



Production, Manufacturing and Logistics

## Confidence-based optimisation for the newsvendor problem under binomial, Poisson and exponential demand

Roberto Rossi <sup>a,\*</sup>, Steven Prestwich <sup>b,2</sup>, S. Armagan Tarim <sup>c,3</sup>, Brahim Hnich <sup>d</sup><sup>a</sup> Business School, University of Edinburgh, United Kingdom<sup>b</sup> Insight Centre for Data Analytics, University College Cork, Ireland<sup>c</sup> Institute of Population Studies, Hacettepe University, Turkey<sup>d</sup> Department of Computer Engineering, Izmir University of Economics, Turkey

## ARTICLE INFO

## Article history:

Received 5 October 2013

Accepted 9 June 2014

Available online 23 June 2014

## Keywords:

Inventory control

Newsvendor problem

Confidence interval analysis

Demand estimation

Sampling

## 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 produces an upper and a lower bound for the associated cost. We apply this approach to three demand distributions 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.

© 2014 Elsevier B.V. All rights reserved.

### 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, Pyke, & Peterson, 1998). Most of the literature on the newsvendor 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 ideal 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 problem under limited historical demand data are the maximum likelihood (see e.g. Scarf, 1959; Gupta, 1960; Fukuda, 1970) 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 may have little relevance in practice. Unfortunately, both the maximum

\* Corresponding author. Address: Business School, University of Edinburgh, 29 Buccleuch Place, EH8 9JS, Edinburgh, United Kingdom. Tel.: +44 (0)131 6515239; fax: +44 (0)131 6508077.

E-mail addresses: [roberto.rossi@ed.ac.uk](mailto:roberto.rossi@ed.ac.uk) (R. Rossi), [s.prestwich@4c.ucc.ie](mailto:s.prestwich@4c.ucc.ie) (S. Prestwich), [armagan.tarim@hacettepe.edu.tr](mailto:armagan.tarim@hacettepe.edu.tr) (S.A. Tarim), [brahim.hnich@ieu.edu.tr](mailto:brahim.hnich@ieu.edu.tr) (B. Hnich).

<sup>1</sup> R. Rossi is supported by the University of Edinburgh CHSS Challenge Investment Fund.

<sup>2</sup> This publication has emanated from research supported in part by a research Grant from Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289.

<sup>3</sup> S. Armagan Tarim is supported by the Scientific and Technological Research Council of Turkey (TUBITAK) under Grant No. MAG-110M500.

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, Biller, and Tayur (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 demand 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.

Our contributions to inventory management are the following.

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 1) and in the rate  $\lambda$  of a Poisson demand (Theorem 2); we also establish that the newsvendor cost function is quasi-convex in the expected value  $1/\lambda$  of an exponential distribution (Theorem 3). These results, to the best of our knowledge, 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 when demand is censored.

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) and 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.

Several authors (see e.g. Scarf, Arrow, & Karlin, 1958; Gallego & Moon, 1993; Perakis & Roels, 2008) investigated the so-called *distribution free* newsvendor, in which 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.

However, 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 (see e.g. Hayes, 1971; Levi, Roundy, & Shmoys, 2007; Huh, Janakiraman, Muckstadt, & Rusmevichientong, 2009). 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 (see Perakis & Roels, 2008, p. 190).

When the class of the distribution is known but its parameters must be estimated from a set of samples non-parametric approaches are not appropriate, since they would produce conservative results. For this reason, in this work we focus on parametric approaches. According to Berk, Gurler, and Levine (2007) there are two general parametric 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 observed demand data. 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. A further 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 censored demand (see e.g. Hill, 1992; Lau & Lau, 1996).

In the literature we find several works on Bayesian approaches under fully observed demand (see e.g. Hill, 1997; Lee, 2008; Bensoussan, Çakanyildirim, Royal, & Sethi, 2009). Similarly, we also find several Bayesian approaches under censoring induced by lost sales (see e.g. Lu, Song, & Zhu, 2008; Chen, 2010; 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. 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.

We shall now turn our attention to frequentist approaches. 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 optimisation problem of finding the optimal 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 demonstrate that this approach is superior to traditional approaches that separate estimation and optimisation. Klabjan, Simchi-Levi, and Song (2013) integrate distribution fitting and robust optimisation by identifying a set of demand distributions that fit historical data according to a given criterion; they then characterise an optimal policy that minimises the maximum expected total cost against such set of demand distributions.

Bayesian and frequentist strategies surveyed so far share a key limitation. Given a limited set of past demand observations, these techniques would analyse these data and provide a single most-promising order quantity and a point estimate of the associated cost. However, they do not answer two fundamental questions: (i) we do not know at what confidence level we can claim the quantity selected to be optimal within a given margin of error and (ii) 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 Keivork (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; thus a first limitation of this work is that demand censoring is not addressed. 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 Akcay et al. (2011), the inventory manager rarely finds 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, Akcay et al. (2011) exploit the “expected one-period cost associated with operating under an estimated inventory policy” (ETOC), originally introduced in Hayes (1969). Hayes (1969) discussed applications of ETOC to exponentially and normally distributed demands and identified the optimal biased order quantity that minimises ETOC in presence of limited historical demand data. This was one of the first works blending statistical estimation with inventory optimisation. Akcay 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. Despite their ability to quantify the inaccuracy in the inventory-target estimation as a function of the length of the historical data via the ETOC, the approaches in Hayes (1969) and Akcay et al. (2011) do 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 they are able to produce an exact confidence interval to estimate the expected total cost associated with a given ordering decision. Finally, both Hayes (1969) and Akcay et al. (2011) do not address censored demand.

### 3. The newsvendor problem

In this section, we shall summarise 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  $\mu$  and variance  $\sigma^2$ . Let  $o$  be the unit overage cost, paid for each item left in stock after demand is realised, and let  $u$  be the unit underage cost, paid for each unit of unmet demand; we assume that  $u > o > 0$ . Let  $g(x) = ox^+ + ux^-$ , 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. Let  $F(Q) = \Pr\{d \leq Q\}$  denote the cumulative distribution function of  $d$ . If  $d$  is continuous and  $F$  strictly increasing, there is a unique optimal solution

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

where  $\beta = u/(u + o)$  is usually named “critical fractile.”

In practice, the probability distribution of the random demand  $d$  often has finite support over the set  $\mathbb{N}_0 = \{0, 1, 2, \dots\}$ . We can then work with the forward difference  $\Delta G(Q) = G(Q + 1) - G(Q)$ ,  $Q \in \mathbb{N}_0$ . Since  $\Delta G(Q) = o - (o + u)\Pr\{d > j\}$  is non-decreasing in  $Q$  and  $\lim_{Q \rightarrow \infty} \Delta G(Q) = o > 0$ , an optimal solution is given by  $Q^* = \min\{Q \in \mathbb{N}_0 : \Delta G(Q) \geq 0\}$  or equivalently

$$Q^* = \min\{Q \in \mathbb{N}_0 : F(Q) \geq \beta\}. \quad (2)$$

#### 3.1. A frequentist and a Bayesian approach

Consider the situation in which the decision maker knows the class of the random demand distribution (e.g. binomial), but does

not know the actual values of some or all the (stationary) parameters of such a distribution. Nonetheless, she is given a set of  $M$  past realisations of the demand from which she has to infer the optimal order quantity and, possibly, estimate the cost associated with the quantity selected. We now 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 use of these two approaches to illustrate how our approach can be used to complement existing frequentist or Bayesian strategies. For the sake of brevity, 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 & 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  $p$ . As discussed in Hill (1997), in the context of the newsvendor problem we might have a known and fixed number  $N$  of customers each having a probability  $p$  of purchasing a single unit within each period  $i \in \{1, \dots, M\}$ . Then, the maximum likelihood estimator for parameter  $p$  is  $\hat{p} = \sum_{i=1}^M d_i / (MN)$ . After computing  $\hat{p}$ , the decision maker employs the random variable  $\text{bin}(N; \hat{p})$  in place of the actual unknown demand distribution in Eq. (2) 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). Consider Bayes' theorem  $\mathcal{G}(a|b) = \frac{\mathcal{F}(b|a) \cdot \mathcal{P}(a)}{\int \mathcal{F}(b|u) \cdot \mathcal{P}(u) du}$ , with  $a$  continuous and  $b$  discrete or continuous. 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. The prior distribution of  $a$ ,  $\mathcal{P}(a)$ , describes an estimate of the likely value that  $a$  might take, without considering the observed samples. This estimation may be derived from expert assessment and/or collateral data. The likelihood function  $\mathcal{L}(b|a)$  represents the probability of observing a set of samples  $b$  given  $a$ . The posterior distribution of  $a$ ,  $\mathcal{G}(a|b)$ , is an updated estimate of the values  $a$  is likely to take based on the prior distribution and the observed information. To express an initial state of complete ignorance with respect to  $a$ , Hill adopts an uninformative prior. Then, by employing the conjugate prior for the particular distribution under analysis, he constructs the posterior distribution for the newsvendor demand as follows  $\mathcal{F}(x|b) = \int_a \mathcal{L}(x|a) \mathcal{G}(a|b) da$ . Finally, Hill uses this posterior distribution in place of the unknown true distribution for the demand in Eq. (1) or (2) to obtain an order quantity  $\hat{Q}^*$ .

## 4. Binomial demand

Consider, as discussed in Hill (1997), a newsvendor that has a known and fixed number  $N$  of customers each having a probability

$p$  of purchasing a single unit within each period  $i \in \{1, \dots, M\}$ . It is a well-known fact that any experiment comprising a sequence of  $N$  Bernoulli trials, each having the same “yes” (respectively, “no”) probability  $p$  (respectively,  $1 - p$ ), can be represented by a random variable  $\text{bin}(N; p)$  that follows a binomial distribution with probability mass function

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

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

Let us now consider the situation in which the true value  $\bar{p}$  of parameter  $p \in [0, 1]$  is not known. The decision maker is given a set of  $M$  past realisations of the demand and from these realisations she has to infer the optimal order quantity and, possibly, estimate the associated cost.

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 & Pearson, 1934; Garwood, 1936; Epstein & Sobel, 1954, etc.) exist for building exact confidence intervals associated with particular distributions. Approximate techniques have been also investigated (see e.g. Agresti & Coull, 1998) because, especially with small sample sizes, an exact confidence interval may be unnecessarily conservative.

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  $[p_{lb}, p_{ub}]$ , where

$$p_{lb} = \min\{p | \Pr\{\text{bin}(N; p) \geq X\} \geq (1 - \alpha)/2\},$$

$$p_{ub} = \max\{p | \Pr\{\text{bin}(N; p) \leq X\} \geq (1 - \alpha)/2\},$$

$X$  is the number of successes (or “yes” events) observed in the sample and  $\alpha$  is the confidence probability. Note that we assume  $p_{lb} = 0$  when  $X = 0$  and  $p_{ub} = N$  when  $X = N$ . As discussed by Forbes, Evans, Hastings, and Peacock (2000), 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, discussed in Agresti and Coull (1998), uses quantiles from the  $F$  distribution.

Intuitively, the “quality” of a confidence interval is directly related to its size. When the confidence level is the same, 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 interval (see e.g. Zieliński, 2010) or “good” approximate ones (see e.g. Agresti & Coull, 1998). It should be noted that the discussion that follows is independent of the particular interval adopted.



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  $p$  of the binomial demand  $\text{bin}(N; p)$  is simply  $[p_{lb}, p_{ub}]$  where

$$p_{lb} = \min\{p \mid \Pr\{\text{bin}(MN; p) \geq X\} \geq (1 - \alpha)/2\},$$

$$p_{ub} = \max\{p \mid \Pr\{\text{bin}(MN; p) \leq X\} \geq (1 - \alpha)/2\},$$

and  $X = \sum_{i=1}^M d_i$ . Let  $Q_{lb}^*$  be the optimal order quantity for the newsvendor problem under a  $\text{bin}(N, p_{lb})$  demand and  $Q_{ub}^*$  be the optimal order quantity for the newsvendor problem under a  $\text{bin}(N, p_{ub})$  demand. As discussed,  $\Delta G(Q)$  is non-decreasing in  $Q$ . Furthermore, it is clear that the optimal order quantity  $Q^*$  must increase if the value of  $p$  increases. It immediately follows that, according to past demand realisations, with confidence probability greater or equal to  $\alpha$  the optimal order quantity  $Q^*$  is a member of the set  $\{Q_{lb}^*, \dots, Q_{ub}^*\}$  (Fig. 1)—note that this probability may be greater than  $\alpha$  because of the discrete nature of the demand.

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  $Q$ ,

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

Then, consider the function

$$G_Q(p) = \underbrace{o \sum_{i=0}^Q \Pr\{\text{bin}(N; p) = i\} (Q - i)}_{C_Q^o(p)} + \underbrace{u \sum_{i=Q}^N \Pr\{\text{bin}(N; p) = i\} (i - Q)}_{C_Q^u(p)}, \quad (3)$$

in which the order quantity  $Q$  is fixed and in which we vary the “success” probability  $p \in [0, 1]$ . It can be proved that  $G_Q(p)$  is convex in the continuous parameter  $p$ . Firstly, we rewrite Eq. (3) as

$$G_Q(p) = o(Q - Np) + (o + u) \sum_{i=Q}^N (1 - \Pr\{\text{bin}(N; p) \leq i\}). \quad (4)$$

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




Fig. 1. Determining the candidate set  $\{Q_{lb}^*, \dots, Q_{ub}^*\}$ ; note that the family of functions obtained by increasing  $Q$  is discrete, and in some cases it may not be possible to find  $Q_{lb}$  (resp.  $Q_{ub}$ ) such that  $\Pr\{\text{bin}(N; p) \leq Q_{lb}\}$  (resp.  $\Pr\{\text{bin}(N; p) \leq Q_{ub}\}$ ) is exactly  $\beta$ , e.g.  $Q_{lb}^*$  and  $Q^*$  in the picture.

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

**Theorem 1.** For  $Q \leq N$ ,

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

is a positive function of  $p \in [0, 1]$ .

**Proof.** The proof is given in the supplementary material, Appendix A. □

Although it is possible to prove that  $G_Q(p)$  is convex in  $p$ , there is no closed form expression for finding the  $p^*$  that minimises this function. We shall therefore use numerical convex optimisation approaches to find the  $p^*$  that minimises or maximises this function over a given interval.

Let us consider the confidence interval  $[p_{lb}, p_{ub}]$  for the parameter  $p$  of the binomial demand. For a given order quantity  $Q$ , consider the values

$$p_{Q,\min}^* = \arg \min_{p \in [p_{lb}, p_{ub}]} G_Q(p) \quad p_{Q,\max}^* = \arg \max_{p \in [p_{lb}, p_{ub}]} G_Q(p).$$

With confidence probability  $\alpha$ ,  $c_{lb} = G_Q(p_{Q,\min}^*)$  and  $c_{ub} = G_Q(p_{Q,\max}^*)$  represent a lower and an upper bound, respectively, for the cost associated with  $Q$  (Fig. 2-A). Note that with probability  $1 - \alpha$  the interval  $[p_{lb}, p_{ub}]$  will not cover  $\bar{p}$  (Fig. 2B and C). In the case of (Fig. 2C)  $[p_{lb}, p_{ub}]$  does not cover  $\bar{p}$ , but the true value  $G_Q(\bar{p})$  of the expected total cost associated with  $Q$  is covered by  $[p_{lb}, p_{ub}]$ . Nevertheless, this should not be regarded as a coverage event, because as shown in Fig. 2D the two cost components of  $G_Q(\bar{p})$ —overage  $C_Q^o(\bar{p})$  and underage  $C_Q^u(\bar{p})$  cost—are misclassified, i.e. they are not covered by the respective intervals  $[c_{lb}^o, c_{ub}^o]$  and  $[c_{lb}^u, c_{ub}^u]$ .

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

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

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

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  $p$  of the  $\text{bin}(N; p)$  demand is computed as

$$p_{lb} = \min \left\{ p \mid \Pr \left\{ \text{bin} \left( \sum_{j=1}^M \hat{N}_j; p \right) \geq X \right\} \geq (1 - \alpha)/2 \right\},$$

$$p_{ub} = \max \left\{ p \mid \Pr \left\{ \text{bin} \left( \sum_{j=1}^M \hat{N}_j; p \right) \leq X \right\} \geq (1 - \alpha)/2 \right\},$$

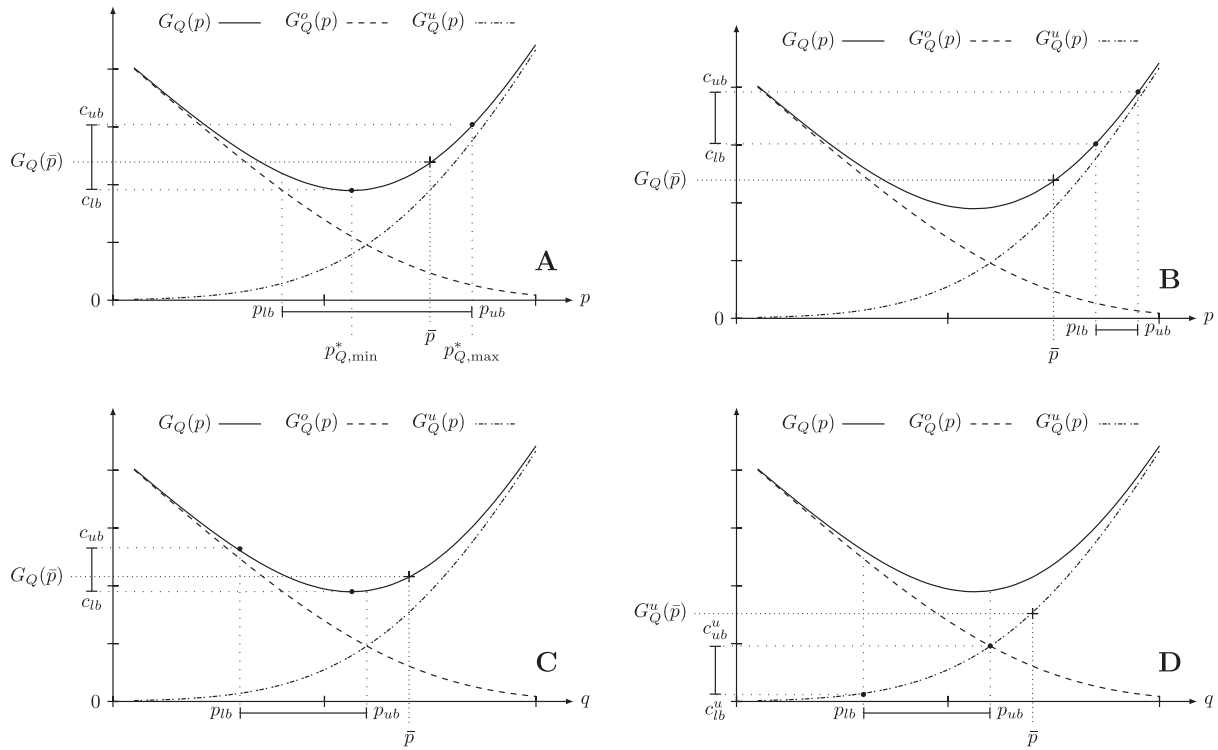


Fig. 2. Coverage of the true cost associated with an order quantity  $Q$ .

where  $\hat{N}_j$  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.

To keep the discussion compact, the computational procedure for binomial demand is presented in the supplementary material, Appendix B.

In general, the set  $\{Q_{lb}^*, \dots, Q_{ub}^*\}$  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 a promising ordering quantity in this set. However, it is worth remarking that information we provide on the minimum and maximum cost associated with each order quantity in  $\mathcal{Q}$  allows not only a better estimation of the actual cost associated with a given decision, but it can also be exploited to make more educated choices than other existing approaches. For instance, if a manager is risk averse, she may select the order quantity for which the  $\alpha$  confidence interval for the estimated cost has the lowest possible upper bound.

### 5. Poisson demand

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

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

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

We shall now consider, also in this case, the situation in which the parameter  $\lambda$  is not known. Instead, the decision maker is given a set of  $M$  past realisations of the demand. As in the previous case,

from these realisations she 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  $\lambda$  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  $\lambda$ . We rewrite  $\bar{d} = \sum_{i=0}^M d_i$ . According to Garwood (1936), the confidence interval for  $\lambda$  is  $(\lambda_{lb}, \lambda_{ub})$ , where

$$\lambda_{lb} = \min\{\lambda \mid \Pr\{\text{Poisson}(M\lambda) \geq \bar{d}\} \geq (1 - \alpha)/2\},$$

$$\lambda_{ub} = \max\{\lambda \mid \Pr\{\text{Poisson}(M\lambda) \leq \bar{d}\} \geq (1 - \alpha)/2\}.$$

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

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

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  $1/M$ . Brown, Cai, and Dasgupta (2003) list 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.

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  $\lambda$  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  $\lambda$  of the Poisson demand. Let  $Q_{lb}^*$  be the optimal order quantity for the newsvendor problem under a  $Poisson(\lambda_{lb})$  demand and  $Q_{ub}^*$  be the optimal order quantity for the newsvendor problem under a  $Poisson(\lambda_{ub})$  demand. With confidence probability  $\alpha$  the optimal order quantity  $Q^*$  is a member of the set  $\{Q_{lb}^*, \dots, Q_{ub}^*\}$ .

Consider the cost associated with an order quantity  $Q$ ,

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

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

$$G_Q(\lambda) = o \sum_{i=0}^Q \Pr\{Poisson(\lambda) = i\}(Q - i) + u \sum_{i=Q}^{\infty} \Pr\{Poisson(\lambda) = i\}(i - Q), \tag{5}$$

is convex in the continuous parameter  $\lambda$ . Firstly, we rewrite Eq. (5) as

$$G_Q(\lambda) = o(Q - \lambda) + (o + u) \sum_{i=Q}^{\infty} (1 - \Pr\{Poisson(\lambda) \leq i\}). \tag{6}$$

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

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

Therefore, we have to prove that

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

**Theorem 2.** For  $Q \geq 0$ ,

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

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

**Proof.** The proof is given in the supplementary material, Appendix C. □

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 optimisation approaches to find the  $\lambda^*$  that minimises or maximises 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  $\lambda$  of the  $Poisson(\lambda)$  demand is computed as

$$\lambda_{lb} = \min\{\lambda | \Pr\{Poisson(\widehat{M}\lambda) \geq \bar{d}\} \geq (1 - \alpha)/2\},$$

$$\lambda_{ub} = \max\{\lambda | \Pr\{Poisson(\widehat{M}\lambda) \leq \bar{d}\} \geq (1 - \alpha)/2\}.$$

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.

To keep the discussion compact, the computational procedure for Poisson demand is presented in the supplementary material, Appendix D.

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

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

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  $\lambda$  can represent the inter-arrival time between two unit demand occurrences in a Poisson process with rate parameter  $\lambda$ . 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  $\lambda$ . Such a situation can be handled by following the discussion in the previous section. In the second case, by using Eq. (1), we easily obtain the optimal order quantity  $Q^*$  for  $\exp(\lambda)$ . This is simply

$$Q^* = -\frac{1}{\lambda} \ln\left(\frac{h}{h+p}\right). \tag{7}$$

Furthermore, consider the cost function

$$G(Q) = o \int_0^Q (Q - i) \Pr\{\exp(\lambda) = i\} di + u \int_Q^{\infty} (i - Q) \Pr\{\exp(\lambda) = i\} di.$$

Rewrite

$$G(Q) = o(Q - \frac{1}{\lambda}) + (o + u) \int_Q^{\infty} (1 - \Pr\{\exp(\lambda) \leq i\}) di.$$

By noting that

$$G(Q) = \frac{o + u}{\lambda} \left( \frac{o}{o + u} (\lambda Q - 1) + e^{-\lambda Q} \right), \tag{8}$$

the optimal cost  $G(Q^*)$  is immediately obtained.

Also in this case we consider the situation in which the parameter  $\lambda$  is not known and the decision maker is given a set of  $M$  past realisations of the demand. As in the previous case, from these realisations she 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  $\lambda$  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  $\lambda$ . 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  $\lambda$  is a random variable  $\text{gamma}(M, 1/\lambda)$  that follows a gamma distribution with shape parameter  $M$  and scale parameter

$1/\lambda$ , the  $\alpha$  confidence interval for the unknown parameter  $\lambda$  is  $(\lambda_{lb}, \lambda_{ub})$ , where

$$\lambda_{lb} = \min\{\lambda | \Pr\{\text{gamma}(M, 1/\lambda) \geq \bar{d}\} \geq (1 - \alpha)/2\},$$

$$\lambda_{ub} = \max\{\lambda | \Pr\{\text{gamma}(M, 1/\lambda) \leq \bar{d}\} \geq (1 - \alpha)/2\}.$$

A closed form expression for this confidence interval—that employs quantiles from the  $\chi^2$  distribution—was proposed by Epstein and Sobel (1954) and takes the following form. Let  $\chi_n^2$  denote the chi-square distribution with  $n$  degrees of freedom, and  $G^{-1}(\chi_n^2, \cdot)$  denote the inverse cumulative distribution function of  $\chi_n^2$ . 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 Trivedi (2001, chap. 10). 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  $1/\bar{d}$ . Finally, a simple approximation can be derived by using a normal approximation to the chi-square distribution.

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  $\lambda$  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  $\lambda$  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  $[Q_{lb}^*, Q_{ub}^*]$ .

It is well-known that the expected total cost  $G(Q)$  associated with an order quantity  $Q$  for a given demand rate  $\lambda > 0$  is convex. Consider the function

$$G_Q(\lambda) = o\left(Q - \frac{1}{\lambda}\right) + (o + u) \int_Q^\infty (1 - \Pr\{\exp(\lambda) \leq i\}) di, \quad (9)$$

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  $\lambda$ . Nevertheless, we shall prove the following property for this function.

**Theorem 3.** Function  $G_Q(\lambda)$  is quasi-convex in the continuous parameter  $\lambda$ .

**Proof.** The proof is given in the supplementary material, Appendix E. □

Because of the property introduced in Theorem 3, we can employ a simple line search procedure in order to find the  $\lambda^*$  that minimises or maximises this function over a given interval.

Since the optimal order quantity  $Q^*$  is, with confidence probability  $\alpha$ , a value in the interval  $[Q_{lb}^*, Q_{ub}^*]$ , it is possible to compute upper ( $c_{ub}^*$ ) and lower ( $c_{lb}^*$ ) bounds for the cost that a manager will face, with confidence probability  $\alpha$ , whatever order quantity she chooses in this interval.

**Theorem 4.** The lower and upper bounds are

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

**Proof.** The proof is given in the supplementary material, Appendix E. □

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.

To keep the discussion compact, the computational procedure for exponential demand is presented in the supplementary material, Appendix F.

7. Computational results

We conducted an extensive computational study to demonstrate how our approach can complement other existing approaches in the literature. Our test set comprises a total of 420 instances under a binomial demand and a total of 200 instances under Poisson and exponential demands. Parameter values investigated in our full factorial analysis are illustrated in Table 1.

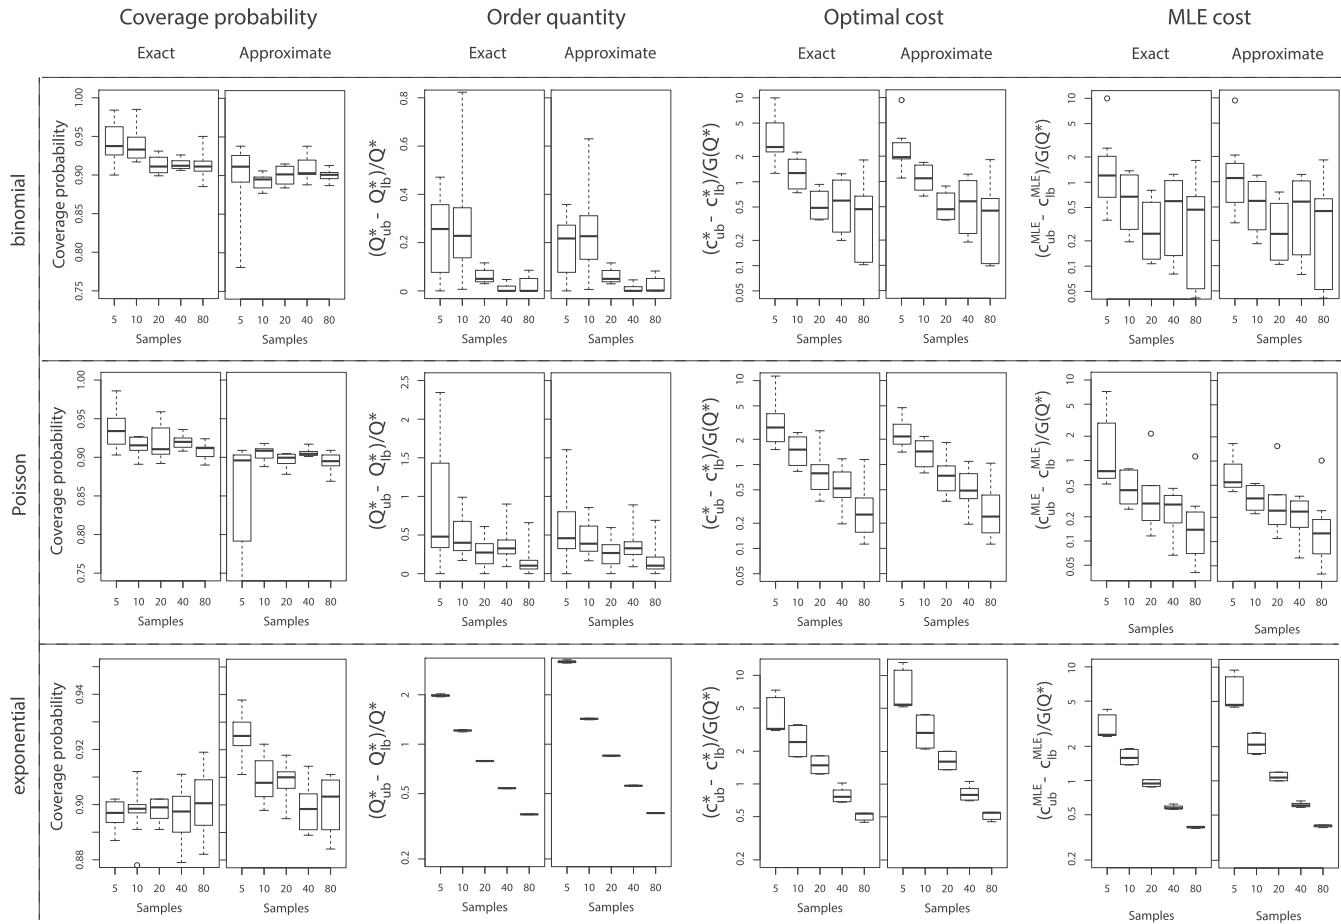
In Fig. G.3, presented in Appendix F of our electronic addendum, we report box plots illustrating the mean absolute percentage error (MAPE) with respect to the true optimal cost for the Bayesian (Bayes) and for the maximum likelihood estimation (MLE) approaches discussed in Section 3.1. It is clear that both these approaches are very good, even under small samples, in determining near-optimal order quantities. Their average MAPE for the real cost associated with the order quantity selected decreases, in fact, from about 6%, if 5 samples are used, to 0.15%, if 80 samples are used. However, the estimation of the cost associated with a given order quantity is problematic. Under small samples, MAPE for the estimated cost reported by these approaches is likely to exceed 100% of the true cost and on average amounts to about 45%. Even when the sample size increases to 80 samples, the average MAPE remains as high as 10%. This is clearly an issue, since by using these approaches the decision maker is not able to reliably estimate the degree of uncertainty associated with the estimated order quantity or cost. Intuitively, a cost or an order quantity estimated by using more samples will be more accurate, the question that these approaches do not answer is how to quantify these different degrees of accuracy.

Our approach answers this question by identifying a range of candidate order quantities that, with prescribed confidence probability, includes the real optimal order quantity for the underlying stochastic demand process. In addition, for each candidate order quantity that is identified, our approach produces an upper and a lower bound for the associated cost. In Fig. 3 we report computational results for our approach over the test set previously discussed. Coverage probabilities for order quantity and cost intervals produced by our approach were estimated by averaging one thousand Monte Carlo (MC) simulations; since this computation

**Table 1**  
Test set employed in our computational analysis.

Parameter	Values	Description
$o$	1	Per unit overage cost
$u$	2, 4, 8, 16	Per unit underage cost
$M$	5, 10, 20, 40, 80	Number of demand samples
$\alpha$	0.9	Confidence level
$N$	1, 2, 4, 8, 16, 32, 64	Binomial trials
$p$	0.5, 0.75, 0.95	Binomial success probability
$\lambda$	0.125, 0.25, 0.5, 1, 2, 4, 8, 16, 32, 64	Poisson/exponential rate





**Fig. 3.** Simulated coverage probabilities; size of the candidate set of order quantities  $[Q_{lb}^*, Q_{ub}^*]$  in relation to the true optimal order quantity  $Q^*$ ; size of the confidence interval  $[c_{lb}^*, c_{ub}^*]$  for the optimal cost in proportion to the true optimal cost  $G(Q^*)$ ; and size of the confidence interval  $[c_{lb}^{MLE}, c_{ub}^{MLE}]$  for the cost associated with the MLE solution in proportion to  $G(Q^*)$ .

turned out to be cumbersome—each MC iteration deals with multiple convex optimisation problems that are solved numerically to determine order quantity and cost interval bounds—we limited our analysis to a randomly selected subset of instances covering approximately 10% of the whole test set; results are summarised again using box plots. We conducted experiments using both exact and approximate confidence interval for the distribution parameter. More specifically, the approximate intervals employed are Agresti–Coull interval (Agresti & Coull, 1998) for the binomial, Wald’s interval (Brown et al., 2003) for the Poisson, and an approximate interval constructed by exploiting a normal approximation to the chi-square distribution for the exponential distribution.

Our results demonstrate that true coverage probabilities for order quantity and cost intervals produced by our approach are in line with the prescribed confidence level  $\alpha$ . When approximate intervals are employed intervals are generally smaller, coverage probability may however fall below expectation if sample size is small. Surprisingly, when demand distribution is exponential, coverage probability for approximate intervals far exceeds the prescribed nominal one for small samples and approximate intervals are larger than the exact ones. This is due to the known fact that a normal approximation to the chi-square distribution performs poorly when sample size is too small. To understand what exactly happens, one should recall that the expected value of an exponential random variable with parameter  $\lambda$  is  $1/\lambda$ ; since we are trying to estimate  $\lambda$ , a smaller confidence interval for  $\lambda$  means a larger interval for the expected demand  $1/\lambda$ , which leads to the conservative results discussed above.

Order quantity and cost intervals reflect manager’s uncertainty about the optimal decision to be taken, its cost, or about the cost associated with a given decision, e.g. ordering the amount suggested by an MLE strategy. Uncertainty about the optimal order quantity, reflected by the size of the interval  $[Q_{lb}^*, Q_{ub}^*]$ , can be quite high when sample size is small; e.g. up to 2.5 times the size of the optimal order quantity  $Q^*$  under a Poisson distribution sampled 5 times. This uncertainty quickly decreases as the same size increases. The same is true for the uncertainty about the cost of an optimal decision, reflected by the size of the interval  $[c_{lb}^*, c_{ub}^*]$ , which can be as large as 10 times the true cost  $G(Q^*)$  of the optimal decision  $Q^*$ . By using our approach, a manager can quantify this uncertainty. For instance, one may determine the cost confidence interval for an optimal decision determined by an MLE strategy, i.e. interval  $[c_{lb}^{MLE}, c_{ub}^{MLE}]$ , which provides an effective tool for addressing the issue illustrated in Fig. G.3 of the high MAPE that might be associated with the cost predicted by MLE or Bayes solutions.

**8. 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.

Likewise, we left as a future research direction the investigation of the case in which parameter  $N$ , rather than parameter  $q$ , of a binomial distribution is unknown—confidence regions that can be employed in this case are discussed in Tang and Sindler (1987)—as well as the case in which both parameters of a binomial distribution are unknown.

Finally, in principle it may be possible to extend the analysis to other distributions such as the multinomial, for which confidence intervals are surveyed in Lee, Nyangoma, and Seber (2002) and 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.

## 9. 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 optimisation problems. Our strategy analytically combines confidence interval analysis and inventory optimisation. More specifically, 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. Finally, we presented a comprehensive numerical study in which we showed that existing strategies based on maximum likelihood estimators or on Bayesian analysis may provide—particularly under small samples—unreliable point estimates of the cost of a given ordering decision; and that our novel approach can be employed to better assess the expected total cost associated with such a decision, as well as the potential discrepancy between the true optimal decision and the one selected.

## Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.ejor.2014.06.007>.

## References

- Agrawal, N., & Smith, S. A. (1996). Estimating negative binomial demand for retail inventory management with unobservable lost sales. *Naval Research Logistics*, 43, 839–861.
- Agresti, A., & Coull, B. A. (1998). Approximate is better than “exact” for interval estimation of binomial proportions. *The American Statistician*, 52, 119–126.
- Akçay, A., Biller, B., & Tayur, S. (2011). Improved inventory targets in the presence of limited historical demand data. *Manufacturing & Service Operations Management*, 13, 297–309.
- Andersen, E. B. (1970). Sufficiency and exponential families for discrete sample spaces. *Journal of the American Statistical Association*, 65, 1248–1255.
- Bensoussan, A., Çakanyildirim, M., Royal, A., & Sethi, S. P. (2009). Bayesian and adaptive controls for a newsvendor facing exponential demand. *Risk and Decision Analysis*, 1, 197–210.
- Berk, E., Gurler, U., & Levine, R. A. (2007). Bayesian demand updating in the lost sales newsvendor problem: A two-moment approximation. *European Journal of Operational Research*, 182, 256–281.
- Brown, L. D., Cai, T. T., & Dasgupta, A. (2003). Interval estimation in exponential families. *Statistica Sinica*, 13, 19–49.
- Chafai, D., & Concordet, D. (2009). Confidence regions for the multinomial parameter with small sample size. *Journal of the American Statistical Association*, 104, 1071–1079.
- Chen, L. (2010). Bounds and heuristics for optimal bayesian inventory control with unobserved lost sales. *Operations Research*, 58, 396–413.
- Clopper, C. J., & Pearson, E. S. (1934). The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26, 404–413.
- Epstein, B., & Sobel, M. (1954). Some theorems relevant to life testing from an exponential distribution. *The Annals of Mathematical Statistics*, 25, 373–381.
- Forbes, C., Evans, M., Hastings, N., & Peacock, B. (2000). *Statistical Distributions*. New York: Wiley.
- Fukuda, Y. (1970). *Estimation problems in inventory control*. Technical Report 64 University of California Los Angeles, CA.
- Gallego, G., & Moon, I. (1993). The distribution free newsboy problem: Review and extensions. *The Journal of the Operational Research Society*, 44, 825–834.
- Garwood, F. (1936). Fiducial limits for the Poisson distribution. *Biometrika*, 28, 437–442.
- Gupta, S. (1960). *A theory of adjusting parameter estimates in decision models*. Ph.D. thesis Case Institute of Technology, Cleveland.
- Harremoes, P. (2001). Binomial and Poisson distributions as maximum entropy distributions. *IEEE Transactions on Information Theory*, 47, 2039–2041.
- Hayes, R. H. (1969). Statistical estimation problems in inventory control. *Management Science*, 15, 686–701.
- Hayes, R. H. (1971). Efficiency of simple order statistic estimates when losses are piecewise linear. *Journal of the American Statistical Association*, 66, 127–135.
- Hill, R. M. (1992). Parameter estimation and performance measurement in lost sales inventory systems. *International Journal of Production Economics*, 28, 211–215.
- Hill, R. M. (1997). Applying bayesian methodology with a uniform prior to the single period inventory model. *European Journal of Operational Research*, 98, 555–562.
- Hill, R. M. (1999). Bayesian decision-making in inventory modeling. *IMA Journal of Mathematics Applied in Business and Industry*, 10, 165–176.
- Huh, W. T., Janakiraman, G., Muckstadt, J. A., & Rusmevichientong, P. (2009). An adaptive algorithm for finding the optimal base-stock policy in lost sales inventory systems with censored demand. *Mathematics of Operations Research*, 34, 397–416.
- Johnson, N. L. (1949). Systems of frequency curves generated by methods of translation. *Biometrika*, 36, 149–176.
- Kevork, I. S. (2010). Estimating the optimal order quantity and the maximum expected profit for single-period inventory decisions. *Omega*, 38, 218–227.
- Khouja, M. (2000). The single-period (newsvendor) problem: Literature review and suggestions for future research. *Omega*, 27, 537–553.
- Klabjan, D., Simchi-Levi, D., & Song, M. (2013). Robust stochastic lot-sizing by means of histograms. *Production and Operations Management*, 22, 691–710.
- Lau, H.-S., & Lau, A. H.-L. (1996). Estimating the demand distributions of single-period items having frequent stockouts. *European Journal of Operational Research*, 92, 254–265.
- Le Cam, L. (1990). Maximum likelihood: An introduction. *International Statistical Review/Revue Internationale de Statistique*, 58.
- Lee, C.-M. (2008). A bayesian approach to determine the value of information in the newsboy problem. *International Journal of Production Economics*, 112, 391–402.
- Lee, A. J., Nyangoma, S. O., & Seber, G. A. F. (2002). Confidence regions for multinomial parameters. *Computational Statistics & Data Analysis*, 39, 329–342.
- Levi, R., Roundy, R. O., & Shmoys, D. B. (2007). Provably near-optimal sampling-based policies for stochastic inventory control models. *Mathematics of Operations Research*, 32, 821–839.
- Liyanage, L. H., & Shanthikumar, J. G. (2005). A practical inventory control policy using operational statistics. *Operations Research Letters*, 33, 341–348.
- Lu, X., Song, J.-S., & Zhu, K. (2008). *Inventory control with unobservable lost sales and Bayesian updates*. Technical Report Duke University’s Fuqua School of Business.
- Mersereau, A. J. (2012). *Demand estimation from censored observations with inventory record inaccuracy*. Technical Report Kenan-Flagler Business School, University of North Carolina.
- Nahmias, S. (1994). Demand estimation in lost sales inventory systems. *Naval Research Logistics*, 41, 739–757.
- Newey, W. K., & McFadden, D. (1986). Large sample estimation and hypothesis testing. In R. F. Engle & D. McFadden (Eds.), *Handbook of Econometrics* (pp. 2111–2245). Elsevier, chapter 36.
- Neyman, J. (1937). 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.

- Neyman, J. (1941). Fiducial argument and the theory of confidence intervals. *Biometrika*, 32, 128–150.
- Perakis, G., & Roels, G. (2008). Regret in the newsvendor model with partial information. *Operations Research*, 56, 188–203.
- Scarf, H. E. (1959). Bayes solutions of the statistical inventory problem. *The Annals of Mathematical Statistics*, 30, 490–508.
- Scarf, H. E., Arrow, K. J., & Karlin, S. (1958). A min–max solution of an inventory problem. *Studies in the Mathematical Theory of Inventory and Production*, 10, 201–209.
- Silver, E. A., Pyke, D. F., & Peterson, R. (1998). *Inventory management and production planning and scheduling*. New York: John-Wiley and Sons.
- Swift, M. B. (2009). Comparison of confidence intervals for a Poisson mean – Further considerations. *Communications in Statistics – Theory and Methods*, 38, 748–759.
- Tang, R., Sindler, V. K. T. (1987). *Confidence interval for parameter  $n$  in a binomial distribution*. Technical Report CRM 86265 Center for Naval Analyses 4401 Ford Avenue, Post Office Box 16268, Alexandria, Virginia 22302-0268.
- Trivedi, K. S. (2001). *Probability and statistics with reliability, queueing, and computer science applications, 2nd edition* (1st ed.). Wiley-Interscience.
- van Dongen, S. (2006). Prior specification in bayesian statistics: Three cautionary tales. *Journal of Theoretical Biology*, 242, 90–100.
- Zieliński, W. (2010). The shortest Clopper–Pearson confidence interval for binomial probability. *Communications in Statistics – Simulation and Computation*, 39, 188–193.