Available online at www.sciencedirect.com
SCIENCE DIRECT®
Int. J. Production Economics 92 (2004) 145–155
international journal of production economics
www.elsevier.com/locate/dsw

Performance of approximations for compound Poisson distributed demand in the newsboy problem

M.J.G. Dominey, R.M. Hill*

School of Mathematical Sciences, University of Exeter, Laver Building, North Park Road, Exeter EX4 4QE, UK

Received 5 February 2003; accepted 2 October 2003

Abstract

This paper considers the classical single period inventory model. The number of customer orders during the period follows a known Poisson distribution and individual customer order sizes are independent random variables. Two costs are incurred: a cost per unit of unsatisfied demand and a cost per unit of stock purchased. The objective is to minimise the expected sum of these two costs. A customer order which cannot be met in full is met to the extent that available stock permits. Although the sizes of all customer orders are known at the time only regularly updated estimates of the first two moments of the customer order size distribution are maintained. Therefore, aggregate demand follows a compound Poisson distribution for which the moments are known but for which the exact distribution is unknown. The immediate objective of this research is to explore the effectiveness of a number of approaches for approximating a compound Poisson distribution in a single period setting. The longer-term objective is to find relatively simple but effective ways of handling a compound Poisson demand process in more general inventory settings.

Keywords: Compound Poisson; Newsboy; Inventory; Single period; Gamma distribution

1. Introduction

The single period inventory (or 'newsboy') model is a standard problem which has been considered at some length in the literature (see, for example, Silver et al., 1998, pp. 385–392 or the fairly comprehensive literature review by Khouja, 1999) and the solution to the standard problem is well known. The newsboy problem is of some interest in its own right but its real value lies in its use as a building block for more complex systems. This paper considers possible approaches when demand is assumed to be compound Poisson but only the moments of the distribution are known.

A plausible model for many demand processes is that customer orders arrive as a Poisson process and the sizes of those orders are independent random variables from some positive integer-valued distribution. The rate of customer arrival may well vary through time but the number of arrivals during any given period may still follow a Poisson distribution with an appropriately estimated parameter. It is reasonable to assume that the distribution of customer order sizes changes less through time.

*Corresponding author. Tel.: 44-1392-264-464; fax: 44-1392-264-460.

E-mail address: r.m.hill@ex.ac.uk (R.M. Hill).

0925-5273/$ - see front matter
doi:10.1016/j.ijpe.2003.10.010

Many customer transaction processing systems, whether retail or industrial, permit customer order sizes to be monitored but few inventory control systems are set up to make full use of the data which become available. This may be because practical limitations on data storage prevent the holding of all past customer order sizes or it may be because control systems which make full use of this data are felt to be too complex or too demanding on processor time to implement. Another reason for wanting to make a smoothing distributional approximation is that raw customer demand data may be sparse and attaching undue emphasis to the few demand sizes which happen to have occurred may give a misleading picture of the nature of the underlying distribution.

In the single period context considered here the parameter of the (Poisson) number of customer arrivals is assumed to be known, as are the mean and the second moment (and hence the variance) of the customer order size distribution. Thus, for any particular stock line just three parameters describing the demand process are available for use.

Six possible approximate approaches are considered:

(i) to assume that aggregate period demand follows a normal distribution, using the available parameters to determine its mean and variance, and employing fairly standard numerical routines to determine the optimal starting stock;
(ii) as (i) but using a gamma distribution;
(iii) as (i) but using a lognormal distribution;
(iv) as (i) but using a mixed-Erlang distribution;
(iii) as (i) but now aggregate demand is assumed to follow a 'batch Poisson' distribution;
(iv) employing the distribution free approach suggested by Scarf in Arrow et al. (1958).

These approaches will be considered more fully in Section 2. The objective of this paper is to determine, through a numerical study, which of these six approaches produces the best policies, in terms of cost minimisation in a newsboy setting, when information on demand is only available in summary form.

2. Analysis

The number of customer arrivals during the period follows a Poisson distribution with mean λ. The sizes of individual customer orders are independent random variables from a distribution with probability function p_i, $i = 1, 2, \ldots$, mean μ_c, variance σ_c^2 and coefficient of variation c_c (a dimensionless measure of variability defined by $c_c^2 = \sigma_c^2/\mu_c^2$). The probability function $q(x)$ for the resulting compound Poisson distribution is then given by the recursive routine (see Adelson, 1966, or Tijms, 1994, pp. 27–29)

$$q(0) = \exp(-\lambda) \tag{1a}$$

and

$$q(x) = \frac{\lambda}{x}\sum_{i=1}^{x} i p_i q(x-i), \quad x = 1, \ldots, \tag{1b}$$

with mean $\mu_X = \lambda\mu_c$, variance $\sigma_X^2 = \lambda(\mu_c^2 + \sigma_c^2)$ and coefficient of variation c_X.

If the cost per unit of unsatisfied demand is L and the cost per unit of stock purchased is h then the expected cost per period given an initial stock of y is given by

$$H(y) = L\sum_{x=y}^{\infty}(x-y)q(x) + hy, \quad y = 0, \ldots, \tag{2a}$$

$$= L\lambda\mu_c + L\sum_{x=0}^{y-1}(y-x)q(x) + (h-L)y, \tag{2b}$$

where, here and elsewhere, a summation takes the value 0 if the lower limit is greater than the upper limit. Other formulations of the single period model which, for example, may assign a unit selling price and/or a cost per unit of stock leftover at the end of the period and have the objective of maximising expected profit can be re-formulated to have the same essential structure as (2), differing only by a fixed term and with a different unit cost ratio.

If h/L is written as c then $H(y)$ is minimised by the smallest value of y satisfying

$$\sum_{x=y+1}^{\infty} q(x) \leqslant c \tag{3a}$$