This means that asymptotic properties often do not hold in practice. Unfortunately, both the maximum likelihood estimation and the Bayesian approaches ignore the uncertainty around the estimated order quantity and its associated expected total cost or profit. Hayes [1969] and, more recently, Akcay et al. [2011] discuss how to quantify this uncertainty by using the concept of expected total operating cost (ETOC), which represents the expected one-period cost associated with operating under an estimated inventory policy. By minimising this performance indicator, they identify an optimal biased order quantity that accounts for the uncertainty around the demand parameters estimated from limited historical data. Their approach, however, does not answer a number of fundamental questions. It does not state at what confidence level we can claim this order quantity to be optimal within a given margin of error; nor does it quantify the probability of incurring an expected total cost substantially different than the estimated one, when such an order quantity is selected. Kevork [2010] exploits the sampling distribution of the estimated demand parameters to study the variability of the estimated optimal order quantity and its expected total profit under a normally distributed de-

mand with unknown parameters. The author shows that these two estimators asymptotically converge to normality. Based on this property, asymptotic confidence intervals are derived for the true optimal order quantity and its expected total profit. Unfortunately, these confidence intervals achieve the prescribed confidence level only asymptotically and they represent an approximation when one operates under finite samples. In this work, we make the following contributions to the inventory management literature:

1. We analytically combine confidence interval analysis and inventory optimisation. By exploiting exact confidence intervals for the parameters of a given distribution, we identify a range of candidate order quantities that, according to a prescribed confidence probability, includes the real optimal order quantity for the underlying stochastic demand process with unknown stationary parameter(s). For each candidate optimal order quantity that is identified, our approach computes an upper and a lower bound for the associated cost. This range covers, once more according to a prescribed confidence probability, the actual cost the decision maker will incur if she selects that particular quantity.
2. To obtain the former result, when the order quantity is fixed, we establish convexity of the newsvendor cost function in the success probability p of a binomial demand (Theorem 3) and in the rate λ of a poisson demand (Theorem 4); we also establish that the newsboy cost function is quasi-convex in the expected value $1/\lambda$ of an exponential distribution (Theorem 1). These results are nontrivial and, to the best of our knowledge, they have not been established before in the literature.
3. For the binomial and the poisson distribution we demonstrate how to extend the discussion to account for unobserved lost sales.

Our strategy is frequentist in nature and based on the theory of statistical estimation introduced by Neyman [1937]. In contrast to Bayesian approaches, no prior knowledge is required to perform the computation, which is entirely data driven. In contrast to [Hayes, 1969, Akcay et al., 2011] we do not introduce new performance indicators, such as the ETOC, and we build our analysis on existing and well established results from inventory theory, i.e. the expected total cost of a policy; and from statistical analysis, i.e. confidence intervals. Finally, in contrast to [Kevork, 2010] our results are valid both asymptotically and under finite samples.

If the identified set of candidate optimal order quantities comprises more than a single value, expert assessment or any existing frequentist or Bayesian approach may be employed to select the most promising of these values according to a given performance indicator. By using our approach, the decision maker may then determine — at a given confidence level and from a limited

set of available data — the exact confidence interval for the expected total cost associated with such a decision, as well as the potential discrepancy between the true optimal decision and the one she selected. For this reason, a further contribution is the following.

4. We effectively complement a number of existing strategies that compute a point estimate of the optimal order quantity and its expected total cost. We demonstrate this fact for the Bayesian approach in [Hill, 1997] and for the frequentist approach based on the maximum likelihood estimator of the demand distribution parameter.

2 Literature survey

A thorough literature review on the newsvendor problem is presented by Khouja [2000]. Among other extensions, the author surveyed those dealing with different states of information about demand. As Silver et al. [1998] point out, in practice demand distribution may be not known. Khouja [2000] points out that several authors relaxed the assumptions of having a specific distribution with known parameters.

One of the earliest approaches to dealing with different states of information about demand is the so-called “maximin approach”, which consists in maximising the worst-case profit. Scarf et al. [1958] applied the maximin approach to the newsvendor problem and assumed that only the mean and the variance of demand are known. Under this assumption, they derived a closed-form expression for the optimal order quantity that maximises the expected total profit against the worst-case demand distribution. Gallego and Moon [1993] provided a simplified form of the rule in [Scarf et al., 1958]; beside this, they also discussed four extensions: a recourse case, a model including fixed ordering cost, a random yield case and a multi-product case. This model was extended to account for balking in Moon and Choi [1995]. Yue et al. [2006] assumed that the demand distribution belongs to a certain class of probability distribution functions with known mean and standard deviation; the authors’ aim is to compute the maximum expected value of distribution information over all possible probability distribution functions with known mean and standard deviation for any order quantity. Perakis and Roels [2008] pointed out that the maximin objective is, generally speaking, conservative, since it focuses on the worst-case scenario for the demand. The authors therefore suggested adopting a less conservative approach: the “minimax regret”, in which the firm minimises its maximum cost discrepancy from the optimal decision. Works mentioned so far focused on a newsvendor setting. In contrast, Notzon [1970] discuss a minimax multi-period inventory model. Gallego et al. [2001] discuss a finite horizon inventory model in which the demand distribution is discrete and partially defined

by selected moments and/or percentiles. Bertsimas and Thiele [2006] analyse a distribution free inventory problem for which demand in each period is a random variable over a given support identified by two values: the lower and the upper estimators. A comparable model is found in Bienstock and Özbay [2008]. Ahmed et al. [2007] minimise a coherent risk measure instead of the total cost in an inventory control model and establish an equivalence between risk aversion described as a coherent risk measure and a minimax formulation of the problem. See and Sim [2010] discuss an inventory control problem under a demand model described by a given support, mean and standard deviation. They then consider a second-order cone optimisation problem that minimises the expected total cost among all distributions satisfying the demand model.

All the aforementioned works operate in a *distribution free* setting. The decision maker has access to partial information about the demand distribution, e.g. mean, variance, symmetry, unimodality etc, but does not know the class of the demand distribution, e.g. Poisson, normal etc. In practice, it is often the case that the decision maker can only access a set of past observations of the demand out of which a near-optimal inventory target must be estimated. Approaches trying to estimate a near-optimal inventory target from observed realisations of the demands can be classified as *non-parametric* or *parametric*.

Non-parametric approaches operate without any access to and assumptions on the true demand distributions. Levi et al. [2007] discuss a non-parametric approach which computes policies based only on observed samples of the demands. The authors determine bounds for the number of samples that are necessary in order to guarantee an arbitrary approximation of the optimal policy defined with respect to the true underlying demand distributions. Huh et al. [2009] discuss an adaptive inventory policy that deals with censored observations, thus effectively relaxing the assumption that past demand data is observable. Other non-parametric approaches based on order statistics were proposed in [Hayes, 1971, Lordahl and Bookbinder, 1994]; approaches based on bootstrapping techniques were discussed in [Bookbinder and Lordahl, 1989, Fricker and Goodhart, 2000].

Parametric approaches assume that demand realisations come from a given probability distribution and make inferences about the parameters of the distribution. The class of the distribution may be determined, for instance, by selecting the maximum entropy distribution that matches the structure of the demand process at hand (see also the discussion in Perakis and Roels [2008] p. 190). According to Jaynes’ “principle of maximum entropy”, introduced in [Jaynes, 1957], subject to known constraints (i.e. testable information), the probability distribution which best represents the current state of knowledge is the one with largest entropy. When the class of the distribution is known but its parameters must be estimated from a set of samples non-parametric approaches such as those discussed so far are not appropriate, since they would produce conservative results.

According to Berk et al. [2007] there are two general approaches for dealing with a stochastic decision making environment in which random variables follow known distributions with unknown parameters: the Bayesian and the frequentist. In the Bayesian approach a “prior” distribution is selected for the demand distribution parameter(s). This selection may be based on collateral data and/or subjective assessment. Subsequently, a “posterior” distribution is derived from the prior distribution by using the demand data that is observed. This posterior distribution is used to construct the posterior distribution of the demand and to determine the optimal order quantity and objective function value. In the frequentist approach a parametric demand distribution is empirically selected and point estimates, e.g. maximum likelihood or moment estimators, for the unknown parameters are obtained according to the observed data; these are then used to derive the optimal order quantity and objective function value.

According to Kevork [2010] another distinction can be made between approaches assuming that demand is fully observed and approaches assuming that demand occurring when the stock level drops to zero is lost and thus not observed. In the latter case, it is necessary to adjust the estimation procedure to account for unobserved demand. Lau and Lau [1996] further distinguish works on estimating demand distributions with stockouts in two groups: estimating the initial demand distribution, e.g. Conrad [1976], Lau and Lau [1996]; and updating the demand distribution parameters Wecker [1978], Bell [1981], Hill [1992].

Bayesian approaches under fully observed demand were proposed in Scarf [1959, 1960], Iglehart [1964], Azoury [1985], Lovejoy [1990], Bradford and Sugrue [1990], Hill [1997], Eppen and Iyer [1997], Hill [1999], Lee [2008], Bensoussan et al. [2009]. Bayesian approaches under censoring induced by lost sales include Lariviere and Porteus [1999], Ding and Puterman [1998], Berk et al. [2007], Chen [2010], Lu et al. [2008], Mersereau [2012]. Bayesian approaches suffer from a number of drawbacks. First, an initial “belief” about the unknown parameters must be expressed as a prior distribution of the demand. It is often assumed that this prior distribution is obtained from collateral data and/or subjective assessment. The need of a prior distribution is structural in the Bayesian approach, which relies on the update of the prior distribution to derive the posterior distribution of the demand given the data. When no supporting information is available, “uninformative” priors can be used, see e.g. Hill [1999], but these tend to introduce a strong bias especially under limited available data to perform Bayesian updating. This fact is well known in the life sciences, see e.g. van Dongen [2006], but it is often ignored in more theoretical settings. At the end of this work, in Section 7, we will provide a practical exemplification of this fact. A second issue that arises with existing Bayesian approaches to the newsvendor problem is that to show that the order quantity derived via the Bayesian approach converges to the optimal order quantity one has to consider an infinite set of samples, see e.g. Bensoussan et al. [2009]. However, in practice it is often the case that available information is very limited. Unfortunately,

Bayesian approaches can be shown to be often biased under small sample sets, especially due to the fact that the choice of the prior may strongly influence the order quantity obtained.

Two early frequentist approaches are Nahmias [1994] and Agrawal and Smith [1996]. Nahmias [1994] analyzed the censored demand case under a normally distributed demand. Agrawal and Smith [1996] considered the estimation of a negative binomial demand under censoring induced by lost sales. However, these two studies consider the stock level as given and thus do not address the associated optimization problem of finding the optimum stock level. More recently, Liyanage and Shanthikumar [2005] introduced the “operational statistics” framework, in which an optimal order quantity, rather than the parameters of the distribution, is directly estimated from the data. The authors consider the case in which it is known that the demand distribution function belongs to a parameterised family of distribution functions. In contrast to the Bayesian approach, they do not assume any prior knowledge on the parameter values. They demonstrate that by combining demand parameter estimation and inventory-target optimisation a higher expected total profit can be achieved with respect to traditional approaches that separate estimation and optimisation. Klabjan et al. [2013] integrate distribution fitting and robust optimisation by identifying a set of demand distributions that fit historical data according a given criterion; they then characterise an optimal policy that minimises the maximum expected total cost against such set of demand distributions.

There is an important limitation that is common to all approaches surveyed so far. Consider a possibly very limited set of past demand observations; strategies based on frequentist analysis, e.g. maximum likelihood estimators and distribution fitting; or on Bayesian analysis, e.g. Hill [1999], would analyze these data and provide a single most-promising order quantity and an estimated cost associated with it. However, given the available data, they do not answer a number of fundamental questions: we do not know at what confidence level we can claim the quantity selected to be optimal within a given margin of error; and we also do not know the probability of incurring a cost substantially higher than the estimated one, when such an order quantity is selected.

To the best of our knowledge Kevork [2010] was the first to exploit the sampling distribution of the estimated demand parameters in order to study the variability of the estimates for the optimal order quantity and associated expected total profit. The author adopts a frequentist approach in which demand is fully observed in each period. By incorporating maximum likelihood estimators for mean and variance of demand into expressions that determine the optimal order quantity and associated expected total profit, the author develops estimators for these latter two variables. These estimators are shown to be consistent and to asymptotically converge to normality. Based on these properties, the author derives confidence intervals for the true optimal order quantity and associated expected total profit. Unfortunately, these estimators are biased

in finite samples and the associated confidence intervals achieve the prescribed confidence level only asymptotically.

As pointed out in [Akçay et al., 2011], the inventory manager finds rarely herself in an asymptotic situation, since an inventory target must be typically estimated from a small sample size. To quantify the uncertainty about distribution parameter estimates and thus about the estimated order quantity, Akçay et al. [2011] adopt the ETOC, originally introduced in [Hayes, 1969], which we recall represents the expected one-period cost associated with operating under an estimated inventory policy. Originally, Hayes [1969] discussed applications of ETOC to exponentially and normally distributed demands. More specifically, they identified the optimal biased order quantity that minimises the ETOC in presence of limited historical demand data. This was one of the first works blending statistical estimation with inventory optimisation. Akçay et al. [2011] extended this analysis to a parameterised family of distributions — the Johnson translation system — that has the ability to match any finite first four moments of a random variable and to capture a broad range of distributional shapes. The authors quantify the inaccuracy in the order quantity estimation by using the expected value of perfect information about the sampling distribution of the demand parameters for the estimated order quantity. By using this interpretation of the ETOC, they seek an order quantity that minimises the ETOC within a class of inventory target-estimators implied by the Johnson translation system. Despite its ability to quantify the inaccuracy in the inventory-target estimation as a function of the length of the historical data via the ETOC, the approach in [Akçay et al., 2011] does not identify a confidence interval that, with prescribed confidence probability, includes the real optimal order quantity for the underlying stochastic demand process with unknown parameter(s); neither it is able to produce a confidence interval for the expected total cost associated with ordering decisions in this interval.

3 The newsvendor problem

In this section, we shall summarize the key features of the newsvendor problem. For more details, the reader may refer to Silver et al. [1998]. Consider a one-period random demand d with mean μ and variance σ^2 . Let h be the unit overage cost, paid for each item left in stock after demand is realized, and let p be the unit underage cost, paid for each unit of unmet demand; we assume $p, h > 0$. Let $g(x) = hx^+ + px^-$, where $x^+ = \max(x, 0)$ and $x^- = -\min(x, 0)$. The expected total cost can be written as $G(Q) = E[g(Q - d)]$, where $E[\cdot]$ denotes the expected value.

It is a well-known fact that, for a continuous demand d ,

$$F(Q) = \Pr\{d \leq Q\} = \frac{p}{p+h} = \beta. \quad (1)$$

If F is continuous there is at least one Q satisfying Eq. 1, that is

$$Q^* = \inf\{Q \geq 0 : F(Q) = \beta\}.$$

For F strictly increasing, there is a unique optimal solution given by

$$Q^* = F^{-1}(\beta). \quad (2)$$

In practice, the probability distribution of the random demand d often has finite support over the set $\mathbb{N}_0 = \{0, 1, 2, \dots\}$. In this case it is useful to work with the forward difference $\Delta G(Q) = G(Q+1) - G(Q)$, $Q \in \mathbb{N}_0$. It is easy to see that

$$\Delta G(Q) = h - (h+p) \Pr\{d > j\}$$

is non-decreasing in Q , and that $\lim_{Q \rightarrow \infty} \Delta G(Q) = h > 0$, so an optimal solution is given by $Q^* = \min\{Q \in \mathbb{N}_0 : \Delta G(Q) \geq 0\}$ or equivalently

$$Q^* = \min\{Q \in \mathbb{N}_0 : F(Q) \geq \beta\}. \quad (3)$$

3.1 A frequentist and a Bayesian approach

Let us now consider the situation in which the decision maker knows the type of the random demand distribution (e.g. binomial), but in which he does not know the actual values of some or all the (stationary) parameters of such a distribution. Nonetheless, the decision maker is given a set of M past realizations of the demand. From these realizations he has to infer the optimal order quantity and, possibly, he has to estimate the cost associated with the optimal Q^* he has selected.

In what follows, we detail the functioning of a frequentist approach, i.e. the maximum likelihood approach, and of a Bayesian approach from the literature Hill [1997]. In the rest of this work we will make ample use of these two approaches, especially to discuss how our approach can be used to complement the results obtained by frequentist or Bayesian approaches. For the sake of brevity, in this work we will focus only on these two strategies from the literature. However, this choice is made without loss of generality. Our approach may in fact also complement any of the other frequentist or Bayesian approaches previously surveyed.

3.1.1 Maximum likelihood approach

A commonly adopted heuristic strategy for order quantity selection under sampled demand information consists in computing, from the available sample set, a point estimate for the unknown demand distribution parameter(s). This may be done by using the maximum likelihood estimator [Le Cam, 1990], thus obtaining the so-called maximum likelihood policy [see e.g. Scarf, 1959, Fukuda, 1970, Gupta, 1960], or the method of moments [Newey and McFadden, 1986]. For instance, assume that the available sample set comprises M observed past demand data, d_1, \dots, d_M , and that the demand is assumed to follow a binomial distribution. The binomial distribution comprises two parameters: the number of trials N and the success probability q . In the context of the newsvendor problem, N may be interpreted as the non-variant number of customers that enter the shop in a given day $i \in \{1, \dots, M\}$, and q may be interpreted as the probability that a customer makes a purchase. Then, by assuming that N is known, the maximum likelihood estimator for parameter q is $\hat{q} = \sum_{i=1}^M d_i / (MN)$. After computing \hat{q} , the decision maker employs the random variable $\text{bin}(N; \hat{q})$ in place of the actual unknown demand distribution in Eq. 3 to compute the estimated optimal order quantity \hat{Q}^* and expected total cost $G(\hat{Q}^*)$.

3.1.2 Hill's Bayesian approach

A Bayesian approach to the Newsvendor problem under partial information is presented by Hill [1997]. Hill's approach considers a "prior" distribution, based on collateral data and/or subjective assessment, for the unknown parameter(s) of the demand distribution. As new data are observed, the prior distribution is updated and a "posterior" distribution is generated. Hill then uses the posterior distribution of the unknown parameter(s) for estimating the posterior distribution of the demand. Finally, the posterior distribution of the demand is used to estimate the order quantity that optimizes the Newsvendor cost function. More formally, we recall that Bayes' theorem tells us that

$$\Pr\{a_j|b\} = \frac{\Pr\{b|a_j\} \cdot \Pr\{a_j\}}{\sum_{i=1}^k \Pr\{b|a_i\} \cdot \Pr\{a_i\}},$$

where $\{a_1, \dots, a_k\}$ is a partition of the sample space and b is an observed event. Bayes actually discusses also the natural extension of the above theorem when a is continuous and b is discrete or continuous,

$$f(a|b) = \frac{f(b|a) \cdot g(a)}{\int f(b|u) \cdot g(u) du}.$$

The denominator is, of course, independent of a , therefore $f(a|b) \propto f(b|a) \cdot g(a)$. In the context of the Newsvendor problem, a represents the unknown parameter

of the demand distribution, b represents the actual set of observed demand samples. According to Hill, the prior distribution of a , $g(a)$, describes an estimate of the likely value that a might take, without considering the observed samples. This estimation is based on subjective assessment and/or collateral data. $f(b|a)$, also known as the likelihood, represents the probability of observing a set of samples b given a . The posterior distribution of a , $f(a|b)$, is an updated estimate of the values a is likely to take based on the prior distribution and the observed information. Hill uses an uninformative, also known as objective, prior to express “an initial state of complete ignorance of the likely values that the parameter a might take.” By employing the conjugate prior for the particular distribution under analysis, he constructs the posterior distribution for the Newsvendor demand as follows,

$$f(d|b) = \int f(d|a)f(a|b)db,$$

where the integral is computed over all permitted values of a . The Bayesian approach proposed by Hill then consists in using this posterior distribution in place of the unknown true distribution for the demand in Eq. 2. This immediately produces a candidate order quantity \hat{Q}^* .

4 Binomial demand

Consider a discrete random variable that follows a Bernoulli distribution. Accordingly, such a variable may produce only two possible outcomes, i.e. “yes” and “no”, with probability q and $1 - q$, respectively. This class of random variables is particularly useful in representing so called “Bernoulli trials”, which are experiments that can have one of two possible outcomes. These events can be phrased into “yes or no” questions, such as “did the customer purchase the newspaper?”

Consider the following situation: a Newsvendor has a pool of N customers that come every day to the stand. Each customer may buy a newspaper with probability q . It is a well-known fact that any experiment comprising a sequence of N Bernoulli trials, each having the same “yes” (respectively, “no”) probability q (respectively, $1 - q$), can be represented by a random variable $\text{bin}(N; q)$ that follows a binomial distribution [Jeffreys, 1961] with probability mass function

$$\Pr\{\text{bin}(N; q) = k\} = \binom{N}{k} q^k (1 - q)^{N-k},$$

where $k = 0, \dots, N$.

According to our previous discussion it is fairly easy to find the optimal order quantity Q^* for a given random demand $\text{bin}(N; q)$. We shall now give

a running example. Consider a random demand $\text{bin}(50, 0.5)$. Let $h = 1$ and $p = 3$, therefore $\beta = 0.75$. From Eq. 3 we compute $Q^* = 27$ and the respective expected total cost $G(Q^*) = 4.4946$.

Let us now consider the situation in which the parameter q is not known. The decision maker is given a set of M past realizations of the demand. From these realizations he has to infer the optimal order quantity and, possibly, he has to estimate the cost associated with the optimal Q^* he has selected.

Since we operate under partial information it may not be possible to uniquely determine “the” optimal order quantity and the exact cost associated with it. Therefore, we argue that a possible approach consists in determining a range of “candidate” optimal order quantities and upper and lower bounds for the cost associated with these quantities. This range will contain the real optimum according to a prescribed confidence probability.

4.1 Confidence intervals for the binomial distribution

Confidence interval analysis [Neyman, 1937, 1941] is a well established technique in statistics for computing, from a given set of experimental results, a range of values that, with a certain confidence level (or confidence probability), will cover the actual value of a parameter that is being estimated. Several techniques [Clopper and Pearson, 1934, Garwood, 1936, Trivedi, 2001, etc.] for building confidence intervals for a given sample set have been proposed.

Approximate techniques for building confidence intervals [see e.g. Agresti and Coull, 1998] become relevant because, especially with small sample sizes, an exact confidence interval may be unnecessarily conservative. In this work, we focus on the exact interval and we leave the analysis of the benefits brought by the use of approximate intervals as future research.

A method to compute “exact confidence intervals” for the binomial distribution has been introduced by Clopper and Pearson [1934]. This method uses the binomial cumulative distribution function in order to build the interval from the data observed. The Clopper-Pearson interval can be written as (q_{lb}, q_{ub}) , where

$$\begin{aligned} q_{lb} &= \min\{q \mid \Pr\{\text{bin}(N; q) \geq X\} \geq (1 - \alpha)/2\}, \\ q_{ub} &= \max\{q \mid \Pr\{\text{bin}(N; q) \leq X\} \geq (1 - \alpha)/2\}, \end{aligned}$$

X is the number of successes (or “yes” events) observed in the sample and α is the confidence probability. Note that we assume $q_{lb} = 0$ when $X = 0$ and $q_{ub} = N$ when $X = N$. As discussed by Agresti and Coull [1998], this interval can be also expressed using quantiles from the beta distribution. More specifically, the lower endpoint is the $(1 - \alpha)/2$ -quantile of a beta distribution with shape parameters X and $N - X + 1$, and the upper endpoint is the $(1 + \alpha)/2$ -quantile of

a beta distribution with shape parameters $X + 1$ and $N - X$. Furthermore, the beta distribution is, in turn, related to the F-distribution so a third formulation of the Clopper-Pearson interval, also discussed by Agresti and Coull [1998], uses quantiles from the F distribution.

It is intuitively clear that the “quality” of a given confidence interval is directly related to its size. The smaller the interval, the better the estimate. In general, confidence intervals that have symmetric tails (i.e. with associated probability $(1 - \alpha)/2$) are not the smallest possible ones. A large literature exists on the topic of determining the smallest possible intervals for a given parameter/distribution combination [see e.g. Zieliński, 2010]. The discussion that follows is independent of the particular interval adopted. For the sake of simplicity, we will adopt intervals having symmetric tails.

4.2 Solution method employing statistical estimation based on classical theory of probability

We shall now employ the Clopper-Pearson interval for computing an upper and a lower bound for the optimal order quantity Q^* in a Newsvendor problem under partial information. The confidence interval for the unknown parameter q of the binomial demand $\text{bin}(N; q)$ is simply (q_{lb}, q_{ub}) where

$$\begin{aligned} q_{lb} &= \min\{q \mid \Pr\{\text{bin}(MN; q) \geq X\} \geq (1 - \alpha)/2\}, \\ q_{ub} &= \max\{q \mid \Pr\{\text{bin}(MN; q) \leq X\} \geq (1 - \alpha)/2\}, \end{aligned}$$

and $X = \sum_{i=1}^M d_i$. Let Q_{lb}^* be the optimal order quantity for the Newsvendor problem under a $\text{bin}(N, q_{lb})$ demand and Q_{ub}^* be the optimal order quantity for the Newsvendor problem under a $\text{bin}(N, q_{ub})$ demand. Since $\Delta G(Q)$ is non-decreasing in Q , according to the available information with confidence probability α the optimal order quantity Q^* is a member of the set $\{Q_{lb}^*, \dots, Q_{ub}^*\}$.

We shall now compute upper (c_{ub}) and lower (c_{lb}) bounds for the cost associated with a solution that sets the order quantity to a value in the set $\{Q_{lb}^*, \dots, Q_{ub}^*\}$. Let us write the cost associated with an order quantity Q ,

$$G(Q) = h \sum_{i=0}^Q \Pr\{\text{bin}(N; q) = i\}(Q - i) + p \sum_{i=Q}^N \Pr\{\text{bin}(N; q) = i\}(i - Q).$$

Then, consider the function

$$G_Q(q) = h \sum_{i=0}^Q \Pr\{\text{bin}(N; q) = i\}(Q - i) + p \sum_{i=Q}^N \Pr\{\text{bin}(N; q) = i\}(i - Q), \quad (4)$$

in which the order quantity Q is fixed and in which we vary the “success” probability $q \in (0, 1)$. It can be proved that $G_Q(q)$ is convex in the continuous

parameter q ; this is trivially true when $Q = N$. The proof for $0 \leq Q < N$ is given in Appendix 8 (Theorem 3).

Although it is possible to prove that $G_Q(q)$ is convex in q , there is no closed form expression for finding the q^* that minimizes this function. Nevertheless, due to its convexity in q , it is clearly possible to use convex optimization approaches to find the q^* that minimizes or maximizes this function over a given interval.

Let us consider the confidence interval (q_{lb}, q_{ub}) for the parameter q of the binomial demand. For a given order quantity Q , consider the value

$$q_{Q,\min}^* = \arg \min_{q \in (q_{lb}, q_{ub})} G_Q(q) \quad \left(q_{Q,\max}^* = \arg \max_{q \in (q_{lb}, q_{ub})} G_Q(q) \right)$$

that minimizes (maximizes) $G_Q(q)$ for $q \in (q_{lb}, q_{ub})$. With confidence probability α , $G_Q(q_{Q,\min}^*)$ and $G_Q(q_{Q,\max}^*)$ represent a lower and an upper bound, respectively, for the cost associated with Q .

By recalling that the optimal order quantity Q^* is, with confidence probability α , a member of the set $\{Q_{lb}^*, \dots, Q_{ub}^*\}$, it is easy to compute upper (c_{ub}^*) and lower (c_{lb}^*) bounds for the cost that a manager will face, with confidence probability α , whatever order quantity he chooses in the candidate set $\{Q_{lb}^*, \dots, Q_{ub}^*\}$. The lower bound is

$$c_{lb}^* = \min_{Q \in \{Q_{lb}^*, \dots, Q_{ub}^*\}} G_Q(q_{Q,\min}^*)$$

and the upper bound is

$$c_{ub}^* = \max_{Q \in \{Q_{lb}^*, \dots, Q_{ub}^*\}} G_Q(q_{Q,\max}^*).$$

It should be emphasized that, when the confidence interval (q_{lb}, q_{ub}) covers the real parameter q of the binomial demand we are estimating, then the set $\{Q_{lb}^*, \dots, Q_{ub}^*\}$ comprises the optimal order quantity Q^* and the interval (c_{lb}^*, c_{ub}^*) comprises the real cost associated with every possible order quantity in $\{Q_{lb}^*, \dots, Q_{ub}^*\}$. Given the way confidence interval (q_{lb}, q_{ub}) is constructed, it is guaranteed that this happens with probability α .

Of course, by increasing the number M of past observations, we can decrease the size of confidence interval (q_{lb}, q_{ub}) . As a direct consequence, the cardinality of the set $\{Q_{lb}^*, \dots, Q_{ub}^*\}$ decreases. In the ideal case, this set comprises a single candidate order quantity Q^* that with confidence probability α represents an optimal solution to the problem and has a cost comprised in the interval $(G_{Q^*}(q_{Q^*,\min}^*), G_{Q^*}(q_{Q^*,\max}^*))$.

Finally, consider the case in which unobserved lost sales occurred and the M observed past demand data, d_1, \dots, d_M , only reflect the number of customers

that purchased an item when the inventory was positive. The analysis discussed above can still be applied provided that the confidence interval for the unknown parameter q of the $\text{bin}(N; q)$ demand is computed as

$$\begin{aligned} q_{lb} &= \min\{q \mid \Pr\{\text{bin}(\sum_{j=1}^M \widehat{N}_j; q) \geq X\} \geq (1 - \alpha)/2\}, \\ q_{ub} &= \max\{q \mid \Pr\{\text{bin}(\sum_{j=1}^M \widehat{N}_j; q) \leq X\} \geq (1 - \alpha)/2\}, \end{aligned}$$

where \widehat{N}_j is the total number of customers that entered the shop in day j — for which a demand sample d_j is available — while the inventory was positive.

4.3 Algorithm

The procedure to compute, under the prescribed confidence probability α , a candidate set \mathcal{Q} of order quantities and upper (c_{ub}^*) and lower (c_{lb}^*) bounds for the cost a manager faces when he selects one of these quantities is presented in Algorithm 1. The code initially computes Clopper-Pearson interval (q_{lb}, q_{ub}) by exploiting the relationship between the binomial distribution and the Beta distribution [Forbes et al., 2000] — `InverseCDF` denotes the inverse cumulative distribution function. Then it computes the critical fractile β and the upper and lower bound for the set \mathcal{Q} of candidate order quantities. Finally, it iterates through the elements of this set to compute the upper (c_{ub}^*) and lower bound (c_{lb}^*) for the estimated cost associated with these candidate order quantities.

In general, the set $\mathcal{Q} = \{Q_{lb}^*, \dots, Q_{ub}^*\}$ may comprise a significant number of elements, especially if a very limited number of samples is available. A decision maker may then employ one of the strategies discussed in Section 3.1 in order to determine the most promising quantity in this set.

4.4 Example

We consider a simple example involving the Newsvendor problem under binomial demand. Assume that, in our problem, $h = 1$, $p = 3$, and the demand follows a $\text{bin}(50, q)$ distribution, in which parameter q is unknown. We are given 10 samples for the demand, which we may use to determine the optimal order quantity Q^* . The samples are $\{28, 28, 24, 27, 25, 26, 28, 28, 23, 27\}$. The real value for parameter q , which is used to generate the 10 samples is 0.5. Accordingly, the optimal order quantity Q^* is equal to 27 and provides a cost equal to 4.4946.

We consider $\alpha = 0.9$. By using Algorithm 1 we compute the set of candidate order quantities $\mathcal{Q} = \{27, 28, \dots, 31\}$ and the confidence interval for the estimated cost $(c_{lb}^*, c_{ub}^*) = (4.4268, 7.2205)$. Among the candidate quantities in \mathcal{Q} ,

Algorithm 1: Newsvendor under incomplete information: binomial demand.

input : M past demand realizations $d_i, i = 1, \dots, M$;
the number of customers per day: N ;
the holding cost: h ;
the penalty cost: p ;
the confidence probability: α .

output: the set \mathcal{Q} of candidate order quantities;
the interval (c_{lb}^*, c_{ub}^*) for the estimated cost.

begin

```

   $a \leftarrow \sum_{i=1, \dots, M} d_i$ ;
   $b \leftarrow MN - \sum_{i=1, \dots, M} d_i$ ;
   $q_{ub} \leftarrow \text{InverseCDF}[\text{BetaDistribution}(a+1, b), (1+\alpha)/2]$ ;
   $q_{lb} \leftarrow \text{InverseCDF}[\text{BetaDistribution}(a, b+1), (1-\alpha)/2]$ ;
   $\beta \leftarrow p/(p+h)$ ;
   $Q_{ub}^* \leftarrow \text{InverseCDF}[\text{BinomialDistribution}(N, q_{ub}), \beta]$ ;
   $Q_{lb}^* \leftarrow \text{InverseCDF}[\text{BinomialDistribution}(N, q_{lb}), \beta]$ ;
   $\mathcal{Q} \leftarrow \{Q_{lb}^*, \dots, Q_{ub}^*\}$ ;
   $c_{ub}^* \leftarrow -\infty$ ;
   $c_{lb}^* \leftarrow \infty$ ;
  for each  $Q \in \mathcal{Q}$  do
     $q_{Q, \max}^* \leftarrow \arg \max_{q \in (q_{lb}, q_{ub})} G_Q(q)$ ;
     $q_{Q, \min}^* \leftarrow \arg \min_{q \in (q_{lb}, q_{ub})} G_Q(q)$ ;
     $c_{ub}^* \leftarrow \max(c_{ub}^*, G_Q(q_{Q, \max}^*))$ ;
     $c_{lb}^* \leftarrow \min(c_{lb}^*, G_Q(q_{Q, \min}^*))$ ;

```

both the strategies presented in Section 3.1 identify $\hat{Q}^* = 29$ as the candidate optimal quantity. By using the approach discussed in Section 4.2, we compute the α confidence interval for the estimated cost, which is

$$\left(G_{\hat{Q}^*}(q_{\hat{Q}^*, \min}^*), G_{\hat{Q}^*}(q_{\hat{Q}^*, \max}^*) \right) = (4.4487, 4.9528).$$

Clearly, the information on the minimum and maximum cost associated with each order quantity in \mathcal{Q} lets the decision maker perform a more educated choice. For instance, if a manager is not a risk-taker, he may decide select the order quantity \hat{Q}^* , for which the α confidence interval for the estimated cost has the lowest possible upper bound $G_{\hat{Q}^*}(q_{\hat{Q}^*, \max}^*)$. In the above example, this is still 29, but in general it may be a different order quantity.

Less conservative, but approximate, confidence intervals may be obtained by replacing the Clopper-Pearson [Clopper and Pearson, 1934] interval with

the Agresti-Coull [Agresti and Coull, 1998] interval for the binomial parameter. The maximum likelihood and the Bayesian approach do not employ confidence intervals for selecting the candidate order quantity. Therefore they are not affected by this choice and the selected order quantity remains $\hat{Q}^* = 29$. The α confidence interval for the estimated cost associated with $\hat{Q}^* = 29$ is $(G_{\hat{Q}^*}(q_{\hat{Q}^*, \min}^*), G_{\hat{Q}^*}(q_{\hat{Q}^*, \max}^*)) = (4.4487, 4.9155)$. This interval is 7.3% smaller than that produced by using the Clopper-Pearson interval.

5 Poisson demand

In many practical contexts, a random demand distributed according to a Poisson law may become relevant. A random demand $\text{Poisson}(\lambda)$ is said to be distributed according to a Poisson law with rate parameter $\lambda > 0$, if its probability mass function is

$$\Pr\{d = k\} = e^{-\lambda} \frac{\lambda^k}{k!}, \quad k = 0, 1, 2, \dots, \infty.$$

The Poisson distribution is the limiting distribution of the binomial distribution when N is large and q is small. In this case, the parameters of the two distributions are linked by the relationship $\lambda = qN$. We recall that the expected value of d is λ and that the standard deviation of d is $\sqrt{\lambda}$.

By using Eq. 3, we easily obtain the optimal order quantity Q for a given demand d . We shall give an example. Consider a demand d that follows a $\text{Poisson}(50)$ distribution. Let $h = 1$ and $p = 3$, therefore $\beta = 0.75$. The optimal order quantity is $Q^* = 55$. Furthermore, by noting that

$$G(Q) = h(Q - \lambda) + (h + p) \sum_{i=Q}^{\infty} (1 - \Pr\{\text{Poisson}(\lambda) \leq i\}),$$

the optimal cost is $G(Q^*) = 9.1222$.

We shall now consider, also in this case, the situation in which the parameter λ is not known. Instead, the decision maker is given a set of M past realizations of d . As in the previous case, from these realizations he has to infer the range of “candidate” optimal order quantities and upper and lower bounds for the cost associated with these quantities. This range will contain the real optimum according to a prescribed confidence probability.

5.1 Confidence intervals for the Poisson distribution

As in the previous case, we discuss the exact confidence interval that can be used to estimate the rate parameter λ of the Poisson distribution. This confidence

interval was proposed by Garwood [1936] and takes the following form. Consider a set of M samples d_i drawn from a random demand d that is distributed according to a Poisson law with unknown parameter λ . We rewrite $\bar{d} = \sum_{i=0}^M d_i$. According to Garwood [1936], the confidence interval for λ is $(\lambda_{lb}, \lambda_{ub})$, where

$$\begin{aligned}\lambda_{lb} &= \min\{\lambda \mid \Pr\{\text{Poisson}(M\lambda) \geq \bar{d}\} \geq (1 - \alpha)/2\}, \\ \lambda_{ub} &= \max\{\lambda \mid \Pr\{\text{Poisson}(M\lambda) \leq \bar{d}\} \geq (1 - \alpha)/2\}.\end{aligned}$$

This interval can be expressed in terms of the chi-square distribution, as shown by Garwood [1936]. Let χ_n^2 denote the chi-square distribution with n degrees of freedom, and $G^{-1}(\chi_n^2, \cdot)$ denote the inverse cumulative distribution function of χ_n^2 . We can write

$$\begin{aligned}\lambda_{lb} &= \frac{G^{-1}(\chi_{2\bar{d}}^2, (1 - \alpha)/2)}{2M}, \\ \lambda_{ub} &= \frac{G^{-1}(\chi_{2\bar{d}+2}^2, (1 + \alpha)/2)}{2M}.\end{aligned}$$

Furthermore, it is possible to express this interval using quantiles from the gamma distribution [Swift, 2009]. More specifically, the lower endpoint is the $(1 - \alpha)/2$ -quantile of a gamma distribution with shape parameter \bar{d} and scale parameter $1/M$, and the upper endpoint is the $(1 + \alpha)/2$ -quantile of a gamma distribution with shape parameter $\bar{d} + 1$ and scale parameter $1/M$. Swift lists a number of existing approaches for building approximate intervals that are less conservative than Garwood's one and he also suggests strategies to shorten Garwood's interval by choosing suitable asymmetric tails [Swift, 2009].

5.2 Solution method employing statistical estimation based on classical theory of probability

The method for computing an upper and a lower bound for the optimal order quantity Q^* in a Newsvendor problem under Poisson demand and partial information on parameter λ can be carried out in a similar fashion to the binomial case given in Section 4.2. Consider Garwood's confidence interval $(\lambda_{lb}, \lambda_{ub})$ for the unknown parameter λ of the Poisson demand. Let Q_{lb}^* be the optimal order quantity for the Newsvendor problem under a $\text{Poisson}(\lambda_{lb})$ demand and Q_{ub}^* be the optimal order quantity for the Newsvendor problem under a $\text{Poisson}(\lambda_{ub})$ demand. With confidence probability α the optimal order quantity Q^* is a member of the set $\{Q_{lb}^*, \dots, Q_{ub}^*\}$.

Consider the cost associated with an order quantity Q ,

$$G(Q) = h \sum_{i=0}^Q \Pr\{\text{Poisson}(\lambda) = i\}(Q - i) + p \sum_{i=Q}^{\infty} \Pr\{\text{Poisson}(\lambda) = i\}(i - Q).$$

Also in this case we can prove that $G_Q(\lambda)$

$$G_Q(\lambda) = h \sum_{i=0}^Q \Pr\{\text{Poisson}(\lambda) = i\}(Q-i) + p \sum_{i=Q}^{\infty} \Pr\{\text{Poisson}(\lambda) = i\}(i-Q), \quad (5)$$

is convex in λ . The proof is given in Appendix 8 (Theorem 4). Therefore upper (c_{ub}) and lower (c_{lb}) bounds for the cost associated with a solution that sets the order quantity to a value in the set $\{Q_{lb}^*, \dots, Q_{ub}^*\}$ can be easily obtained by using convex optimization approaches to find the λ^* that minimizes or maximizes this function over a given interval.

Also in this case, consider the case in which unobserved lost sales occurred and the M observed past demand data, d_1, \dots, d_M , only reflect the number of customers that purchased an item when the inventory was positive. The analysis discussed above can still be applied provided that the confidence interval for the unknown parameter λ of the $\text{Poisson}(\lambda)$ demand is computed as

$$\begin{aligned} \lambda_{lb} &= \min\{\lambda \mid \Pr\{\text{Poisson}(\widehat{M}\lambda) \geq \bar{d}\} \geq (1 - \alpha)/2\}, \\ \lambda_{ub} &= \max\{\lambda \mid \Pr\{\text{Poisson}(\widehat{M}\lambda) \leq \bar{d}\} \geq (1 - \alpha)/2\}. \end{aligned}$$

where $\widehat{M} = \sum_{j=1}^M T_j$, and $T_j \in (0, 1)$ denotes the fraction of time in day j — for which a demand sample d_j is available — during which the inventory was positive.

5.3 Algorithm

The computational procedure for Poisson demand is presented in Algorithm 2. The code initially computes Garwood's interval $(\lambda_{lb}, \lambda_{ub})$ by exploiting the relationship between the Poisson distribution and the gamma distribution [Swift, 2009]. Then it computes the critical fractile β and the upper and lower bound for the set \mathcal{Q} of candidate order quantities. Finally, it iterates through the elements of this set to compute the upper (c_{ub}^*) and lower bound (c_{lb}^*) for the estimated cost associated with these candidate order quantities.

5.4 Example

We consider a simple example involving the Newsvendor problem under Poisson demand. In our problem, $h = 1$, $p = 3$, and the demand follows a $\text{Poisson}(\lambda)$ distribution, in which parameter λ is unknown. We are given 10 samples for the demand, which we may use to determine the optimal order quantity Q^* ; these are $\{51, 54, 50, 45, 52, 39, 52, 54, 50, 40\}$. The real value for parameter λ , which

Algorithm 2: Newsvendor under incomplete information: poisson demand.

input : M past demand realizations $d_i, i = 1, \dots, M$;
the holding cost: h ;
the penalty cost: p ;
the confidence probability: α .

output: the set \mathcal{Q} of candidate order quantities;
the interval (c_{lb}^*, c_{ub}^*) for the estimated cost.

begin

```

 $a \leftarrow \sum_{i=1, \dots, M} d_i$ ;
 $b \leftarrow M$ ;
 $\lambda_{ub} \leftarrow \text{InverseCDF}[\text{GammaDistribution}(a + 1, 1/b), (1 + \alpha)/2]$ ;
 $\lambda_{lb} \leftarrow \text{InverseCDF}[\text{GammaDistribution}(a, 1/b), (1 - \alpha)/2]$ ;
 $\beta \leftarrow p/(p + h)$ ;
 $Q_{ub}^* \leftarrow \text{InverseCDF}[\text{PoissonDistribution}(\lambda_{ub}), \beta]$ ;
 $Q_{lb}^* \leftarrow \text{InverseCDF}[\text{PoissonDistribution}(\lambda_{lb}), \beta]$ ;
 $\mathcal{Q} \leftarrow \{Q_{lb}^*, \dots, Q_{ub}^*\}$ ;
 $c_{ub}^* \leftarrow -\infty$ ;
 $c_{lb}^* \leftarrow \infty$ ;
for each  $Q \in \mathcal{Q}$  do
     $\lambda_{Q, \max}^* \leftarrow \arg \max_{\lambda \in (\lambda_{lb}, \lambda_{ub})} G_Q(\lambda)$ ;
     $\lambda_{Q, \min}^* \leftarrow \arg \min_{\lambda \in (\lambda_{lb}, \lambda_{ub})} G_Q(\lambda)$ ;
     $c_{ub}^* \leftarrow \max(c_{ub}^*, G_Q(\lambda_{Q, \max}^*))$ ;
     $c_{lb}^* \leftarrow \min(c_{lb}^*, G_Q(\lambda_{Q, \min}^*))$ ;

```

is used to generate the samples, is 50. Accordingly, the optimal order quantity Q^* is equal to 55 and provides a cost equal to 9.1222.

We consider $\alpha = 0.9$. By using Algorithm 2 we compute the set of candidate order quantities $\mathcal{Q} = \{50, 51, \dots, 57\}$ and the confidence interval for the estimated cost $(c_{lb}^*, c_{ub}^*) = (8.6803, 14.6220)$. Let us consider a strategy based on the maximum likelihood estimator. In the case of the Poisson distribution, this estimator takes the following convenient form, $\frac{1}{M} \sum_{i=1}^M d_i$, where d_i , for $i = 1, \dots, M$ are the observed samples. Therefore, according to the above samples, the maximum likelihood estimator for λ is 48.7. By using a demand that follows a Poisson distribution with mean rate $\lambda = 48.7$ in Eq. 3 we obtain a candidate optimal order quantity $\hat{Q}^* = 53$ and an estimated expected cost of 9.0035. However, such a strategy does not provide any information on the reliability of the above estimates. In fact, the actual cost associated with this order quantity, when $\lambda = 50$, is 9.3693. Conversely, our approach reports the α confidence interval $(8.9463, 11.0800)$ for the expected cost associated with

$\widehat{Q}^* = 53$, which in this case includes the actual cost a decision maker will face in case he decides to order 53 units.

Similar issues occur for the Bayesian approach presented by Hill [1997]. For the sample set presented above, this approach suggests ordering 54 units and estimates a cost of 9.4764, but it does not provide any information on the reliability of these estimates. In contrast, for an order quantity of 54 units our approach reports the α confidence interval (9.0334, 10.3374) for the expected cost, which include the actual cost 9.1530 associated with this quantity when $\lambda = 50$.

6 Exponential demand

A random demand $\exp(\lambda)$ is said to be distributed according to an exponential law with rate parameter $\lambda > 0$ if its probability density function is

$$f(\lambda, k) = \lambda e^{-\lambda k}, \quad k \geq 0;$$

the expected value of $\exp(\lambda)$ is $1/\lambda$.

In the context of the Newsvendor, the exponential distribution may occur in two cases. An exponentially distributed random variable $\exp(\lambda)$ with rate parameter λ can represent the inter-arrival time between two unit demand occurrences in a Poisson process with rate parameter λ . Alternatively, an exponentially distributed random variable $\exp(\lambda)$ can represent the total demand over the Newsvendor planning horizon. It is clear that the first case can be easily reduced to the case of a random demand that follows a Poisson distribution with rate parameter λ . Such a situation can be handled by following the discussion in the previous section. In the second case, by using Eq. 2, we easily obtain the optimal order quantity Q^* for $\exp(\lambda)$. This is simply

$$Q^* = -\frac{1}{\lambda} \ln \left(\frac{h}{h+p} \right). \quad (6)$$

We shall give an example. Consider a random demand $\exp(1/50)$ with mean 50. Let $h = 1$ and $p = 3$, therefore $h/(h+p) = 0.25$. The optimal order quantity is $Q^* = 69.32$. Furthermore, consider the cost function

$$G(Q) = h \int_0^Q (Q-i)f(\lambda, i)di + p \int_Q^\infty (i-Q)f(\lambda, i)di,$$

where $f(\lambda, \cdot)$ denotes the probability density function of $\exp(\lambda)$. Rewrite

$$G(Q) = h(Q - \frac{1}{\lambda}) + (h+p) \int_Q^\infty (1 - F(\lambda, i))di,$$

where $F(\lambda, \cdot)$ denotes the cumulative distribution function of $\exp(\lambda)$. By noting that

$$G(Q) = \frac{h+p}{\lambda} \left(\frac{h}{h+p} (\lambda Q - 1) + e^{-\lambda Q} \right), \quad (7)$$

the optimal cost is $G(Q^*) = 69.32$.

We shall now consider, also in this case, the situation in which the parameter λ is not known and the decision maker is given a set of M past realizations of the demand. As in the previous case, from these realizations he has to infer the range of candidate optimal order quantities and upper and lower bounds for the cost associated with these quantities. This range will contain the real optimum according to a prescribed confidence probability.

6.1 Confidence intervals for the exponential distribution

We discuss the exact confidence interval that can be used to estimate the rate parameter λ of the exponential distribution. Consider a set of M samples d_i drawn from a random variable that is distributed according to an exponential law with unknown parameter λ . We rewrite $\bar{d} = \sum_{i=0}^M d_i$. Since the sum of M independent and identically distributed exponential random variables with rate parameter λ is a random variable $\text{gamma}(M, 1/\lambda)$ that follows a gamma distribution with shape parameter M and scale parameter $1/\lambda$, the α confidence interval for the unknown parameter λ is $(\lambda_{lb}, \lambda_{ub})$, where

$$\begin{aligned} \lambda_{lb} &= \min\{\lambda \mid \Pr\{\text{gamma}(M, 1/\lambda) \geq \bar{d}\} \geq (1 - \alpha)/2\}, \\ \lambda_{ub} &= \max\{\lambda \mid \Pr\{\text{gamma}(M, 1/\lambda) \leq \bar{d}\} \geq (1 - \alpha)/2\}. \end{aligned}$$

A closed form expression for this confidence interval — that employs quantiles from the χ^2 distribution — was proposed by Trivedi [2001, chap. 10] and takes the following form. Let χ_n^2 denote the chi-square distribution with n degrees of freedom, and $G^{-1}(\chi_n^2, \cdot)$ denote the inverse cumulative distribution function of χ_n^2 . We can write

$$\lambda_{lb} = \frac{G^{-1}(\chi_{2M}^2, (1 - \alpha)/2)}{2\bar{d}}, \quad \lambda_{ub} = \frac{G^{-1}(\chi_{2M}^2, (1 + \alpha)/2)}{2\bar{d}}.$$

Furthermore, it is possible to express this interval using quantiles from the gamma distribution. More specifically, the lower endpoint is the $(1 - \alpha)/2$ -quantile of a gamma distribution with shape parameter M and scale parameter $1/\bar{d}$, and the upper endpoint is the $(1 + \alpha)/2$ -quantile of a gamma distribution with shape parameter M and scale parameter $1/\bar{d}$.

6.2 Solution method employing statistical estimation based on classical theory of probability

Consider the confidence interval $(\lambda_{lb}, \lambda_{ub})$ for the unknown parameter λ of the exponential demand. Let Q_{lb}^* be the optimal order quantity for the Newsvendor problem under an $\exp(\lambda_{ub})$ demand and Q_{ub}^* be the optimal order quantity for the Newsvendor problem under an $\exp(\lambda_{lb})$ demand. Recall that λ is a rate, this is the reason why the optimal order quantity for the Newsvendor problem under an $\exp(\lambda_{lb})$ gives an *upper bound* Q_{ub}^* for the real optimal order quantity. Clearly, the optimal order quantity Q^* lies in the interval (Q_{lb}^*, Q_{ub}^*) .

Let us write the expected total cost associated with an order quantity Q for a given demand rate $\lambda > 0$,

$$G_\lambda(Q) = h(Q - \frac{1}{\lambda}) + (h + p) \int_Q^\infty (1 - F(\lambda, i)) di,$$

it is known that this function is convex. Then, consider the function

$$G_Q(\lambda) = h(Q - \frac{1}{\lambda}) + (h + p) \int_Q^\infty (1 - F(\lambda, i)) di, \quad (8)$$

in which the order quantity Q is fixed and in which we vary the demand rate $\lambda \geq 0$. Unfortunately, $G_Q(\lambda)$ is not convex in the continuous parameter λ . Nevertheless, we prove a number of properties for this function.

Theorem 1. $\lim_{\lambda \rightarrow 0} G_Q(\lambda) = \infty$, $\lim_{\lambda \rightarrow \infty} G_Q(\lambda) = hQ^-$, the function admits a single global minimum λ^* , it is strictly increasing for $\lambda > \lambda^*$ and strictly decreasing for $\lambda < \lambda^*$.

The proof is given in Appendix 8.

Because of the properties of $G_Q(\lambda)$ introduced in Theorem 1 we can employ a simple line search procedure in order to find the λ^* that minimizes or maximizes this function over a given interval.

Since the optimal order quantity Q^* is, with confidence probability α , a value in the interval (Q_{lb}^*, Q_{ub}^*) , it is therefore easy to compute upper (c_{ub}^*) and lower (c_{lb}^*) bounds for the cost that a manager will face, with confidence probability α , whatever order quantity he chooses in this interval.

Theorem 2. *The lower bound is*

$$c_{lb}^* = G_{Q_{lb}^*}(\lambda_{ub})$$

the upper bound is

$$c_{ub}^* = \max\{G_{Q_{lb}^*}(\lambda_{lb}), G_{Q_{ub}^*}(\lambda_{ub})\}.$$

The proof is given in Appendix 8.

Unlike the previous cases, it is not straightforward to extend the above reasoning to the case in which unobserved lost sales occurred and the M observed past demand data, d_1, \dots, d_M , only reflect the number of customers that purchased an item when the inventory was positive. This is due to the fact that the distribution of the general sum of exponential random variables is not exponential, rather it is Hypoexponential. We therefore leave this discussion as a future research direction.

6.3 Algorithm

The computational procedure for exponential demand is presented in Algorithm 3. The code initially computes the confidence interval $(\lambda_{lb}, \lambda_{ub})$ by exploiting the

Algorithm 3: Newsvendor under incomplete information: exponential demand.

input : M past demand realizations $d_i, i = 1, \dots, M$;
the holding cost: h ;
the penalty cost: p ;
the confidence probability: α .

output: the set \mathcal{Q} of candidate order quantities;
the interval (c_{lb}^*, c_{ub}^*) for the estimated cost.

begin

$a \leftarrow M$;
$b \leftarrow \sum_{i=1, \dots, M} d_i$;
$\lambda_{ub} \leftarrow \text{InverseCDF}[\text{GammaDistribution}(a, 1/b), (1 + \alpha)/2]$;
$\lambda_{lb} \leftarrow \text{InverseCDF}[\text{GammaDistribution}(a, 1/b), (1 - \alpha)/2]$;
$\beta \leftarrow p/(p + h)$;
$Q_{ub}^* \leftarrow \text{InverseCDF}[\text{PoissonDistribution}(\lambda_{ub}), \beta]$;
$Q_{lb}^* \leftarrow \text{InverseCDF}[\text{PoissonDistribution}(\lambda_{lb}), \beta]$;
$\mathcal{Q} \leftarrow \{Q_{lb}^*, \dots, Q_{ub}^*\}$;
$c_{ub}^* \leftarrow \max\{G_{Q_{lb}^*}(\lambda_{lb}), G_{Q_{ub}^*}(\lambda_{ub})\}$;
$c_{lb}^* \leftarrow G_{Q_{lb}^*}(\lambda_{ub})$;

relationship between the exponential distribution and the gamma distribution. Then it computes the critical fractile β and the upper and lower bound for the set \mathcal{Q} of candidate order quantities. Finally, it computes the upper (c_{ub}^*) and lower bound (c_{lb}^*) for the estimated cost associated with these candidate order quantities by exploiting Theorem 2.

6.4 Example

We consider a simple example involving the Newsvendor problem under exponential demand. In our problem, $h = 1$, $p = 3$, and the demand is a random variable $\exp(\lambda)$ for which parameter λ is unknown. We are given 10 samples for the demand, which we may use to determine the optimal order quantity Q^* ; these are

$$\{39.79, 39.26, 32.21, 0.51, 107.03, 72.87, 45.23, 20.12, 26.46, 56.80\}.$$

The real value for parameter λ , which is used to generate the samples, is $1/50$. Accordingly, the optimal order quantity Q^* is equal to 69.31 and provides a cost equal to 69.31.

We consider $\alpha = 0.9$. The α confidence interval for the demand rate λ is $(0.0123211, 0.0356664)$. By using Algorithm 3 we compute the range of candidate order quantities $Q = (38.86, 112.51)$ and the confidence interval for the estimated cost $(c_{lb}^*, c_{ub}^*) = (38.86, 158.81)$. A plot of the expected total cost as a function of $\lambda \in (0.0123211, 0.0356664)$ and of $Q \in (38.86, 112.51)$ is shown in Appendix 8. Let us consider a strategy based on the maximum likelihood estimator for the demand rate λ . In the case of the exponential distribution, this estimator takes the following form, $\hat{\lambda} = \frac{N}{\sum_{i=1}^N d_i}$, where d_i , for $i = 1, \dots, N$ are the observed samples. Therefore, according to the above samples, the maximum likelihood estimator for λ is $\hat{\lambda} = 0.0227099$. By using a rate $\lambda = 0.0227099$ in Eq. 3 we obtain a candidate optimal order quantity $\hat{Q}^* = 61.04$ and an estimated expected cost of 61.04. As previously remarked, such a strategy does not provide any information on the reliability of the above estimates. In fact, the actual cost associated with this order quantity, when $\lambda = 1/50$, is 70.03. Conversely, our approach reports the α confidence interval $(45.71, 132.90)$ for the expected cost associated with $\hat{Q}^* = 61.04$, which in this case includes the actual cost a decision maker will face in case he decides to order 61.04 units.

Similarly, the Bayesian approach presented by Hill [1997] suggests an order quantity of 59.14 units and estimates a cost of 65.05. As discussed, this strategy does not provide any information on the reliability of these estimates. Conversely, for an order quantity of 59.14 units, our approach reports the α confidence interval $(44.71, 134.63)$ for the expected cost, which includes the actual cost 70.42 associated with this quantity when $\lambda = 1/50$.

7 Discussion and future works

In this section we first discuss advantages of our strategy, based on Neyman's method of confidence intervals, with respect to existing frequentist and Bayesian

approaches to the Newsvendor problem under sampled demand information. Secondly, we discuss limitations of our work and possible future research directions.

7.1 Comparison with frequentist and Bayesian approaches

Bayesian approaches such as the one proposed by Hill [1997] present several theoretical and practical drawbacks that we are now going to enumerate. Theoretical drawbacks of the Bayesian approach to parameter estimation are illustrated by Neyman [1937, p. 343]. The first issue raised by Neyman is the fact that the unknown parameter a of the demand distribution is not a random variable, therefore assigning a prior or posterior distribution to it has no meaning. In fact, one may try to employ the prior probability distribution — for instance a uniform distribution in Hill’s case — to compute $\Pr\{a < \Delta\}$, where $\Delta \in (0, 1)$. Of course, a is not a random variable therefore this probability should be either 1 or 0 depending on Δ . Therefore, talking about prior or posterior distribution for a can only represent an approximation. It is often stated that in Bayesian probability prior and posterior distributions are meant to represent a “state of knowledge”, that is decision maker’s uncertainty about the unknown quantity a . However, problems arise immediately as soon as one tries to interpret the meaning of the prior and posterior distribution in light of classical probability theory. Neyman, in fact, also points out that, even if the unknown parameter a is a random variable, the posterior distribution $f(d|b)$ for the random demand d , computed as illustrated in Hill [1997], does not generally have the property serving as a definition of the elementary probability law of the observed data b . In particular, this distribution is not compatible with the classical definition of probability, in the sense that, if we repeat an experiment an infinite number of times, the observed frequency does not converge to the probability predicted by such a distribution. For instance, consider once more the example presented in Section 4.4. Instead of having a single set of demand observations, we now consider M experiments, with M large, in each of which we observe 10 demand realizations. Let b_i be the demand set observed in experiment $i = 1, \dots, M$. For each experiment i , we construct the posterior distribution, $f_i(d_i|b_i)$, of the random demand d_i from 10 demand observations, according to the Bayesian strategy discussed by Hill [1997]. Neyman points out that, in general, when we select two values Δ_1 and Δ_2 , and we compute $p_i = \Pr\{\Delta_1 < d_i < \Delta_2\}$, the quantity $1/M \sum_{i=1}^M p_i$ will not converge — as it should, according to the law of large numbers — to its real value $\Pr\{\Delta_1 < \text{bin}(50; q) < \Delta_2\}$, where q is the real value of a . In our example, we set $\Delta_1 = 26$ and $\Delta_2 = 28$. $\Pr\{\Delta_1 < \text{bin}(50; 0.5) < \Delta_2\} = 0.1747$. Nevertheless, if we estimate $p_i = \Pr\{\Delta_1 < d_i < \Delta_2\}$ in each experiment i by using the posterior distribution $f(d_i|b)$ discussed by Hill [1997] for a binomial demand, the estimated probability $1/M \sum_{i=1}^M p_i$ converges to 0.2654, when the experiment is repeated $M = 100000$ times. This essentially differs from the real value, as

Neyman remarks. For this reason, using the posterior distribution in place of the original distribution in the problem of interest is a strategy that may lead to misleading results.

In practice, a direct consequence is that, although for large samples asymptotic results may be obtained (see e.g. Bensoussan et al. [2009]), it becomes hard to assess the quality of estimates produced for small sample set. It is fairly simple to observe this latter fact by considering once more the example presented in Section 4.4. By using the approach in Hill [1997], we obtain the posterior distribution for the random demand out of the 10 samples, and then we use this distribution in order to compute an estimate of Q^* , which in the particular example we consider is 29. We also employ the posterior distribution in the Newsvendor cost function in order to estimate the cost associated with the optimal order quantity selected; this turns out to be 4.6692. Clearly, Hill's approach does not give any hint on the "quality" of the estimates produced, which in this particular case are relatively poor. In particular, based on the available data, we do not know with what frequency the order quantity may substantially over- or underestimate the real optimum order quantity, and how far the prescribed order quantity is likely to be from the real optimum order quantity. The same, of course, holds for the estimated optimum cost. Our approach based on Neyman's framework, in contrast, suggests that, with 90% confidence, the optimal order quantity — that is 27 — lies between 27 and 31, and that the optimum cost — that is 4.4946 — lies in the interval (4.4268, 7.2205). Then, when a given heuristic suggests ordering 29 units — for instance according to Hill's Bayesian approach — our approach can be used to derive a 90% confidence interval for the cost associated with this decision, that is (4.4487, 4.9528). This interval actually covers, in this specific case, the real cost associated with the decision of ordering 29 units, i.e. 4.8904. In general, the interval will cover the real cost according to the prescribed confidence probability. Conversely, Hill's approach suggests that the cost associated with ordering 29 units is 4.6692, but it provides no indication on the reliability of this estimate. This shows a practical exemplification of how our approach can be employed to effectively complement existing Bayesian approaches under small sample sets. In fact, we must also underscore the fact that Bayesian approaches such as the one in [Hill, 1997] represent very effective and practical heuristics for order quantity selection.

Similar issues arise in classical frequentist approaches. For the example in Section 4.4, an approach based on the maximum likelihood estimator, as remarked, suggests an order quantity of 29 units. The estimated cost according to this strategy is 4.4614. Nevertheless, also in this case we have no indication on the reliability of this estimate. Kevork [2010] derives maximum likelihood estimators for the optimal order quantity and for the maximum expected profit. The asymptotic distribution for these estimators are then derived and asymptotic confidence intervals are extracted for the corresponding true quantities. Unfortunately, as the author remarks, these intervals are only asymptotically

exact and do not provide the prescribed confidence, i.e. they are biased, for small sample sets.

By using our approach based on Neyman’s framework, the confidence interval produced for the unknown parameter of the binomial demand is always guaranteed to cover its actual value according to the prescribed confidence probability. This probability is controlled by the decision maker and influences the size of the interval. Neyman’s method differs advantageously from the Bayesian approach by being independent of a priori information about the unknown parameter. This approach remains valid even if the unknown parameter is a random variable. By using our approach, it is possible to immediately translate the confidence interval for the unknown parameter into a confidence interval for the order quantity and for the actual cost. Intuitively, in the example presented in Section 4.4, if we observe a set of 10 samples over and over again and we repeat our analysis, the intervals produced will cover the real optimum order quantity and the associated cost according to the prescribed probability. In contrast to other existing frequentist or Bayesian approaches our approach provides explicit and exact likelihood guarantees that can be easily interpreted in the context of classical probability theory. Furthermore, by using confidence intervals, the decision maker has a better control on the risk of exceeding a certain cost and a better outlook on the range of order quantities that may be optimal according to the observed demands, especially when a limited set of samples is employed.

7.2 Limitations and future works

Our analysis is limited to three maximum entropy probability distributions in the exponential family [Andersen, 1970], each of which features a single parameter that must be estimated. As shown by Harremoes [2001], the binomial and the Poisson are maximum entropy probability distributions for the case in which all we know about the distribution of a random demand is that it has positive mean and discrete support that goes from 0 to a maximum value N (binomial) or to infinity (Poisson). The exponential distribution is the maximum entropy probability distribution for the case in which all we know about the distribution of a random demand is that it has positive mean and continuous support that goes from 0 to infinity. These considerations show how broadly applicable the results in this work are. In this work, the normal distribution — which is part of the exponential family and which is also a maximum entropy probability distribution — has not been considered. The analysis on the normal distribution is complicated by the fact that two parameters, mean and variance, must be considered. Then a number of cases naturally arise: unknown mean and known variance, unknown variance and known mean, etc. For this reason, in order to keep the size and the scope of the discussion limited, we decided to leave this discussion as a future work. Furthermore, in principle it may be possible to extend the analysis to other distributions such as the multinomial, for which confi-

dence intervals are surveyed in [Lee et al., 2002, Chafai and Concordet, 2009]; or the Johnson translation system [Johnson, 1949], if exact or approximate expressions for the confidence regions of its unknown parameters were available. Unfortunately, we are not aware of any work that investigated these confidence regions.

8 Conclusions

We considered the problem of controlling the inventory of a single item with stochastic demand over a single period. We introduced a novel strategy to address the issue of demand estimation in single-period inventory optimization problems. Our strategy is based on the theory of statistical estimation. We employed confidence interval analysis in order to identify a range of candidate order quantities that, with prescribed confidence probability, includes the real optimal order quantity for the underlying stochastic demand process with unknown parameter(s). In addition, for each candidate order quantity that is identified, our approach can produce an upper and a lower bound for the associated cost. We applied our novel approach to three demand distribution in the exponential family: binomial, Poisson, and exponential. For two of these distributions we also discussed the case in which the decision maker faces unobserved lost sales. Numerical examples are presented in which we showed how our approach complements existing strategies based on maximum likelihood estimators or on Bayesian analysis. In particular, we showed that our approach does not provide a single order quantity recommendation and a point estimate for the associated cost, but — according to a prescribed confidence level — a set of candidate optimal order quantities and, for each of these, a confidence interval for the associated cost. This advanced information can be employed, together with existing frequentist or Bayesian approaches, to better assess the impact of a given decision.

Appendix I: proofs of statements for Binomial demand

Consider Eq. 4 in Section 4.2, it can be proved that $G_Q(q)$ is convex in the continuous parameter q . Firstly, we rewrite Eq. 4 as

$$G_Q(q) = h(Q - Nq) + (p + h) \sum_{i=Q}^N (1 - \Pr\{\text{bin}(N; q) \leq i\}). \quad (9)$$

We now show that the second derivative of this function is positive. Of course, this is equivalent to proving that

$$\frac{d^2}{dq^2}(p+h) \sum_{i=Q}^N (1 - \Pr\{\text{bin}(N; q) \leq i\}) \geq 0.$$

Theorem 3. For $Q \leq N$,

$$\frac{d^2}{dq^2} \sum_{i=Q}^N (1 - \Pr\{\text{bin}(N; q) \leq i\})$$

is a positive function of $q \in (0, 1)$.

Proof (Theorem 3). We introduce the following notation:

$$\begin{aligned} f(i; N, q) &= \Pr\{\text{Bin}(N; q) = i\} = \binom{N}{i} q^i (1-q)^{N-i}, \\ F(i; N, q) &= \Pr\{\text{Bin}(N; q) \leq i\}. \end{aligned}$$

We reduce the convexity of $G_Q(q)$ to showing that

$$\frac{d^2}{dq^2} \sum_{i=Q}^N F(i; N, q) \leq 0.$$

Using the regularized incomplete beta function:

$$F(i; N, q) = (N-i) \binom{N}{i} \int_0^{1-q} t^{N-i-1} (1-t)^i dt;$$

differentiating under the integral sign by Leibniz's rule:

$$\begin{aligned} \frac{d}{dq} F(i; N, q) &= -(N-i) \binom{N}{i} (1-q)^{N-i-1} q^i \\ &= -N \binom{N-1}{i} q^i (1-q)^{N-i-1} \\ &= -N f(i; N-1, q) \\ &= -N [F(i; N-1, q) - F(i-1; N-1, q)]; \end{aligned}$$

using this recursive relationship:

$$\begin{aligned} \frac{d^2}{dq^2} F(i; N, q) &= -N [-(N-1)(F(i; N-2, q) - F(i-1; N-2, q)) + \\ &\quad (N-1)(F(i-1; N-2, q) - F(i-2; N-2, q))] \\ &= N(N-1)[f(i; N-2, q) - f(i-1; N-2, q)]; \end{aligned}$$

and summing over i , all terms cancel out except the first and last:

$$\frac{d^2}{dq^2} \sum_{i=Q}^N F(i; N, q) = N(N-1)[f(N; N-2, q) - f(Q-1; N-2, q)].$$

However, $f(N; N-2, q) = 0$ because it represents the probability of N successes in $N-2$ trials, so

$$\frac{d^2}{dq^2} \sum_{i=Q}^N F(i; N, q) = -N(N-1)f(Q-1; N-2, q) \leq 0$$

□

Appendix II: proofs of statements for Poisson demand

Consider Eq. 5 in Section 5.2, it can be proved that $G_Q(\lambda)$ is convex in the continuous parameter λ . Firstly, we rewrite Eq. 5 as

$$G_Q(\lambda) = h(Q - \lambda) + (h + p) \sum_{i=Q}^{\infty} (1 - \Pr\{\text{Poisson}(\lambda) \leq i\}). \quad (10)$$

We now show that the second derivative of this function is positive. Of course, this is equivalent to proving that

$$\frac{d^2}{d\lambda^2} (h + p) \sum_{i=Q}^{\infty} (1 - \Pr\{\text{Poisson}(\lambda) \leq i\}) \geq 0.$$

Therefore, we have to prove that

$$\frac{d^2}{d\lambda^2} - \sum_{i=Q}^{\infty} \Pr\{\text{Poisson}(\lambda) \leq i\} \geq 0.$$

Theorem 4. For $Q \geq 0$,

$$\frac{d^2}{d\lambda^2} - \sum_{i=Q}^{\infty} \Pr\{\text{Poisson}(\lambda) \leq i\}$$

is a positive function of $\lambda \geq 0$.

Proof (Theorem 4). The following derivations prove convexity for the above expression.

$$\begin{aligned} & \frac{d^2}{d\lambda^2} - \sum_{i=Q}^{\infty} \Pr\{\text{Poisson}(\lambda) \leq i\} = \\ & \frac{d^2}{d\lambda^2} - \sum_{i=Q}^{\infty} e^{-\lambda} \sum_{k=0}^i \frac{\lambda^k}{k!} = \\ & - \left(\frac{d^2}{d\lambda^2} e^{-\lambda} \sum_{i=Q}^{\infty} \sum_{k=0}^i \frac{\lambda^k}{k!} \right) = \\ & - \left(\frac{d}{d\lambda} - e^{-\lambda} \sum_{i=Q}^{\infty} \sum_{k=0}^i \frac{\lambda^k}{k!} + \frac{d}{d\lambda} e^{-\lambda} \sum_{i=Q}^{\infty} \sum_{k=1}^i \frac{\lambda^{k-1}}{(k-1)!} \right) = \\ & - \left(e^{-\lambda} \sum_{i=Q}^{\infty} \sum_{k=0}^i \frac{\lambda^k}{k!} - e^{-\lambda} \sum_{i=Q}^{\infty} \sum_{k=1}^i \frac{\lambda^{k-1}}{(k-1)!} - e^{-\lambda} \sum_{i=Q}^{\infty} \sum_{k=1}^i \frac{\lambda^{k-1}}{(k-1)!} + e^{-\lambda} \sum_{i=Q}^{\infty} \sum_{k=2}^i \frac{\lambda^{k-2}}{(k-2)!} \right) = \end{aligned}$$

$$e^{-\lambda} \left(-\sum_{i=Q}^{\infty} \sum_{k=0}^i \frac{\lambda^k}{k!} + \sum_{i=Q}^{\infty} \sum_{k=1}^i \frac{\lambda^{k-1}}{(k-1)!} + \sum_{i=Q}^{\infty} \sum_{k=1}^i \frac{\lambda^{k-1}}{(k-1)!} - \sum_{i=Q}^{\infty} \sum_{k=2}^i \frac{\lambda^{k-2}}{(k-2)!} \right) =$$

$$\sum_{i=Q}^{\infty} \left(-e^{-\lambda} \sum_{k=0}^i \frac{\lambda^k}{k!} + 2e^{-\lambda} \sum_{k=1}^i \frac{\lambda^{k-1}}{(k-1)!} - e^{-\lambda} \sum_{k=2}^i \frac{\lambda^{k-2}}{(k-2)!} \right).$$

For convenience, we rewrite this expression as

$$\sum_{i=Q}^{\infty} (-\text{CDF}(\text{Poisson}(\lambda), i) + 2\text{CDF}(\text{Poisson}(\lambda), i-1) - \text{CDF}(\text{Poisson}(\lambda), i-2))$$

where CDF denotes the cumulative distribution function. By expanding, we obtain

$$\begin{aligned} & -\text{CDF}(\text{Poisson}(\lambda), Q) + 2\text{CDF}(\text{Poisson}(\lambda), Q-1) - \text{CDF}(\text{Poisson}(\lambda), Q-2) + \\ & -\text{CDF}(\text{Poisson}(\lambda), Q+1) + 2\text{CDF}(\text{Poisson}(\lambda), Q) - \text{CDF}(\text{Poisson}(\lambda), Q-1) + \\ & -\text{CDF}(\text{Poisson}(\lambda), Q+2) + 2\text{CDF}(\text{Poisson}(\lambda), Q+1) - \text{CDF}(\text{Poisson}(\lambda), Q) + \dots = \end{aligned}$$

$$\text{CDF}(\text{Poisson}(\lambda), Q-1) - \text{CDF}(\text{Poisson}(\lambda), Q-2) =$$

$$\begin{aligned} & e^{-\lambda} \sum_{k=0}^{Q-1} \frac{\lambda^k}{k!} - e^{-\lambda} \sum_{k=0}^{Q-2} \frac{\lambda^k}{k!} = \\ & e^{-\lambda} \left(\sum_{k=0}^{Q-1} \frac{\lambda^k}{k!} - \sum_{k=0}^{Q-2} \frac{\lambda^k}{k!} \right) = \\ & \frac{e^{-\lambda} \lambda^{Q-1}}{(Q-1)!} \geq 0 \end{aligned}$$

□

Appendix III: proofs of statements for exponential demand

In this section we provide the proofs for the two theorems introduced in Section 6.2.

Proof (Theorem 1). *The first fact, that is $\lim_{\lambda \rightarrow 0} G_Q(\lambda) = \infty$ can be easily verified by simple algebraic derivations. We shall therefore prove that $\lim_{\lambda \rightarrow \infty} G_Q(\lambda) = hQ^-$ and that the function admits a single global minimum. Let us split Eq. 7 into two parts*

$$G_Q(\lambda) = \frac{h+p}{\lambda} e^{-\lambda Q} + \left(hQ - \frac{h}{\lambda} \right). \quad (11)$$

We shall consider the first term

$$\frac{h+p}{\lambda} e^{-\lambda Q}, \quad (12)$$

and the second term

$$hQ - \frac{h}{\lambda}, \quad (13)$$

on the right hand side of Eq. 7, separately.

Firstly, we observe that, when $\lambda \rightarrow \infty$, $G_Q(\lambda) \rightarrow hQ$ from below, that is $\lim_{\lambda \rightarrow \infty} G_Q(\lambda) = hQ^-$. This is due to the fact that (12) approaches zero faster than h/λ does, i.e.

$$\lim_{\lambda \rightarrow \infty} \frac{\frac{h+p}{\lambda} e^{-\lambda Q}}{\frac{h}{\lambda}} = 0.$$

From this fact we immediately infer that the derivative of $G_Q(\lambda)$ must be equal to zero for at least one value λ other than infinity. Furthermore, the derivative of (12) is negative, strictly increasing for $\lambda > 0$. The derivative of (13) is positive strictly decreasing for $\lambda > 0$. Therefore there exists only a single value of λ for which the derivative of (12) and the derivative of (13) add up to zero. This immediately implies that $G_Q(\lambda)$ admits a single global minimum, it is strictly increasing for $\lambda > \lambda^*$ and strictly decreasing for $\lambda < \lambda^*$ \square

Proof (Theorem 2). Firstly, let us consider c_{lb}^* . By definition, this is the expected total cost associated with the optimal order quantity Q_{lb}^* for the largest possible value λ_{ub} that the demand rate takes in the confidence interval. Consider a demand rate λ and the associated optimal order quantity Q_λ^* . By substituting Q in Eq. 7 with the expression of the optimal order quantity in Eq. 6 we immediately see that the expected total cost associated with an optimal order quantity Q_λ^* is decreasing in the respective demand rate λ — i.e. it is increasing w.r.t. the expected value $1/\lambda$ of the demand — it immediately follows that there exists no other pair $\langle Q_\lambda^*, \lambda \rangle$, where $\lambda \in (\lambda_{lb}, \lambda_{ub})$ that ensures a lower expected total cost.

Let $\lambda \in (\lambda_{lb}, \lambda_{ub})$ and $Q \in (Q_{lb}, Q_{ub})$. Consider a point in the two dimensional space $\lambda \times Q$, for which $\lambda = \bar{\lambda}$ and $Q = \bar{Q}$. For any of such points, two cases can be observed, that is (i) $\bar{Q} > Q_{\bar{\lambda}}^*$, or (ii) $\bar{Q} < Q_{\bar{\lambda}}^*$. A strict equality can be reduced to any of these two cases. If we are in case (i), then $G_{\bar{\lambda}}(\bar{Q}) < G_{\bar{\lambda}}(Q_{ub})$ because of Theorem 1. \bar{Q} was already an order quantity larger than the optimal one, therefore Q_{ub} is also an order quantity larger than the optimal one for a demand rate $\bar{\lambda}$. Consequently, if we increase the demand rate λ (i.e. we decrease the expected demand $1/\lambda$) our cost can only increase; this means that $G_{\bar{\lambda}}(Q_{ub}) < G_{\lambda_{ub}}(Q_{ub})$. If we are in case (ii), then $G_{\bar{\lambda}}(\bar{Q}) < G_{\bar{\lambda}}(Q_{lb})$ because of Theorem 1. \bar{Q} was already an order quantity smaller than the optimal one, therefore Q_{lb} is also an order quantity smaller than the optimal one for a demand rate $\bar{\lambda}$. Consequently, if we decrease the demand rate λ (i.e. we increase the expected demand $1/\lambda$) our cost can only increase; this means that $G_{\bar{\lambda}}(Q_{lb}) < G_{\lambda_{lb}}(Q_{lb})$. Therefore, the maximum cost, when we let λ vary in $(\lambda_{lb}, \lambda_{ub})$ and Q vary in (Q_{lb}^*, Q_{ub}^*) , can be either observed at $\langle Q_{lb}^*, \lambda_{lb} \rangle$ or at $\langle Q_{ub}^*, \lambda_{ub} \rangle$ \square

Appendix IV: plot for the expected total cost of the example in Section 6.4

In Fig. 1 we provide a graphical outlook of the cost function discussed in Section 6.4.

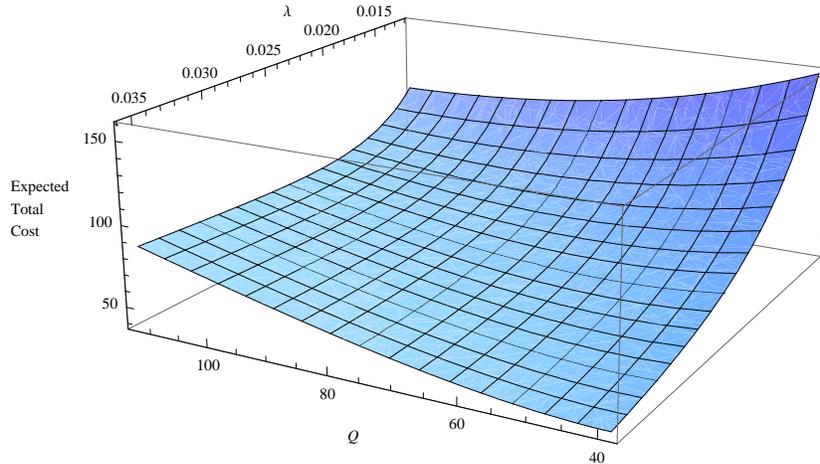


Figure 1: Expected total cost as a function of $\lambda \in (\lambda_{lb}, \lambda_{ub}) \equiv (0.0123211, 0.0356664)$ and of $Q \in (Q_{lb}^*, Q_{ub}^*) \equiv (38.86, 112.51)$. Note that $c_{lb}^* = G_{Q_{lb}^*}(\lambda_{ub})$ and that $c_{ub}^* = \max\{G_{Q_{lb}^*}(\lambda_{lb}), G_{Q_{ub}^*}(\lambda_{ub})\} = G_{Q_{lb}^*}(\lambda_{lb})$.

References

- N. Agrawal and S. A. Smith. Estimating negative binomial demand for retail inventory management with unobservable lost sales. *Naval Research Logistics*, 43(6):839–861, 1996.
- A. Agresti and B. A. Coull. Approximate is better than "exact" for interval estimation of binomial proportions. *The American Statistician*, 52(2):119–126, 1998. ISSN 00031305.
- S. Ahmed, U. Çakmak, and A. Shapiro. Coherent risk measures in inventory problems. *European Journal of Operational Research*, 182(1):226–238, October 2007. ISSN 03772217.
- A. Akcay, B. Biller, and S. Tayur. Improved inventory targets in the presence of limited historical demand data. *Manufacturing & Service Operations Management*, 13(3):297–309, June 2011. ISSN 1526-5498.
- E. B. Andersen. Sufficiency and exponential families for discrete sample spaces. *Journal of the American Statistical Association*, 65(331):1248–1255, 1970. ISSN 01621459.
- K. S. Azoury. Bayes solution to dynamic inventory models under unknown demand distribution. *Management Science*, 31(9):pp. 1150–1160, 1985. ISSN 00251909.
- P. C. Bell. Adaptive sales forecasting with many stockouts. *The Journal of the Operational Research Society*, 32(10):pp. 865–873, 1981. ISSN 01605682.
- A. Bensoussan, M. Çakanyildirim, A. Royal, and S. P. Sethi. Bayesian and adaptive controls for a newsvendor facing exponential demand. *Risk and Decision Analysis*, 1(4):197–210, 2009.
- E. Berk, U. Gurler, and R. A. Levine. Bayesian demand updating in the lost sales newsvendor problem: A two-moment approximation. *European Journal of Operational Research*, 182(1):256–281, 2007.
- D. Bertsimas and A. Thiele. A robust optimization approach to inventory theory. *Operations Research*, 54(1):150–168, January 2006. ISSN 1526-5463.
- D. Bienstock and N. Özbay. Computing robust basestock levels. *Discrete Optimization*, 5(2):389–414, May 2008. ISSN 15725286.
- J. H. Bookbinder and A. E. Lordahl. Estimation of inventory Re-Order levels using the bootstrap statistical procedure. *IIE Transactions*, 21(4):302–312, December 1989.
- J. W. Bradford and P. K. Sugrue. A bayesian approach to the two-period style-goods inventory problem with single replenishment and heterogeneous poisson demands. *The Journal of the Operational Research Society*, 41(3):pp. 211–218, 1990. ISSN 01605682.

- D. Chafaï and D. Concordet. Confidence regions for the multinomial parameter with small sample size. *Journal of the American Statistical Association*, 104 (487):1071–1079, September 2009.
- L. Chen. Bounds and heuristics for optimal bayesian inventory control with unobserved lost sales. *Operations Research*, 58(2):396–413, March 2010. ISSN 1526-5463.
- C. J. Clopper and E. S. Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934. ISSN 00063444.
- S. A. Conrad. Sales data and the estimation of demand. *Operational Research Quarterly (1970-1977)*, 27(1):pp. 123–127, 1976. ISSN 00303623.
- X. Ding and M. L. Puterman. The censored newsvendor and the optimal acquisition of information. *Operations Research*, 50:517–527, 1998.
- G. D. Eppen and A. V. Iyer. Improved fashion buying with bayesian updates. *Operations Research*, 45(6):pp. 805–819, 1997. ISSN 0030364X.
- C. Forbes, M. Evans, N. Hastings, and B. Peacock. *Statistical Distributions*. Wiley, New York, 2000.
- R. D. Fricker and C. A. Goodhart. Applying a bootstrap approach for setting reorder points in military supply systems. *Naval Research Logistics*, 47(6): 459–478, September 2000.
- Y. Fukuda. Estimation problems in inventory control. Technical Report 64, University of California, Los Angeles, CA, 1970.
- G. Gallego and I. Moon. The distribution free newsboy problem: Review and extensions. *The Journal of the Operational Research Society*, 44(8):pp. 825–834, 1993. ISSN 01605682.
- G. Gallego, J. K. Ryan, and D. Simchi-Levi. Minimax analysis for finite-horizon inventory models. *IIE Transactions*, 33(10):861–874, 2001.
- F. Garwood. Fiducial limits for the poisson distribution. *Biometrika*, 28:437–442, 1936.
- S. Gupta. *A theory of adjusting parameter estimates in decision models*. PhD thesis, Case Institute of Technology, Cleveland, 1960.
- P. Harremoës. Binomial and poisson distributions as maximum entropy distributions. *IEEE Transactions on Information Theory*, 47(5):2039–2041, July 2001. ISSN 00189448.
- R. H. Hayes. Statistical estimation problems in inventory control. *Management Science*, 15(11):pp. 686–701, 1969. ISSN 00251909.

- R. H. Hayes. Efficiency of simple order statistic estimates when losses are piecewise linear. *Journal of the American Statistical Association*, 66(333):127–135, March 1971.
- R. M. Hill. Parameter estimation and performance measurement in lost sales inventory systems. *International Journal of Production Economics*, 28(2): 211–215, November 1992.
- R. M. Hill. Applying bayesian methodology with a uniform prior to the single period inventory model. *European Journal of Operational Research*, 98(3): 555–562, May 1997.
- R. M. Hill. Bayesian decision-making in inventory modeling. *IMA Journal of Mathematics Applied in Business and Industry*, 10:165–176, 1999.
- W. T. Huh, G. Janakiraman, J. A. Muckstadt, and P. Rusmevichientong. An adaptive algorithm for finding the optimal Base-Stock policy in lost sales inventory systems with censored demand. *Mathematics of Operations Research*, 34(2):397–416, May 2009. ISSN 1526-5471.
- D. L. Iglehart. The dynamic inventory problem with unknown demand distribution. *Management Science*, 10(3):pp. 429–440, 1964. ISSN 00251909.
- E. T. Jaynes. Information Theory and Statistical Mechanics. *Physical Review*, 106:620–630, May 1957.
- H. Jeffreys. *Theory of Probability*. Clarendon Press, Oxford, UK, 1961.
- N. L. Johnson. Systems of frequency curves generated by methods of translation. *Biometrika*, 36(1/2):pp. 149–176, 1949. ISSN 00063444.
- I. S. Kevork. Estimating the optimal order quantity and the maximum expected profit for single-period inventory decisions. *Omega*, 38:218–227, 2010.
- M. Khouja. The single-period (newsvendor) problem: Literature review and suggestions for future research. *Omega*, 27:537–553, 2000.
- D. Klabjan, D. Simchi-Levi, and M. Song. Robust stochastic Lot-Sizing by means of histograms. *Production and Operations Management*, 22(3):691–710, May 2013. ISSN 10591478.
- M. A. Lariviere and E. L. Porteus. Stalking information: Bayesian inventory management with unobserved lost sales. *Management Science*, 45(3):346–363, March 1999. ISSN 1526-5501.
- H.-S. Lau and A. H.-L. Lau. Estimating the demand distributions of single-period items having frequent stockouts. *European Journal of Operational Research*, 92(2):254–265, July 1996. ISSN 03772217.
- L. Le Cam. Maximum likelihood: An introduction. *International Statistical Review / Revue Internationale de Statistique*, 58(2), 1990. ISSN 03067734.

- A. J. Lee, S. O. Nyangoma, and G. A. F. Seber. Confidence regions for multinomial parameters. *Computational Statistics & Data Analysis*, 39(3):329–342, May 2002. ISSN 01679473.
- C.-M. Lee. A bayesian approach to determine the value of information in the newsboy problem. *International Journal of Production Economics*, 112(1):391–402, March 2008. ISSN 09255273.
- R. Levi, R. O. Roundy, and D. B. Shmoys. Provably Near-Optimal Sampling-Based policies for stochastic inventory control models. *Mathematics of Operations Research*, 32(4):821–839, November 2007. ISSN 1526-5471.
- L. H. Liyanage and J. G. Shanthikumar. A practical inventory control policy using operational statistics. *Operations Research Letters*, 33(4):341–348, July 2005. ISSN 01676377.
- A. E. Lordahl and J. H. Bookbinder. Order-statistic calculation, costs, and service in an (s, q) inventory system. *Naval Research Logistics*, 41(1):81–97, February 1994.
- W. S. Lovejoy. Myopic policies for some inventory models with uncertain demand distributions. *Management Science*, 36(6):pp. 724–738, 1990. ISSN 00251909.
- X. Lu, J.-S. Song, and K. Zhu. Inventory control with unobservable lost sales and bayesian updates. Technical report, Duke University’s Fuqua School of Business, 2008.
- A. J. Mersereau. Demand estimation from censored observations with inventory record inaccuracy. Technical report, Kenan-Flagler Business School, University of North Carolina, 2012.
- I. Moon and S. Choi. The distribution free newsboy problem with balking. *The Journal of the Operational Research Society*, 46(4):pp. 537–542, 1995. ISSN 01605682.
- S. Nahmias. Demand estimation in lost sales inventory systems. *Naval Research Logistics*, 41(6):739–757, 1994.
- W. K. Newey and D. McFadden. Large sample estimation and hypothesis testing. In R. F. Engle and D. McFadden, editors, *Handbook of Econometrics*, volume 4 of *Handbook of Econometrics*, chapter 36, pages 2111–2245. Elsevier, 03 1986.
- J. Neyman. Outline of a theory of statistical estimation based on the classical theory of probability. *Philosophical Transactions of the Royal Society of London*, 236:333–380, 1937.
- J. Neyman. Fiducial argument and the theory of confidence intervals. *Biometrika*, 32:128–150, 1941.

- E. Notzon. Minimax inventory and queueing models. Technical Report 7, Department of Operations Research, Stanford University, Stanford, CA, 1970.
- G. Perakis and G. Roels. Regret in the newsvendor model with partial information. *Operations Research*, 56(1):188–203, January 2008. ISSN 1526-5463.
- H. E. Scarf. Bayes solutions of the statistical inventory problem. *The Annals of Mathematical Statistics*, 30(2):pp. 490–508, 1959. ISSN 00034851.
- H. E. Scarf. Some remarks on bayes solutions to the inventory problem. *Naval Research Logistics Quarterly*, 7(4):591–596, 1960. ISSN 1931-9193.
- H. E. Scarf, K. J. Arrow, and S. Karlin. A min-max solution of an inventory problem. *Studies in the mathematical theory of inventory and production*, 10: 201–209, 1958.
- C.-T. See and M. Sim. Robust approximation to multiperiod inventory management. *Operations Research*, 58(3):583–594, May 2010. ISSN 1526-5463.
- E. A. Silver, D. F. Pyke, and R. Peterson. *Inventory Management and Production Planning and Scheduling*. John-Wiley and Sons, New York, 1998.
- M. B. Swift. Comparison of confidence intervals for a poisson mean - further considerations. *Communications in Statistics - Theory and Methods*, 38:748–759, 2009.
- K. S. Trivedi. *Probability and Statistics with Reliability, Queueing, and Computer Science Applications, 2nd Edition*. Wiley-Interscience, 1 edition, October 2001. ISBN 0471333417.
- S. van Dongen. Prior specification in bayesian statistics: Three cautionary tales. *Journal of Theoretical Biology*, 242(1):90–100, September 2006. ISSN 00225193.
- W. E. Wecker. Predicting demand from sales data in the presence of stockouts. *Management Science*, 24(10):pp. 1043–1054, 1978. ISSN 00251909.
- J. Yue, B. Chen, and M.-C. Wang. Expected value of distribution information for the newsvendor problem. *Operations Research*, 54(6):1128–1136, November 2006. ISSN 1526-5463.
- W. Zieliński. The shortest clopper-pearson confidence interval for binomial probability. *Communications in Statistics - Simulation and Computation*, 39 (1):188–193, 2010.