Omega 38 (2010) 218–227

Contents lists available at ScienceDirect

Omega

journal homepage: www.elsevier.com/locate/omega

Estimating the optimal order quantity and the maximum expected profit for single-period inventory decisions

Ilias S. Kevork *

Department of Economics, University of Thessaly, Korai 43, Volos 38333, Greece

ARTICLE INFO

Article history:
Received 16 August 2008
Accepted 14 September 2009
Available online 23 September 2009

Keywords:
Inventory control
Single-period inventory model
Demand estimation
Asymptotic distribution
Monte-Carlo simulations

ABSTRACT

The paper considers the classical single-period inventory model, also known as the Newsboy Problem, with the demand normally distributed and fully observed in successive inventory cycles. The extent of applicability of such a model to inventory management depends upon demand estimation. Appropriate estimators for the optimal order quantity and the maximum expected profit are developed. The statistical properties of the two estimators are explored for both small and large samples, analytically and through Monte-Carlo simulations. For small samples, both estimators are biased. The form of distribution of the optimal order quantity estimator depends upon the critical fractile, while the distribution of the maximum expected profit estimator is always left-skewed. Small samples properties of the estimators indicate that, when the critical fractile is set over a half, the optimal order quantity is underestimated and the maximum expected profit is overestimated with probability over 50%, whereas the probability of overestimating both quantities exceeds again 50% when the critical fractile is below a half. For large samples, based on the asymptotic properties of the two estimators, confidence intervals are derived for the corresponding true population values. The validity of confidence intervals using small samples is tested by developing appropriate Monte-Carlo simulations. In small samples, these intervals attain acceptable confidence levels, but with high unit shortage cost, for the case of maximum expected profit, significant reductions in their precision and stability are observed.

1. Introduction

A single-period inventory model handles inventory policies for products whose demand lasts one inventory cycle (or one period). In its classical form, which is described in Khouja [14], the inventory decision aims at determining a single order quantity that maximizes the expected profit per inventory cycle. The term "expected" refers to two mutually exclusive events that occur at the end of the inventory cycle: (a) demand is overestimated, and so a stock for the product remains at the end of the period, or (b) demand is underestimated, in which case some customers cannot be satisfied. To derive the expected profit, together with the selling price and the salvage value, the following two cost elements per unit of product are considered: the purchase cost and the shortage cost. The optimal order quantity that maximizes expected profit is determined by equating the probability of the demand not to exceed the order quantity to the critical fractile. The latter is a function of the price, the salvage value and the two cost elements. Schweitzer and Cachon [23] state that when the critical fractile is greater (less) than 0.5, the product is classified as high-profit (low-profit) product.

For single-period inventory models, most of the research has focused on specifying optimal inventory policies under the assumption that parameters of the demand distribution are known. However, the extent of applicability of such models to managerial aspects of inventories depends on the estimation of demand parameters. Berk et al. [5] recognize two general approaches for demand estimation: the Frequentist and the Bayesian. According to the Frequentist approach, point estimates for the unknown parameters of the known parametric demand distribution are obtained using historical data. On the other hand, Bayesian methodology is based on a "prior" distribution for the demand parameter, which is constructed based on collateral data or subjective assessment. From the prior distribution, the corresponding posterior distribution is generated as new data of demand become available. This posterior distribution of the unknown parameter is used to estimate, first, the posterior distribution of demand and, second, the optimal order quantity and the optimum value of the objective function.

Another classification concerns the capability of observing demand for inventory cycles for which stockouts occur. So, we have methods that assume that demand is fully observed and methods that consider the demand over the stocking level as not observed and therefore being considered as lost. In the latter case, appropriate adjustments in the estimation procedure are made to account for the unobserved component of demand. Lau and Lau

* Tel.: +30 24210 74921; fax: +30 24210 74772.
E-mail address: kevork@uth.gr

0305-0483/$ - see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.omega.2009.09.005

[18] classify works in lost sales environments to two groups. Works in the first group estimate the parameters of the demand distribution using sales data, which, for some inventory cycles, are censored by stockouts. Works in the second group offer recursive form(s) to update forecasts for demand parameters (such as mean, variance or mean absolute deviation) and are suitable for non-stationary demand patterns.

In the area of demand estimation, the majority of papers being classified in at least one of the above taxonomies focus their interest on studying how specific problems of demand estimation affect the quality of the estimated parameters, *without addressing any optimization problem*. Indicative examples constitute the works of Nahmias [22] and Agrawal and Smith [2]. Nahmias [22] considered *the order quantity as given* at the stage of developing and evaluating a sequential procedure for estimating and updating the parameters of a normally distributed demand when sales are lost due to stockouts. Modeling demand by the negative binomial distribution, Agrawal and Smith [2] developed a parameter estimation methodology, first when demand is fully observed, and second when demand equals the inventory whenever there are stockouts. The methodology effectiveness was demonstrated by "*artificially*" truncating sales data and comparing the resulting estimates to those obtained using full data.

Research on studying the effects of demand estimation on optimal inventory policies for single-period inventory models is limited. Conrad [7] presented how to estimate from sales data the unknown parameter of Poisson distributed demand and showed that using sales data as a proxy for demand can result in order quantities other than optimal, especially with a significant number of stockouts. Hill [10] performed analytical and numerical comparisons of the Frequentist approach against the Bayesian for three distributions (Exponential, Poisson and Binomial), using expected values of the optimal order quantity and the minimum expected cost, whose functional forms included estimates of the demand parameter. The same direction had the paper of Hill [11], which modeled the number of demands per customer per inventory cycle by a Poisson distribution with an unknown mean λ, and the size of an individual customer demand by an unknown distribution with known mean, second moment about the origin and variance. Bell [4] related the optimum order quantity to a set of recursive formulae, which updated forecasts for the average demand and the mean absolute deviation using exponential smoothing. To handle high percentages of stockouts, in the recursive form of the mean absolute deviation, Bell used the expected variance conditional upon the observed stockout.

Using a Bayesian Markov decision process, Ding et al. [8] showed that, with a general continuous demand distribution, the optimal inventory level in the presence of lost sales is higher as compared to the uncensored case. So, having available a higher inventory level in early inventory cycles, the additional information available for demand estimation can lead to better decisions in later inventory cycles. Similar conclusions were reported by Lariviere and Porteus [17] who considered Bayesian updating of demand distributions with unobserved lost sales using dynamic programming and modeling demand as a special case of an exponential distribution with a gamma conjugate prior. Wang and Webster [25] reported that in practice managers may have preferences other than the assumption of risk-neutrality, the latter based upon selecting the order quantity to maximize expected profit. So, using a single "kinked" piecewise-linear loss-aversion utility function to model managers' decision-making, Wang and Webster found that a loss-averse newsvendor facing low (high) shortage cost will order less (more) than the risk-neutral newsvendor. On the other hand, Wu et al. [26] showed that with stockout cost the risk-averse newsvendor does not necessarily order less than the risk-neutral newsvendor when the objective function has a mean-variance form, deriving the exact conditions under which this will happen when demand follows the power distribution. Finally, Chahar and Taaffe [6] showed that a conditional value-at-risk approach results in fewer worst-case profit scenarios as compared to the expected profit solution in the case where an order will either come in at a predefined level or will not come in at all.

In the current paper, we consider the classical single-period inventory model with demand sizes in consecutive inventory cycles to be independent and identically distributed normal random variables with the same mean and variance. Assuming that demand is fully observed for each period, a different approach for studying the effects of demand estimation on the optimal inventory policies is followed. This approach is focused on exploring the *variability* of estimates for the optimal order quantity and the maximum expected profit, which is *caused* by the sampling distribution of the estimated parameters of demand. Incorporating the maximum likelihood estimators for the mean and variance into the forms that determine the optimal order quantity and the maximum expected profit, we develop for the latter two variables appropriate estimators whose statistical properties are investigated. Although in finite samples both estimators are biased, we show that they are consistent, and asymptotically they converge to normality. Based on their asymptotic properties, we derive confidence intervals for the true optimal order quantity and the true maximum expected profit, whose validity is tested in small samples by developing appropriate Monte-Carlo simulations. Symbols and acronyms, which are used in the analysis, are explained in the text as and when required. Besides, for the reader's convenience, Appendix A provides a list of all the symbols and acronyms being used throughout this paper.

2. Estimators for the optimal order quantity and the maximum expected profit

Let $D_t = \mu + \varepsilon_t$ be the demand size for period t, with μ the average demand and ε_t's a sequence of independent and identically distributed normal random variables with zero mean and constant variance σ^2. Denote also by p the selling price per unit, c the purchase (or production) cost per unit, v the salvage value, and s the shortage penalty cost per unit. Under the previous specifications, the profit function per period for the classical newsboy problem is given in [14] as

$$\pi = \begin{cases} (p-c)Q - (p-v)(Q-D_t) & \text{if } Q - D_t \geq 0, \\ (p-c)Q + s(Q-D_t) & \text{if } Q - D_t < 0, \end{cases} \quad (1)$$

where Q is the order quantity. Denoting by φ_z and Φ_z the density function and the distribution function, respectively, of the standard normal evaluated at $z=(Q-\mu)/\sigma$, the expected value of (1) is derived in Silver et al. [24], and is given by

$$\begin{aligned} E(\pi) &= (p-c)\mu - (c-v)Q - (p-v+s)\sigma G(z) \\ &= (p-c)Q + s(Q-\mu) - (p-v+s)\{(Q-\mu)\Phi_z + \sigma\varphi_z\}, \end{aligned} \quad (2)$$

where $G(z)=\varphi_z-(Q-\mu)(1-\Phi_z)/\sigma$.

The optimal order quantity, Q^*, maximizing (2), satisfies the equation

$$\Phi_z = Pr\left(Z \leq \frac{Q-\mu}{\sigma}\right) = Pr(Z \leq z_R) = \frac{p-c+s}{p-v+s} = R,$$

which leads to $Q^*=\mu+z_R\sigma$, with z_R to be the inverse function of the standard normal evaluated at R. The ratio $R=(p-c+s)/(p-v+s)$ is called *critical fractile* [23]. Replacing Q with Q^* into (2), the