Production, Manufacturing and Logistics

The distribution-free newsboy problem under the worst-case and best-case scenarios

Jerzy Kamburowski *

College of Business and Innovation, The University of Toledo, Mail Stop #103, Toledo, OH 43606, USA

ABSTRACT

New theoretical foundations for analyzing the newsboy problem under incomplete information about the probability distribution of random demand are presented. Firstly, we reveal that the distribution-free newsboy problem under the worst-case and best-case demand scenarios actually reduces to the standard newsboy problem with demand distributions that bound the allowable distributions in the sense of increasing concave order. Secondly, we provide a theoretical tool for seeking the best-case and worst-case order quantities when merely the support and the first k moments of the demand are known. Using this tool we derive closed form formulas for such quantities in the case of known support, mean and variance, i.e. $k = 2$. Consequently, we generalize all results presented so far in literature for the worst-case and best-case scenarios, and present some new ones. Extensions of our findings to the cases of the known mode of a unimodal demand distribution, the known median, and to other stochastic inventory problems are indicated.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

Since the pioneer works of Arrow, Harris, and Marshak (1951) and Morse and Kimball (1951), the classic single-period, single-item inventory problem with random demand, commonly referred to as the newsboy or newsvendor problem, has attracted a great deal of attention and played a central role at the conceptual foundations of stochastic inventory theory; Porteus (2002, chap. 1). Typically, it is formulated as follows. Each day the proverbial newsboy has to decide how many newspapers to stock before observing demand. He purchases them from a publisher at a unit cost c, and sells them at a price p to customers whose uncertain demand is described by a random variable X. Any unsold items are recycled with a unit salvage value s; it is assumed that $p > c > s$. The newsboy problem is to find the order (purchase) quantity that maximizes the expected profit, and this quantity is known to be the smallest q such that $F(q) \geqslant \frac{p-c}{p-s}$, where F is the cumulative distribution function of X. Numerous extensions of the newsboy problem were reviewed in Khouja (1999) and Qin, Wang, Vakharia, Chen, and Seref (2011). Note here that we adopt a classic approach by assuming that the decision-maker is risk-neutral. The models in which he/she is risk-averse, risk-seeking, or uses a maximum entropy approach can be found in Wang, Webster, and Suresh (2009),

Du Andersson, Jörnsten, Nonås, Sandal, and Ubøe (2012), Wu, Zhu, and Teunter (2013), and Wu et al. (2013).

Scarf (1958) was the first who addressed the distribution-free newsboy problem. He assumed that merely the mean $\mu = E(X)$ and the variance $\sigma^2 = Var(X)$ are known, and derived a closed form formula for the order quantity that maximizes the minimum expected profit over all demand distributions with given μ and σ^2. For this reason, this worst-case order quantity is also referred to as being found under the maximin criterion. Gallego and Moon (1993) disseminated the rather unnoticed result of Scarf, modified it by restricting the demand distribution to non-negative values, provided a simpler proof with economic interpretations, and showed some applications to other stochastic inventory problems. Alfares and Elmorra (2005) extended the modified Scarf formula to the case of shortage penalty.

Gallego, Ryan, and Simchi-Levy (2001) assumed that X is a discrete random variable taking a finite number of given values. When selected moments and percentiles of X are known, they showed that the worst-case order quantity can be found by solving a linear program. They also demonstrated that if the demand distribution is characterized by μ and σ^2 then the maximin policy mostly performs quite well, relative to a policy based on the normally distributed demand. Similar results concerning the performance of the modified Scarf formula were earlier reported by Gallego and Moon (1993). Perakis and Roels (2008) and Andersson et al. (2013) observed, however, that in some situations the maximin policy may lead to an expected profit much lower than the optimal one.

* Tel.: +1 419 530 4361; fax: +1 419 530 2290.
E-mail address: jerzy.kamburowski@utoledo.edu

The order quantity found under the worst-case demand scenario is pessimistic (conservative). A less pessimistic policy based on minimizing the expected maximum regret was proposed as an alternative; see e.g. Yue, Chen, and Wang (2006), Perakis and Roels (2008). Yue et al. (2006) found the corresponding closed form formula for the minimax regret order quantity under the Scarf assumption that μ and σ are known without imposing the non-negativity constraint on the demand distribution. Perakis and Roels (2008) extended their result by adding this obvious constraint, and also solved some cases involving the known median and mode, the symmetricity and unimodality of the demand distribution. Another alternative to the maximin policy was proposed in Andersson et al. (2013). Assuming only the knowledge of μ and σ, the authors demonstrated empirically that finding the most likely distribution in the sense of the maximum entropy leads on average to better results.

To the best of our knowledge, the maximax criterion, which maximizes the expected profit under the best-case demand scenario, has been much less examined. Only trivial cases, resulting from Jensen's inequality, of the best-case order quantities have been mentioned in Gallego and Moon (1993) and Yue et al. (2006). This observation also refers to other distribution-free stochastic inventory problems discussed in literature; see e.g. Godfrey and Powell (2001), Wu, Li, and Tsai (2002), Lin and Chu (2006), Ho (2009), and Kwon and Cheong (2014).

In this paper we present new theoretical foundations for analyzing the distribution-free newsboy problem under the best-case and worst-case scenarios. Firstly, by using a simple expression for the expected profit, we reveal that this problem actually reduces to the standard newsboy problem with demand distributions that bound the allowable distributions in the sense of increasing concave order. Secondly, we provide a theoretical tool for seeking the best-case and worst-case order quantities when merely the support and the first k moments of the demand are known. Using this tool we derive closed form formulas for such quantities in the case of known support $[a, b]$, mean μ, and variance σ^2, i.e. $k = 2$. Consequently, we generalize all results presented so far in literature for the worst-case and best-case scenarios, and present some new ones. Theoretical tools are also provided for examining the problem when the median, or the mode of the unimodal distribution are additionally available.

The paper is organized as follows. In Section 2 we formulate the problem under study and show its main theoretical result. Section 3 includes a theoretical tool for seeking the sharp lower and upper bounds on the expected met demand needed for deriving the worst-case and the best-case order quantities. Using this tool we find, in particular, closed form formulas for such bounds in the case of known support, mean, and variance. The resulting worst-case and best-case order quantities are presented in Section 4. Section 5 illustrates our findings by the use of a numerical example taken from literature. In Section 6 we present some extensions of our results, while final remarks are made in Section 7.

2. Problem formulation

In the classic newsboy problem, no cost is assumed if the order quantity does not meet the demand. Although this cost might be difficult to define in practice, we adopt a more general model considered by Alfares and Elmorra (2005) and Perakis and Roels (2008), in which a known unit lost sales (shortage) cost of ℓ is assumed. Therefore, if q is an order quantity and X denotes the random demand, $\min(X, q)$ represents the demand that is met, $X - \min(X, q)$ the demand that is unmet, and $q - \min(X, q)$ is the salvage amount. Thus, $\min(X, q)$ units will be sold at a price p, $q - \min(X, q)$ units at s, and the ordering and lost sales costs will

be cq and $\ell[X - \min(X, q)]$, respectively. Consequently, the expected profit is expressed by

$$\pi(q) = pE[\min(X, q)] + sE[q - \min(X, q)] - cq - \ell\{\mu - E[\min(X, q)]\}$$
$$= (p + \ell - s) \int_{-\infty}^{\infty} \min(x, q) dF(x) - (c - s)q - \ell\mu,$$

where $\mu = E[X]$ and F is the cumulative distribution function of X, for short called later the distribution, that is, $F(x) = P(X \leqslant x)$.

Since for every distribution F (defined as a right-continuous function) with a finite mean

$$\frac{d_+}{dq}\left[\int_{-\infty}^{\infty} \min(x, q) dF(x)\right] = \frac{d_+}{dq}\left[q - \int_{-\infty}^{q} F(x)dx\right] = 1 - F(q),$$

the first right derivative of $\pi(q)$ is $\frac{d_+}{dq}[\pi(q)] = (p + \ell - c) - (p + \ell - s)F(q)$, and the optimal order quantity q^* can be defined as the smallest q such that $F(q) \geqslant r$, where $r = \frac{p+\ell-c}{p+\ell-s}$. In particular, if X is a continuous random variable, $q^* = F^{-1}(r)$. Note here that the inclusion of the expected lost sales cost $\ell\{\mu - E[\min(X, q)]\}$ can make the maximum expected profit $\pi(q^*)$ negative. Assuming then $q^* = 0$ does not help because $\pi(0) = -\ell\mu$ for a non-negative X.

We would like to emphasize here that the expected profit $\pi(q)$ has been mostly presented by using more complex (but equivalent) expressions. For example, in Alfares and Elmorra (2005) and Perakis and Roels (2008) this profit is shown as

$$\pi(q) = pE[\min(X, q)] + sE[\max(q - X, 0)]] - \ell E[\max(X - q, 0)] - cq.$$

Suppose only a partial information about the distribution F of X is available in the sense that $F \in \mathcal{F}$, where $\mathcal{F}$ is a certain family of distributions with a finite mean. For every q, let $L(q)$ and $U(q)$ be lower and upper bounds on the expected met demand $E[\min(X, q)]$, that is,

$$L(q) \leqslant \min_{F \in \mathcal{F}} \int_{-\infty}^{\infty} \min(x, q) dF(x) \leqslant \max_{F \in \mathcal{F}} \int_{-\infty}^{\infty} \min(x, q) dF(x)$$
$$\leqslant U(q).$$

The bounds $L(q)$ and $U(q)$ are sharp, if there exist distributions $\underline{F}_q, \overline{F}q \in \mathcal{F}$ such that

$$L(q) = \int_{-\infty}^{\infty} \min(x, q) d\underline{F}_q(x) = \min_{F \in \mathcal{F}} \int_{-\infty}^{\infty} \min(x, q) dF(x),$$
$$U(q) = \int_{-\infty}^{\infty} \min(x, q) d\overline{F}_q(x) = \max_{F \in \mathcal{F}} \int_{-\infty}^{\infty} \min(x, q) dF(x).$$

The sharp bounds $L(q)$ and $U(q)$ lead to the following sharp lower and upper bounds on the expected profit $\pi(q)$:

$$\underline{\pi}(q) = (p + \ell - s)L(q) - (c - s)q - \ell\mu,$$
$$\overline{\pi}(q) = (p + \ell - s)U(q) - (c - s)q - \ell\mu.$$

When $\underline{\pi}(q)$ and $\overline{\pi}(q)$ are maximized over q, one defines the worst-case and best-case order quantities denoted by $\underline{q}^*$ and $\overline{q}^*$, respectively. Consequently, for any distribution $F \in \mathcal{F}$, the corresponding optimal order quantity q^* satisfies $\underline{\pi}(\underline{q}^*) \leqslant \pi(q^*) \leqslant \overline{\pi}(\overline{q}^*)$. It does not mean, however, that $\underline{q}^* \leqslant q^* \leqslant \overline{q}^*$, and one can get, for example, $\overline{q}^* < \underline{q}^* < q^*$.

Recall that F is said to be smaller than G in the sense of increasing concave order, written $F \leqslant_{icv} G$, if for every non-decreasing and concave function $\varphi(x)$, $\int_{-\infty}^{\infty} \varphi(x)dF(x) \leqslant \int_{-\infty}^{\infty} \varphi(x)dG(x)$ provided that these integrals are finite; see e.g. Müller and Stoyan (2002, chap. 1) and Shaked and Shanthikumar (2007, chap. 4). Furthermore, if F and G have finite means then $F \leqslant_{icv} G$ if, and only if, $\int_{-\infty}^{\infty} \min(x, q)dF(x) \leqslant \int_{-\infty}^{\infty} \min(x, q)dG(x)$ for all q.

Let $\mathcal{F}$ be a certain family of distributions with the same finite mean. A distribution $\underline{F}(\overline{F})$ is called the infimum (supremum) of $\mathcal{F}$ with respect to $\leqslant_{icv}$ if $\underline{F}(\overline{F})$ is the greatest (smallest) distribution with respect to $\leqslant_{icv}$ such that $\underline{F} \leqslant_{icv} F$ ($F \leqslant_{icv} \overline{F}$) for all $F \in \mathcal{F}$; see e.g. Müller and Stoyan (2002, chap. 1). Although $\underline{F}$ and $\overline{F}$ bound