RC23805 (W0511-123) November 16, 2005
Mathematics

IBM Research Report

Improved Demand Parameter Estimators in Inventory Problems*

Kaan Katircioglu
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Derek Atkins
Faculty of Commerce and Business Administration
University of British Columbia
Vancouver, B.C. V6T 1Y8

(revised on October, 2005)

* This research was partially supported by a grant from the Natural Sciences and Engineering Research Council, Canada.

Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

Improving Traditional Inventory Policies Through Biased Estimation

Abstract

In stochastic inventory problems with unknown demand parameters, unknown parameters are typically replaced with their estimates before solving the inventory policy optimization problem. The hoep is that the policy determined this way is satisfactory. This paper presents an analytical method that can be used to enhance unbiased estimators that are commonly used. The method takes in to account the impact of estimation errors on the inventory cost function. The method is developed generally for a class of inventory problems including problems with cost minimization objectives and problems with service achievement objectives. The method is applicable to problems where the demand has a probability distribution which belongs to the scale-location family. Examples with Normal and Gamma demand cases are presented.

1 Introduction

The common approach to solving inventory problems presented in the majority of texts is to estimate demand by such means as the sample mean or sample standard deviation and then use these estimators as a substitute for the unknown paramater values.

The estimation typically derives from statistical procedures that equally weigh over and under estimation. Moreover, the costs of either having too much or too little inventory are seldom symmetrical.

The natural question arrives as to whether we might not be able to improve total costs by combining the two steps of estimation and optimization. This paper is dedicated to exploring this idea. Although we are not the first to tread this path, there has been little work that addresses and clarifies this matter. The most relevant papers in the literature are Hayes [6], Weerahandi [14], Silver and Rahmana [12, 13], and Ritchken and Sankar [11].

Hayes [6] essentially studies a newsboy problem. He derives the expected total operating cost (ETOC) for exponential demand and then calculates an optimal bias (a multiplier) for the sample mean of demand observations and uses this as a biased estimator for the mean demand so as to minimize the ETOC. He also works with Normal demand with unknown mean, known variance and, both unknown mean and variance. In both cases, he approximates the ETOC by Taylor's expansion. Then, using this approximate cost, he calculates an optimal bias. Unlike the traditional method, this approach combines both estimation and optimization. He has more details and examples in his thesis [7]. His work is particularly thorough, but unhappily seems to have been ignored or overlooked by the papers that followed, which studied many of the same topics but added little.

Weerahandi [14] also considers the newsboy problem. He, too (but without any approximation), determines an optimal biasing factor for the standard deviation of demand, when demand is normally distributed. Although Hayes uses an approximate cost, his bias is the same. Weerahandi also provides an optimal bias for the sample mean if demand is gamma.

Silver and Rahmana [12] examined the cost penalties caused by estimating the demand distribution parameters in the case of a normal distribution. They use a variant of a (Q, r) model with Q being a predetermined constant. Later, they presented a numerical algorithm to calculate the optimal bias that should be used to overestimate the reorder point so as to minimize the cost penalty.

Ritchken and Sankar [11] consider a newsboy problem with two types of constraints. In case of a probability of stock-out constraint, they give a correction factor for the sample standard deviation to achieve the targeted service level for a normally distributed demand. They develop a regression approach for the second type of constraint which ensures that a desired percentage of demand (fill rate) is satisfied with a certain level of significance.

There have been numerous papers dealing just with the estimation of demand which are not directly relevant here. There are a few papers somewhat relevant to the context of this work.

Bookbinder and Lordahl [2] use a distribution free bootstrapping method which bypasses the need for parameter estimates to estimate the fractiles of demand distribution, in order to identify inventory reorder levels. In a more recent paper by Lordahl and Bookbinder [10], order-statistics are used for the same objective.

Jacobs and Wagner [8] showed that using exponentionally smoothed estimators of the mean and variance gives a smaller total system cost for the (S,s) model as compared to using classical unbiased estimators, sample mean and variance, when the demand variability is large. In another paper [9], they presented some simulation results which show that if there is a significant functional relationship between the mean and variance of the demand, a regression-based estimator of the variance may perform even better than the sample variance in the *power approximation* suggested by Ehrhardt [3].

Biggs and Campion [1] showed by simulating an MRP-based production system with capacity constraints that utilizing a biased (a positive forecast error) demand forecast yields better performance measures (these are the quantities that constitute the system cost) than using an unbiased forecast.

In this work, we attempt to place all of the work in as general a framework as possible, where necessary reproving existing results and then generalizing and extending them. We consider problems with either cost or service objectives when the demand has a scale-location family type of distribution.

First, a summary of three inventory problems whose cost functions have similar forms for our purpose is presented. Then, Hayes' alternative approach to the traditional method of determining the optimal policy is developed in a more general context. Finally, the general theory is applied to a demand whose distribution is in the scale-location family and some analytical results are derived for Normal and Gamma demand.

For this method to be useful, the proposed biasing factor to be applied to the usual estimates needs to be independent of any unknown parameters. We prove this to be true for the optimal bias when using the scale-location family in terms of both cost and service objectives. This independence had already been obtained for the Normal and Gamma cases by Weerahandi.

When a service objective is considered, we add to the results of Ritchken and Sankar by showing that for a demand in the scale-location family, one can calculate a biased estimate which will guarantee the achievement of a target service level without having to know the true demand parameters.

Finally, we give exact expressions for the optimal bias and cost in a (Q, r) model with normal demand eliminating the need for Silver and Rahmana's numerical algorithm.

2 Inventory Models

To enhance the generality of our results we shall use the following cost function, where y is the inventory level (decision variable), X is the random demand with mean μ.

$$\begin{aligned} F(y;\theta) &= AE_X[y-X]^+ + B(\mu-y) + Cy + D \qquad (1) \\ &= A\int_0^y \Phi(x|\theta)\, dx + B(\mu-y) + Cy + D \end{aligned}$$

where $\phi(.)$ and $\Phi(.)$ are *pdf* and *cdf* of X, respectively, and constants A, B, C, D are such that $A > B-C > 0$. We first give three examples of this general form.

2.1 Newsboy Problem

In this single-period An initial inventory y is purchased at a price c, and sold at a price p. Unsold units are salvaged at a value s per unit ($p > c > s$). It can readily be found that maximizing expected profit for this problem is equivalent to minimizing the expected cost of the form in (1) with $A = p - s$, $C = c - p$, and $B = D = 0$.

2.2 A Base-Stock Model

This is a continuous review inventory model in which demands arrive in a continuous stream and unmet demands are fully backlogged. Replenishment orders maintain the inventory position (stock on hand + on order - backorders) at a critical level (base-stock). This policy always maintains the inventory position at this base-stock level. There is a positive replenishment lead time. Inventory holding (h) and backlogging costs (p) are linear with no ordering cost. The objective is to determine an optimal base-stock level so as to minimize the long run average holding and backlogging costs.

Let us denote the critical base-stock level by y and let X be a continuous random demand during the lead time, with mean μ. At any time, all outstanding orders a lead time ago will have arrived. Therefore, the inventory on hand at any time will have the same distribution as $(y-X)^+$, and the backorders will have the same distribution as $(X-y)^+$. The long run average cost is in the general form (1) with $A = h + p$, $B = p$, and $C = D = 0$.

2.3 (Q, y) Model With Fixed Q

In this model, When the inventory position is at or below y (reorder level) we make an order of fixed size Q. The demand during lead time is X. The annual demand λ is assumed to be fixed and known. Each order costs K dollars. Unfilled demand is backlogged and each unit backordered costs π (\$/unit). The annual inventory holding cost is h (\$/unit).

Using Hadley and Whitin's [5] approximation (which is also used by Silver and Rahmana [12, 13]), the general cost function in (1) can be recovered by letting $A = \pi\lambda/Q$, $B = (\pi\lambda/Q) - h$, $D = (K\lambda/Q) + (hQ/2)$, and $C = 0$.

3 An Alternative Approach

3.1 Cost Objective

The objective for the control of inventory is to minimize the cost function $F(y;\theta)$, given the demand parameter θ.

$$\min_{y} \quad F(y;\theta) \ . \tag{2}$$

$F(y;\theta)$ is easily proved convex, so using the first derivative and letting $M \equiv (B - C)/A$, the solution to (2) can be written as

$$y^*(\theta) \quad = \Phi^{-1}(M|\theta) \ . \tag{3}$$

If θ is unknown, the standard practical approach is to sequentially first estimate θ by $\hat{\theta}$ and then optimize to find $y^*(\hat{\theta})$. Clearly, $F(y;\theta)$ need not attain its minimum at $y = y^*(\hat{\theta})$, because $y^*(\hat{\theta})$ can be different from $y^*(\theta)$ at which it is minimized.

The question is whether it is possible to improve this solution by trying different estimators for θ. Since $F(y^*(\hat{\theta});\theta)$ is a random variable, the expected value of $F(y^*(\hat{\theta});\theta)$ will be the expected cost associated with the policy $y^*(\hat{\theta})$. Thus, we want to solve

$$\min_{\hat{\theta}\in\Theta} E_{\hat{\theta}} \left[F(y^*(\hat{\theta});\theta) \right]$$

where Θ is a class of estimators in which we are interested. Since we do not want Θ to be too general to work with, we restrict it as follows

$$\Theta = \{\theta : \theta = \omega\hat{\theta}, \omega \in R\} \ ,$$

where $\hat{\theta}$ is a given estimator of θ. Therefore, we choose to focus on the following problem.

$$\min_{\omega} E_{\hat{\theta}} \left[F(y^*(\omega\hat{\theta});\theta) \right] . \tag{4}$$

Problem 4 can be interpreted as follows: Find the optimal biasing factor for $\hat{\theta}$ (that is ω) so as to obtain a better solution $y^*(\omega\hat{\theta})$ at which the expectation of $F(y^*(\omega\hat{\theta});\theta)$ is as close as possible to the unknown minimum $F(y^*(\theta);\theta)$.

Let $G(\hat{\theta})$ be the distribution function of $\hat{\theta}$, that is, $G(\hat{\theta})$ gives the sampling distribution of $\hat{\theta}$. Then, the expectation is

$$E_{\hat{\theta}} \left[F(y^*(\omega\hat{\theta});\theta) \right] = \int_{\hat{\theta}} F(y^*(\omega\hat{\theta});\theta)\, dG(\hat{\theta}).$$

Observe that if $y^*(\omega\hat{\theta})$ is linear in ω, and if the expectation $E_{\hat{\theta}}\left[F(y^*(\omega\hat{\theta});\theta)\right]$ exists, then it is convex in ω for all θ. This follows because $F(y;\theta)$ is convex in y for all θ and $y^*(\omega\hat{\theta})$ is linear in ω, $F(y^*(\omega\hat{\theta});\theta)$ is convex in ω for all θ as the convexity of a function is preserved by a linear transformation of its argument.

Proposition 1 *Let $F(y;\theta)$ be given by (1) with $A > B - C > 0$, and let $y^*(\omega\hat{\theta})$ be linear in ω. If $E_{\hat{\theta}}\left[F(y^*(\omega\hat{\theta});\theta)\right]$ exists, then, the optimal solution to (4), say ω^*, satisfies*

$$\int_{\hat{\theta}} \Phi(y^*(\omega\hat{\theta})|\theta)\,\hat{\theta}\,dG(\hat{\theta}) - M\,E_{\hat{\theta}}[\hat{\theta}] = 0.$$

Proof: Since $y^*(\omega\hat{\theta})$ is linear in ω, let $\frac{\partial}{\partial\omega}y^*(\omega\hat{\theta}) = a\hat{\theta}$, where a is a real number. The first derivative of the expected cost in (4) is

$$\begin{aligned}
\frac{\partial}{\partial\omega}E_{\hat{\theta}}\left[F(y^*(\omega\hat{\theta});\theta)\right] &= \frac{\partial}{\partial\omega}\int_{\hat{\theta}} F(y^*(\omega\hat{\theta});\theta)\,dG(\hat{\theta}) \\
&= \int_{\hat{\theta}} \frac{\partial}{\partial\omega}\left(A\int_0^{y^*(\omega\hat{\theta})} \Phi(x|\theta)\,dx + B(\mu - y^*(\omega\hat{\theta})) + Cy^*(\omega\hat{\theta}) + D\right) dG(\hat{\theta}) \\
&= \int_{\hat{\theta}} \left(A\Phi(y^*(\omega\hat{\theta})|\theta)a\hat{\theta} + (C-B)a\hat{\theta}\right) dG(\hat{\theta}) \\
&= aA\int_{\hat{\theta}} \left(\Phi(y^*(\omega\hat{\theta})|\theta)\hat{\theta} - M\hat{\theta}\right) dG(\hat{\theta}).
\end{aligned}$$

Setting it equal to zero, we get the optimality condition. The second derivative of the expected cost is positive and thus the solution is a minimum. □

3.2 Service Objective

Service objective refers to the situation where stockouts are controlled by a service level constraint rather than a stockout cost. There are a number of different service level definitions - here we use the probability that the demand is satisfied in a certain period of time (usually lead time) is at least α. We assume that the cdf of demand, $\Phi(y|\theta)$ is both continuous and strictly increasing in y. Then, the minimum inventory, y, that achieves a service level α will be unique and given by

$$\Phi(y|\theta) = \alpha.$$

The solution is $y_c^*(\theta) = \Phi^{-1}(\alpha|\theta)$. Now, if θ is known, $\Phi(y_c^*(\theta)|\theta) = \alpha$ is the actual service level that would be realized when $y_c^*(\theta)$ is the inventory policy. However, if θ is unknown, the traditional approach has been to replace θ by $\hat{\theta}$ and use $y_c^*(\hat{\theta}) = \Phi^{-1}(\alpha|\hat{\theta})$. Thus, the actual but unknown service level in this case would be

$$\Phi(y_c^*(\hat{\theta})|\theta).$$

which does not necessarily equal α. As before, we can modify this traditional method by introducing a bias ω to $\hat{\theta}$ so as to achieve the targeted service level α. That is,

$$E_{\hat{\theta}}[\Phi(y_c^*(\omega\hat{\theta})|\theta)] = \alpha, \tag{5}$$

Thus, denoting the solution of (5) by ω_c^*, if it exists, using $\omega_c^*\hat{\theta}$ instead of just $\hat{\theta}$ would result in a long run average service level of α, as required.

4 Scale-Location Family

In order to get more specific results, we make some assumptions for the demand distribution.

Definition 1 *Let $\Phi(x|\beta,\theta)$ be the cdf of a random demand X with parameters β and θ. Then, we say that X belongs to the scale-location family with scale parameter θ and location parameter β if its cdf can be written as*

$$\Phi(x|\beta,\theta) = \Phi(\tfrac{x-\beta}{\theta}|0,1), \tag{6}$$

where $\Phi(.|0,1)$ is called the standardized cdf, which does not depend on (β,θ).

It is easy to see from Definition 1 that

$$\phi(x|\beta,\theta) = \frac{1}{\theta}\phi(\tfrac{x-\beta}{\theta}|0,1). \tag{7}$$

The scale-location family includes the Normal, Gamma and Weibull with fixed shape parameters, Cauchy, Uniform, Logistic, Student's t and Laplace distributions.

Cost objective: Now, Problem (2) becomes

$$\min_{y} \quad F(y;\beta,\theta). \tag{8}$$

Noting in (6) that $\Phi(M|\beta,\theta) = \Phi(\frac{M-\beta}{\theta}|0,1)$, the solution in (3) becomes $y^*(\beta,\theta) = k\theta+\beta$, where $k = \Phi^{-1}(M|0,1)$ is a known constant.

If θ and β are replaced by their estimates $\hat{\theta}$ and $\hat{\beta}$, then the solution will be

$$y^*(\hat{\beta},\hat{\theta}) = k\hat{\theta}+\hat{\beta}. \tag{9}$$

We restrict our class of estimators to

$$\Theta = \{(\hat{\beta},\omega\hat{\theta}) : \omega \in R,\ \hat{\beta},\hat{\theta} \text{ are given}\},$$

although other choices are possible. This satisfies the linearity assumption of $y^*(.)$ in ω.

From Proposition 1, the optimality condition for ω will be

$$\int_{\hat{\beta}} \int_{\hat{\theta}} \hat{\theta} \left(\Phi(\frac{k\omega\hat{\theta} + \hat{\beta} - \beta}{\theta} | \, 0, 1) - M \right) \, dG_{\hat{\theta},\hat{\beta}}(\hat{\theta}, \hat{\beta}) = 0, \tag{10}$$

where $G_{\hat{\theta},\hat{\beta}}(\hat{\theta}, \hat{\beta})$ is the joint distribution function of $\hat{\theta}$ and $\hat{\beta}$, and thus implementable.

The following proposition shows that the optimal bias is independent of θ and β.

Proposition 2 *Let the cdf of X, $\Phi(x|\beta, \theta)$, belong to the scale-location family with location parameter β and scale parameter $\theta > 0$. Note that*

$$E(X) = a\theta + \beta, \text{ and } Var(X) = b\theta^2,$$

where a and $b > 0$ are constants. Suppose we estimate $E(X)$ by the sample mean $\bar{x}$ and $Var(X)$ by the sample variance s^2 (if one of θ and β is known, we estimate only $E(X)$).

Then, if ω satisfies (10), it is independent of (θ, β).

Proof: We have $\hat{\theta} = s/\sqrt{b}$ and $\hat{\beta} = \bar{x} - a\hat{\theta}$. Let $z_i = (x_i - \beta)/\theta$ and $\bar{z} = \frac{1}{n}\sum_{i=1}^{n} z_i$. Note that both z_i and $\bar{z}$ are independent of (θ, β). Now,

$$Prob\{\frac{\bar{x} - \beta}{\theta} \leq t\} \; = \; Prob\{\frac{1}{n}\sum_{i=1}^{n}\left(\frac{x_i - \beta}{\theta}\right) \leq t\} \; = \; Prob\{\bar{z} \leq t\}.$$

$$\begin{aligned} Prob\{\frac{\hat{\theta}}{\theta} \leq t\} &= Prob\{\frac{1}{\theta}\sqrt{\sum_{i=1}^{n}\frac{(x_i - \bar{x})^2}{n-1}} \leq t\sqrt{b}\} \\ &= Prob\{\sqrt{\sum_{i=1}^{n}\left(\frac{x_i - \beta}{\theta} - \frac{\bar{x} - \beta}{\theta}\right)^2} \leq t\sqrt{b(n-1)}\} \\ &= Prob\{\sqrt{\sum_{i=1}^{n}(z_i - \bar{z})^2} \leq t\sqrt{b(n-1)}\}. \end{aligned}$$

Thus, $(\bar{x} - \beta)/\theta$ and $\hat{\theta}/\theta$ are independent of (θ, β). The optimality condition in (10) is

$$\int_{\hat{\beta}} \int_{\hat{\theta}} \hat{\theta} \left(\Phi((k\omega - a)\frac{\hat{\theta}}{\theta} + \frac{\bar{x} - \beta}{\theta} | \, 0, 1) - M \right) dG_{\hat{\theta},\hat{\beta}}(\hat{\theta}, \hat{\beta}) = 0.$$

Let $u = \hat{\theta}/\theta$, $v = (\bar{x} - \beta)/\theta$ and let $H_{u,v}(u, v)$ be the joint distribution of (u, v). Then, dividing both sides by θ, the optimality condition becomes

$$\int_{v} \int_{u} u \left(\Phi((k\omega - a)u + v | \, 0, 1) - M \right) dH_{u,v}(u, v) = 0$$

which is independent of (θ, β). $\square$

An important implication of Proposition 2 is stated in the following corollary.

Corollary 1 *Let $\mathcal{S}$ be the set of all possible values of (θ,β) for which $F(.\ ;\theta,\beta)$ is defined. Then, if ω^* satisfies (10), it solves*

$$\min_{\omega\in R}\max_{(\theta,\beta)\in\mathcal{S}} E_{\hat\theta,\hat\beta}[F(y^*(\hat\beta,\omega\hat\theta);\theta,\beta)].$$

Proof: From (9), $\frac{\partial}{\partial\omega}y^*(\omega\hat\theta)=k\hat\theta$. Then, from Proposition 1, we have

$$\frac{\partial}{\partial\omega}E_{\hat\theta,\hat\beta}\left[F(y^*(\hat\beta,\omega\hat\theta);\theta,\beta)\right] = kA\int_{\hat\beta}\int_{\hat\theta}\hat\theta\left(\Phi(y^*(\hat\beta,\omega\hat\theta);\theta,\beta)-M\right)dG_{\hat\theta,\hat\beta}(\hat\theta,\hat\beta).$$

Note that the integral is equivalent to the one in (10). Then, using the last equation in the proof of Proposition 2, we can write

$$\frac{\partial}{\partial\omega}E_{\hat\theta,\hat\beta}\left[F(y^*(\hat\beta,\omega\hat\theta);\theta,\beta)\right] = kA\theta\int_v\int_u u\left(\Phi((k\omega-a)u+v|\,0,1)-M\right)dH_{u,v}(u,v).$$

Let the term in brackets be $g(\omega)$. Proposition 2 implies that $g(\omega)$ is independent of (θ,β). Let $G(\omega)$ be such that $dG(\omega)=g(\omega)d\omega$. Then, we can write the expected cost as

$$E_{\hat\theta,\hat\beta}\left[F(y^*(\omega\hat\theta);\theta)\right]=kA\theta G(\omega)+l(\beta,\theta),$$

where $l(\beta,\theta)$ is a function independent of ω. Let $(\theta^*(\omega),\beta^*(\omega))$ maximize this expectation. Substituting $\theta^*(\omega)$ and $\beta^*(\omega)$, and differentiating with respect to ω, we have

$$\begin{aligned}\frac{\partial}{\partial\omega}\max_{(\theta,\beta)\in\mathcal{S}}E_{\hat\theta,\hat\beta}[F(y^*(\hat\beta,\omega\hat\theta);\theta,\beta)] &= kA\frac{\partial}{\partial\omega}\theta^*(\omega)G(\omega)+kA\theta^*(\omega)g(\omega)\\ &+ \frac{\partial}{\partial\theta}l(\beta^*(\omega),\theta^*(\omega))\frac{\partial}{\partial\omega}\theta^*(\omega)+\frac{\partial}{\partial\beta}l(\beta^*(\omega),\theta^*(\omega))\frac{\partial}{\partial\omega}\beta^*(\omega).\end{aligned}$$

Now, if $\beta^*(\omega)$ is an interior solution, then $\frac{\partial}{\partial\beta}l(\beta^*(\omega),\theta^*(\omega))=0$. If it is a boundary solution, this will imply $\frac{\partial}{\partial\omega}\beta^*(\omega)=0$. In both cases, the last term vanishes in the equation. Similarly, if $\theta^*(\omega)$ is an interior solution, then $kAG(\omega)+\frac{\partial}{\partial\theta}l(\beta^*(\omega),\theta^*(\omega))=0$. If it is a boundary solution, this implies $\frac{\partial}{\partial\omega}\theta^*(\omega)=0$. In both cases, the equation will be

$$\frac{\partial}{\partial\omega}\max_{(\theta,\beta)\in\mathcal{S}}E_{\hat\theta,\hat\beta}[F(y^*(\hat\beta,\omega\hat\theta);\theta,\beta)]=kA\theta^*(\omega)g(\omega).$$

The second derivative is

$$\begin{aligned}\frac{\partial^2}{\partial\omega^2}\max_{(\theta,\beta)\in\mathcal{S}}E_{\hat\theta,\hat\beta}[F(y^*(\hat\beta,\omega\hat\theta);\theta,\beta)] &= kA\theta^*(\omega)\frac{\partial}{\partial\omega}g(\omega)\\ &= kA\theta^*(\omega)\int_v\int_u uk\phi((k\omega-a)u+v|\,0,1)dH_{u,v}(u,v).\end{aligned}$$

Note that for ω satisfying the first order condition, $g(\omega)=0$, the second derivative is positive. Hence the optimality occurs when $g(\omega)=0$, which is equivalent to (10). □

Corollary 1 implies that ω^* is also a *minimax* policy among the policies of the form $y^*(\hat\beta,\omega\hat\theta)=k\omega\hat\theta+\hat\beta$. Thus, ω^* not only minimizes the average cost but also the worst possible average cost. This could be useful for problems with budget limitations.

Service objective: Let $l = \Phi^{-1}(\alpha|0,1)$. Then, the target service achieving bias ω_c^* which solves (5) will satisfy

$$\int_{\hat{\beta}}\int_{\hat{\theta}} \Phi(\frac{l\omega\hat{\theta} + \hat{\beta} - \beta}{\theta}|\,0,1)\, dG_{\hat{\theta},\hat{\beta}}(\hat{\theta},\hat{\beta}) = \alpha, \tag{11}$$

It follows from the proof of Proposition 2 that ω_c^* is also independent of (θ,β).

4.1 Normal Demand Case

Let $\phi(.|\mu,\sigma)$ and $\Phi(.|\mu,\sigma)$ be the *pdf* and *cdf* of the Normal distribution with mean μ and standard deviation σ. We denote their standard versions by $\phi(.)$ and $\Phi(.)$. Suppose we have a sample of demand observations for n periods. We use sample mean $\bar{x}$ and variance s^2 to estimate μ and σ^2, respectively. From (9), the inventory policy is

$$y^*(\bar{X},s) = ks + \bar{X}. \tag{12}$$

Now, we want to bias s by ω and use ωs as an estimator of σ, whereas $\bar{x}$ estimates μ. From (10), the optimal ω satisfies

$$\int_s\int_{\bar{x}} s\Phi(\frac{k\omega s + \bar{x} - \mu}{\sigma})\, dG_{\bar{x},s}(\bar{x},s) = M\int_s\int_{\bar{x}} s\, dG_{\bar{x},s}(\bar{x},s), \tag{13}$$

where $G_{\bar{x},s}(\bar{x},s)$ is the joint *cdf* of $\bar{x}$ and s.

Using lemmas 4 and 5, after some algebra, it can be shown that the optimality condition reduces to

$$\omega^* = \frac{T_n^{-1}(M)}{\Phi^{-1}(M)}\sqrt{1-\frac{1}{n^2}} \tag{14}$$

where $T_n^{-1}(.)$ is the inverse cdf of *Student's t* distribution with n degrees of freedom. It is easy to see that as n becomes large, ω^* tends to 1, as one would expect. This expression was obtained by Weerahandi for the newsboy problem and is identical to that found by Hayes [6], although there, an approximation to the expectation of $F(y^*(\bar{X},\omega s);\mu,\sigma)$ was made with a Taylor's expansion about $y^*(\mu,\sigma)$. It is also important to note that ω^* is independent of σ and μ and hence implementable as predicted in Theorem 2 above. The sensitivity of ω^* to values of M and n is shown in Table 1. Note that for small sample sizes, a significant amount of bias is required to minimize the expected costs especially when M is high. In all models we consider, a high value of M corresponds to a high value of the unit shortage cost compared to unit inventory holding cost (for instance, in the base-stock model, $M = p/(p+h)$ is high when backlogging cost p is high). Therefore, when sample sizes are small and shortages cost significantly more than holding inventory, the amount of bias is high. For example, an over 40% increase in the sample standard deviation is required for $M = 0.99$ when the sample size is 5.

From (12), the optimal inventory policy adjusted for estimation errors is thus

$$y^*(\bar{X},\omega^* s) = \bar{X} + s\,T_n^{-1}(M)\sqrt{1-\frac{1}{n^2}}.$$

Now, we state the following lemma for the expectation of the cost function.

	Optimal Bias ω^* for Normal			
M	n=5	n=10	n=15	n=20
0.10	1.128	1.065	1.044	1.033
0.30	1.045	1.027	1.019	1.015
0.50	0.980	0.995	0.998	0.999
0.90	1.128	1.065	1.044	1.033
0.95	1.200	1.096	1.063	1.047
0.99	1.417	1.182	1.116	1.085

Table 1: The value of optimal bias for normal distribution.

Lemma 1 *For any bias ω,*

$$E_{\bar{X},s}[F(y^*(\bar{X},\omega s);\mu,\sigma)] = a_n(\omega)A\sigma + C\mu + D \ ,$$

where

$$a_n(\omega) = \sqrt{\frac{n+1}{2\pi n}}\left(1+\frac{nk^2\omega^2}{n^2-1}\right)^{-\frac{n-1}{2}} + \sqrt{\frac{2}{n-1}}\frac{\Gamma(\frac{n}{2})}{\Gamma(\frac{n-1}{2})}k\omega\left[T_n(\frac{nk\omega}{\sqrt{n^2-1}}) - M\right].$$

Proof: See Appendix A. □

Note that $a_n(\omega)$ depends only on M and n and that the expectation of the cost function is linear in σ and μ.

When $\omega = \omega^*$, $a_n(\omega)$ reduces to

$$a_n(\omega^*) = \sqrt{\frac{n+1}{2\pi n}}\left(1+\frac{[T_n^{-1}(M)]^2}{n}\right)^{-\frac{n-1}{2}}.$$

The percent reduction in cost is given by

$$\frac{[a_n(1) - a_n(\omega^*)]A\sigma}{a_n(1)A\sigma + C\mu + D}(100).$$

Note that for the base-stock model, since $C = D = 0$, the percent reduction in the total cost is independent of μ and σ.

Since D is a constant, the percent reduction in the controllable portion of the cost (*i.e.* $F(y;\theta) - D$ in (1)) can be calculated by letting $D = 0$, which depends only on the ratio σ/μ. That is

$$\frac{[a_n(1) - a_n(\omega^*)]A\delta}{a_n(1)A\delta + C}(100).$$

where $\delta \equiv \sigma/\mu$ is the coefficient of variation. In the (Q, r) model we have $C = 0$. Thus, the percent reduction in the controllable cost does not depend on the unknown parameters.

For the service objective, the following lemma gives the optimal bias ω_c^*.

Lemma 2 *If the demand is normally distributed with mean μ and variance σ^2, then Equation 11 is satisfied by*

$$\omega_c^* = \frac{T_{n-1}^{-1}(\alpha)}{\Phi^{-1}(\alpha)}\sqrt{1+\frac{1}{n}}\ .$$

Proof: See Appendix A. □

This is previously obtained by Ritchken and Sankar [11]. The expected cost of $y^*(\omega_c^* s)$ can be calculated by using Lemma 1 with $\omega = \omega_c^*$.

As can be seen in the proof of Lemma 2, letting $l = \Phi^{-1}(\alpha)$, the expected service level is

$$E_{\bar{x},s}[S(\bar{x},\omega s)] = T_{n-1}(\frac{l\omega n}{n+1}), \tag{15}$$

which equals α for $\omega = \omega_c^*$. The usefulness of Lemma 2 is that the bias ω_c^* can be applied without any knowledge of the parameters μ and σ, and the long run average service level is guaranteed to be α.

4.2 Gamma Demand Case

Let $\phi_r(.|\gamma)$ and $\Phi_r(.|\gamma)$ denote the *pdf* and *cdf*, respectively, for the Gamma distribution with shape parameter r, scale parameter γ and location parameter 0. We use the notation $X \sim Gamma(r,\gamma)$ to mean that

$$\phi_r(x|\gamma) = \frac{(1/\gamma)^r}{\Gamma(r)} x^{r-1} e^{-x/\gamma}, \quad x > 0.$$

Note that $E(X) = r\gamma$, and that $\hat{\gamma} = \frac{1}{nr}\sum_{i=1}^n x_i$ is an unbiased estimator for γ. Further, $\sum_{i=1}^n x_i \sim Gamma(nr,\gamma)$. Thus, $\hat{\gamma} \sim Gamma(nr, \gamma/nr)$.

If we use a biased estimator of the form $\omega\hat{\gamma}$ for γ, from (10), the optimal value of ω will satisfy

$$\int_{\hat{\gamma}} \hat{\gamma}\, \Phi_r(\frac{k\omega\hat{\gamma}}{\gamma}|1)\, \phi_{nr}(\hat{\gamma}|\gamma/nr)\, d\hat{\gamma} = M\gamma\ ,$$

where $k = \Phi_r^{-1}(M|1)$. Using Lemma 6, it can be shown that the optimal bias is

$$\omega^* = \frac{nr\mathcal{B}_{r,nr+1}^{-1}(M)}{k(1-\mathcal{B}_{r,nr+1}^{-1}(M))}\ . \tag{16}$$

Weerahandi [14] has the same bias except that in his, $nr+1$ appears to be incorrectly stated as nr. Again note that ω^* is independent of γ and is thus implementable. From (9), the optimal inventory policy adjusted for estimation errors would be

$$y^*(\omega^*\hat{\gamma}) = \frac{\mathcal{B}_{r,nr+1}^{-1}(M)}{1-\mathcal{B}_{r,nr+1}^{-1}(M)} n\bar{x}. \tag{17}$$

The optimal bias is shown in Table 2 for various values of n, r and M. We see similar results as in the case of normal distribution. For small sample sizes and high M values (indicating high shortage costs relative to inventory holding costs), the bias is significantly bigger than 1, which means that an increased

	Optimal Bias ω^* for Gamma					
	r=1		*r=3*		*r=8*	
M	n=5	n=20	n=5	n=20	n=5	n=20
0.10	0.841	0.955	0.913	0.977	0.950	0.987
0.50	0.883	0.968	0.958	0.989	0.984	0.996
0.90	1.016	1.007	1.039	1.012	1.033	1.009
0.95	1.081	1.024	1.072	1.019	1.048	1.013
0.99	1.254	1.065	1.147	1.037	1.086	1.022

Table 2: The value of optimal bias for gamma distribution.

value of the sample mean should be used. However, for problems with higher inventory costs (*i.e.* small M value), the bias is smaller than 1, in which case a decreased sample mean should be used.

Now, the following lemma states the expected cost.

Lemma 3 *For any bias ω, $E_{\hat{\gamma}}[F(y^*(\omega\hat{\gamma});\gamma)] = a_n(\omega)\gamma r + D$, where*

$$a_n(\omega) = \frac{Ak\omega}{r}\left(\mathcal{B}_{r,nr+1}\left(\frac{k\omega}{k\omega+nr}\right) - M\right) - A\mathcal{B}_{r+1,nr}\left(\frac{k\omega}{k\omega+nr}\right) + B\ .$$

Proof: See Appendix A. □

For $\omega = \omega^*$, $a_n(\omega)$ simplifies to

$$a_n(\omega^*) = B - A\mathcal{B}_{r+1,nr}\left(\frac{k\omega^*}{k\omega^*+nr}\right).$$

The percent reduction is, then, given by

$$\frac{[a_n(\omega^*) - a_n(1)]\,\gamma r}{a_n(1)\gamma r + D}(100).$$

Note that when $D = 0$ (as in newsboy and base-stock models) this is independent of γ, which means the reduction in the controllable cost is not affected by the unknown parameter γ, as in the case of Normal distribution.

As for the service objective, using Equation 11 we have

$$\int_{\hat{\gamma}} \Phi_r(\frac{l\omega\hat{\gamma}}{\gamma}|1)\, dG_{\hat{\gamma}}(\hat{\gamma}) = \alpha\,.$$

Using Lemma 6 the integral is simply

$$E_{\hat{\gamma}}\left[\Phi_r(\frac{l\omega\hat{\gamma}}{\gamma}|1)\right] = \mathcal{B}_{r,nr}(\frac{l\omega}{l\omega+nr}).$$

Then, the bias that will yield an average service level of exactly α will be

$$\omega_c^* = \frac{nr\mathcal{B}_{r,nr}^{-1}(\alpha)}{l(1-\mathcal{B}_{r,nr}^{-1}(\alpha))}, \tag{18}$$

where $l = \Phi_r^{-1}(\alpha|1)$.

5 (Q, r) Model with Normal Daily Demand

In this section we will give exact closed form expressions for the optimal bias and the expected cost if one has a (Q, r) model with a given fixed Q. We assume daily demand data are used for estimation as in Silver and Rahmana [12]. The derivations of these expressions follow from Lemmas (4), (5), and are done very similarly to the previous calculations.

The optimal reorder level, from (9), will be

$$y_d^*(\bar{d}, \omega s_d) = L\bar{d} + k\omega\sqrt{L}s_d \ ,$$

where $\bar{d}$, s_d are the sample mean and standard deviation of daily demand data, and L is a fixed lead time in days. Assuming that the demand during a day has normal distribution with mean μ and variance σ^2, the optimal bias and the expected cost can be shown to be

$$\omega_d^* = \frac{T_n^{-1}(M)}{k}\sqrt{\frac{(n-1)(n+L)}{n^2}},$$

$$E_{\bar{d},s_d}[F(y_d^*(\bar{d}, \omega s_d); \mu, \sigma)] = a_{d,n}(\omega)A\sigma\sqrt{L} + C\mu L + D$$

where

$$\begin{aligned} a_{d,n}(\omega) \quad = \quad & \sqrt{\frac{n+L}{2\pi n}}\left(1 + \frac{nk^2\omega^2}{(n-1)(n+L)}\right)^{-\frac{n-1}{2}} \\ & + \sqrt{\frac{2}{n-1}}\frac{\Gamma(\frac{n}{2})}{\Gamma(\frac{n-1}{2})}k\omega\left[T_n\left(\frac{nk\omega}{\sqrt{(n-1)(n+L)}}\right) - M\right]. \end{aligned}$$

These closed form expressions eliminate the need for Silver and Rahmana's algorithm.

5.1 An Example

We conclude with an example in which we will examine a (Q, r) model with fixed Q, no order cost ($K = 0$) and Normal demand as in Silver and Rahmana [13].

a) Cost objective: We let the annual demand $\lambda = 1000$, and the annual holding cost $h = \$1$ per unit. It is assumed that the estimation of lead time demand parameters is done using a daily demand data for a fixed lead time L, and that the actual daily demand parameters are $\mu = 3$, $\sigma = 3/4$, when calculating the percent reduction in total cost.

Table 3 shows the effect of biasing the standard deviation. We have previously seen that the percent reduction in controllable cost does not depend on the unknown demand parameters (if the demand is normal). Note that as the cost of a backorder, π, increases the effect of using a bias factor is more significant. The percentage gain in cost is more significant for larger lead times, and intuitively enough, this gain decreases with sample size. In short, biasing is the most effective when we have small sample size, high backorder

n	L	π	*Optimal Bias and % Reduction*					
			Q=15			*Q=30*		
			ω_d^*	$\%R_1$	$\%R_2$	ω_d^*	R_1	R_2
n=5	L=1	$\pi = 1$	1.36	11.4	3.4	1.26	5.6	0.8
		$\pi = 5$	1.63	34.7	15.7	1.50	23.2	5.4
		$\pi = 15$	1.87	54.2	32.6	1.71	41.9	14.2
	L=5	$\pi = 1$	1.75	31.2	19.0	1.63	20.3	7.3
		$\pi = 5$	2.10	59.0	46.7	1.94	47.2	26.4
		$\pi = 15$	2.41	74.7	66.3	2.21	65.3	46.5
n=10	L=1	$\pi = 1$	1.16	3.9	0.9	1.12	1.8	0.2
		$\pi = 5$	1.26	14.2	4.5	1.21	8.7	1.4
		$\pi = 15$	1.33	26.4	10.2	1.28	18.2	3.7
	L=5	$\pi = 1$	1.35	15.1	7.3	1.31	9.3	2.6
		$\pi = 5$	1.47	34.8	21.4	1.42	25.4	9.7
		$\pi = 15$	1.56	51.2	36.5	1.50	40.8	19.6
n=20	L=1	$\pi = 1$	1.08	1.1	0.2	1.06	0.5	0.1
		$\pi = 5$	1.12	4.1	1.1	1.10	2.5	0.3
		$\pi = 15$	1.15	8.2	2.4	1.13	5.4	0.9
	L=5	$\pi = 1$	1.17	5.3	2.2	1.16	3.2	0.8
		$\pi = 5$	1.22	13.2	6.5	1.20	9.2	2.7
		$\pi = 15$	1.25	21.6	11.9	1.23	16.0	5.5

Table 3: The effect of biasing on the cost for (Q, r) model. $R_{1(2)}$: % reduction in controllable (total) cost.

cost and long lead time. In such cases, the amount of biasing one must use is the greatest. Our further calculations revealed that there is no clear effect of Q, although larger Q may seem to decrease the effect of biasing from the table.

b) Service objective: We let $Q = 20$, $\pi = 1$, $\mu = 4$, $\sigma = 2$, and the other parameters remain as in part (a). The objective is to achieve a long run average service level α during the lead time. Table 4 shows the expected total costs and service levels for both the traditional method and the biasing approach.

The costs for the traditional method and the biasing approach were calculated by the formula given in Lemma 1 by letting ω equal 1 and ω_c^*, respectively. The long run average service level of the traditional method was calculated using (15) for $\omega = 1$. Note that in the traditional method, the actual average service levels in the long run are consistently lower than expected. In other words, an inventory manager using the traditional method will deliver a service level below, and for small sample size, significantly below the level being targeted. For example, with a sample size of $n = 5$, a manager planning a service level of 90% will find that the long run service actually achieved is more than 5% lower than this.

		(Q,r) model with M=0.80, μ=4,σ=2			
		Traditional Method		*Biasing Alternative*	
n	α	*Service*	*Cost*	ω_c^*	Cost
5	0.80	0.757	1.601	1.225	1.608
	0.90	0.847	1.671	1.311	1.865
	0.95	0.896	1.844	1.420	2.329
	0.99	0.950	2.322	1.764	3.883
20	0.80	0.789	1.448	1.048	1.448
	0.90	0.887	1.552	1.062	1.591
	0.95	0.938	1.766	1.077	1.860
	0.99	0.982	2.330	1.119	2.587

Table 4: The effect of biasing on the service for (Q,r) model.

6 Conclusions and Remarks

When the demand parameters are estimated in inventory control problems, there is an increase in the expected operating cost or a degradation in expected service level of the system due to estimation errors. It is possible, however, to adjust the traditional estimators of sample mean and standard deviation by introducing biases which result in a lower operating cost.

This work has generalized the idea of biasing for the scale-location family of distributions and has shown that the newsboy problem, base-stock model and (Q, r) model with fixed Q are identical in terms of applying the biasing approach. We have attempted to bring together and extend a variety of papers dealing with the introduction of bias to compensate for estimation errors in inventory control. There seems to have been some 'reinventing the wheel' happening on this topic, with the early paper of Hayes being missed by later authors who have ended up in some cases delivering less than he had already done. Because of the common practice of setting customer service levels for inventory, the paper of Ritchken and Sankar [11] with the extension given here is particularly noteworthy, because, for small samples, major errors in the actual service levels delivered from those targeted can occur.

Some other interesting observations we made are: In all three models with either normal or gamma demand, the amount of savings due to biasing is a multiple of the standard deviation of demand and thus increases as the standard deviation increases. That is, biasing gives more savings for problems with large variations in demand. In particular, for a base-stock model, if the effectiveness of biasing is based on percent reduction in cost rather than the amount of savings, one can calculate the *exact* effect of biasing. This implies that if biasing is effective for a problem of certain cost parameters, it will remain equally effective regardless of how demand parameters change.

Since our analysis assumes stationary demand, the results here should be used with caution in problems where the demand shows seasonal patterns or a steady change. However, such problems are generally difficult to handle by other methods, too. A practical approach can be to apply the results developed here myopically, especially when simple estimates such as sample mean and variance from a sample of the most recent observations are used to forecast demand and the past observations are not effective forecasters. That is, biases can be calculated based on a recent sample and inventory policies can be modified accordingly by assuming stationary demand. Since sudden dramatic changes in demand parameters is not very typical even in problems with non-stationary demand, if the bias calculation is updated frequently, this method can provide a good approximate solution to such problems, as an alternative.

A Appendix

Lemma 4 *Let Z have a standard Normal distribution. Then, for any real numbers a and b,*

$$E_Z\left[\Phi(aZ+b)\right] = \Phi(\frac{b}{\sqrt{1+a^2}}).$$

Proof 4: Grundy *et al* [4]. $\square$

Lemma 5 *Let Y have a Chi square distribution with ν degrees of freedom and m be a real number such that $r = 2m + \nu > 0$ is an integer. Then, for any real number b we have*

$$E_Y\left[Y^m \Phi(b\sqrt{Y})\right] = 2^m \frac{\Gamma(r/2)}{\Gamma(\nu/2)} T_r(b\sqrt{r}) \ ,$$

where T_r is the cdf of Student's t distribution with r degrees of freedom.

Proof 5: (This lemma is given without any proof in Weerahandi [14] and m is unnecessarily restricted to be an integer)

$$E_Y[Y^m \Phi(b\sqrt{Y})] = \int_0^\infty \int_{-\infty}^{b\sqrt{y}} y^m \phi(z)\chi(y)\, dz\, dy$$

where $\chi(y)$ is the pdf of Y. Note that $t \equiv z\sqrt{\nu/y}$ is a Student's t random variable with ν degrees of freedom. Then,

$$E_Y[Y^m \Phi(b\sqrt{Y})] = \int_0^\infty \int_{-\infty}^{b\sqrt{\nu}} y^m \phi(t\sqrt{y/\nu})\sqrt{y/\nu}\,\chi(y)\, dt\, dy.$$

Interchanging the integrals, and rearranging we have

$$E_Y[Y^m \Phi(b\sqrt{Y})] = \sqrt{\frac{1}{\nu}} \int_{-\infty}^{b\sqrt{\nu}} \int_0^\infty y^{m+1/2}\, \phi(t\sqrt{y/\nu})\, \chi(y)\, dy\, dt.$$

Now,

$$\begin{aligned}
\phi(t\sqrt{y/\nu})\chi(y)\, dt\, dy &= \frac{1}{\sqrt{2\pi}} e^{-\frac{yt^2}{2\nu}} \frac{1}{\Gamma(\nu/2)} (1/2)^{\nu/2} y^{\frac{\nu}{2}-1} e^{-y/2} \\
&= \frac{1}{\sqrt{2\pi}} \left(\frac{2\nu}{t^2+\nu}\right)^{\nu/2} \left(\frac{t^2+\nu}{2\nu}\right)^{\nu/2} \frac{(1/2)^{\nu/2}}{\Gamma(\nu/2)} y^{\frac{\nu}{2}-1} e^{-\left(\frac{t^2+\nu}{2\nu}\right)y} \\
&= \frac{1}{\sqrt{2\pi}} \left(\frac{\nu}{t^2+\nu}\right)^{\nu/2} \phi_{\nu/2}(y),
\end{aligned}$$

where $\phi_{\nu/2}(y)$ is the *pdf* of a Gamma random variable with shape parameter $\nu/2$ and scale parameter $2\nu/(t^2+\nu)$. Then,

$$\begin{aligned}
E_Y[Y^m \Phi(b\sqrt{Y})] &= \int_{-\infty}^{b\sqrt{\nu}} \int_0^\infty \sqrt{\frac{1}{2\pi\nu}}\, y^{m+1/2} \left(\frac{\nu}{t^2+\nu}\right)^{\nu/2} \Phi_{\nu/2}(y)\, dy\, dt \\
&= \int_{-\infty}^{b\sqrt{\nu}} \sqrt{\frac{1}{2\pi\nu}}\, E_Y[Y^{m+1/2}] \left(\frac{\nu}{t^2+\nu}\right)^{\nu/2} dt \\
&= \int_{-\infty}^{b\sqrt{\nu}} \sqrt{\frac{1}{2\pi\nu}} \left(\frac{2\nu}{t^2+\nu}\right)^{m+1/2} \frac{\Gamma(\frac{\nu+2m+1}{2})}{\Gamma(\nu/2)} \left(\frac{\nu}{t^2+\nu}\right)^{\nu/2} dt.
\end{aligned}$$

Letting $u = t\sqrt{\frac{\nu+2m}{\nu}}$ and noting that $\nu + 2m = r$, we get

$$\begin{aligned} E_Y[Y^m \Phi(b\sqrt{Y})] &= \frac{\Gamma(r/2)}{\Gamma(\nu/2)} 2^m \int_{-\infty}^{b\sqrt{r}} \frac{\Gamma(\frac{r+1}{2})}{\Gamma(r/2)} \frac{1}{\sqrt{r\pi}} \left(\frac{r}{u^2+r}\right)^{\frac{r+1}{2}} du \\ &= \frac{\Gamma(r/2)}{\Gamma(\nu/2)} 2^m T_r(b\sqrt{r}). \end{aligned}$$

Lemma 6 *If $X \sim Gamma(r_1, \gamma_1)$ and $Y \sim Gamma(r_2, \gamma_2)$ then for $b \geq 0$*

$$E_X\left[\Phi_{r_2}(bX|\gamma_2)\right] = \mathcal{B}_{r_2,r_1}(\frac{b\gamma_1}{b\gamma_1+\gamma_2})$$

where $\mathcal{B}_{r_2,r_1}$ is the cdf of a Beta distribution with parameters r_2 and r_1.

Proof 6: We generalize the proof in [14] for any two scale parameters γ_1 and γ_2.

$$\begin{aligned} E_X[\Phi_{r_2}(bX|\gamma_2)] &= E_X[\Phi_{r_2}(\frac{bX}{\gamma_2}|1)] \\ &= E_X[Prob\{Z \leq \frac{bX}{\gamma_2}\} \mid X] \end{aligned}$$

where $Z \sim Gamma(r_2, 1)$. Let $M = X/\gamma_1$. Then, $M \sim Gamma(r_1, 1)$. Thus,

$$\begin{aligned} E_X[\Phi_{r_2}(bX|\gamma_2)] &= E_M[Prob\{Z \leq \frac{b\gamma_1}{\gamma_2}M\} \mid M] \\ &= Prob\{\frac{Z}{M} \leq \frac{b\gamma_1}{\gamma_2}\} \\ &= Prob\{\frac{Z}{Z+M} \leq \frac{b\gamma_1}{b\gamma_1+\gamma_2}\} \\ &= \mathcal{B}_{r_2,r_1}\left(\frac{b\gamma_1}{b\gamma_1+\gamma_2}\right). \end{aligned}$$

The last equality follows from the relationship between the Gamma and the Beta distributions, which is a known result from distribution theory.

Proof of Lemma 1: We first find

$$E_{\bar{X},s}[E_X[y^*(\bar{X}, \omega s) - X]^+].$$

The inner expectation is actually conditional on $(\bar{X}, s)$, with X being independent of both $\bar{X}$ and s. Replacing $y^*(\bar{X}, \omega s)$ by $\bar{X} + k\omega s$, we have

$$E_{\bar{X},s}[E_X[\bar{X} + k\omega s - X]^+ \mid \bar{X}, s] = E_{\bar{X},s,X}[\bar{X} + k\omega s - X]^+.$$

Let $u = (n-1)s^2/\sigma^2$ and $v = \bar{x} - x$. Note that u and v are independent, u having a Chi-square distribution with $n-1$ degrees of freedom and v having a Normal distribution with mean 0 and variance $(1+1/n)\sigma^2$. Then, the expectation is

$$E_u[E_v[v + \frac{k\omega\sigma}{\sqrt{n-1}}\sqrt{u}]^+ \mid u].$$

The partial expectation inside is messy but straight forward to calculate and equals the following.

$$E_v[v + \frac{k\omega\sigma}{\sqrt{n-1}}\sqrt{u}]^+ = \frac{\sqrt{n+1}\sigma}{\sqrt{2\pi n}} e^{-\frac{\omega^2 k^2 un}{2(n-1)(n+1)}} + \frac{\omega k\sigma}{\sqrt{n-1}}\sqrt{u}\Phi\left(\sqrt{\frac{n}{n^2-1}}k\omega\sqrt{u}|0,1\right).$$

Now, the outer expectation with respect to u can be calculated by noting that

$$\begin{aligned} E_u[e^{au}] &= (1-2a)^{-\frac{n-1}{2}}, \ a < 1/2, \\ E_u[u^{1/2}] &= \frac{\sqrt{2}\Gamma(n/2)}{\Gamma(\frac{n-1}{2})}. \end{aligned}$$

Using the above equations and Lemma 5 for $m = 1/2$, the outer expectation is

$$\left[\sqrt{\frac{n+1}{2\pi n}}\left(1 + \frac{nk^2\omega^2}{n^2-1}\right)^{-\frac{n-1}{2}} + \sqrt{\frac{2}{n-1}}\frac{\Gamma(\frac{n}{2})}{\Gamma(\frac{n-1}{2})}k\omega T_n\left(\frac{nk\omega}{\sqrt{n^2-1}}\right)\right]\sigma.$$

Given the above, the expectation of $F(y^*(\bar{X}, \omega s); \mu, \sigma)$ can easily be found.

Proof of Lemma 2: (taken from Ritchken and Sankar [11]) Equation 11 is simply

$$\begin{aligned} \alpha &= E_{\bar{X},s}[\Phi(\frac{l\omega s + \bar{X} - \mu}{\sigma}) \mid \bar{X}, s] \\ &= Prob\{Z \le \frac{l\omega s + \bar{X} - \mu}{\sigma}\} \\ &= Prob\{\frac{\sigma Z + \mu - \bar{X}}{s} \le l\omega\} \end{aligned}$$

where Z is a standard Normal variable. It can easily be verified that $t \equiv \frac{\sigma Z + \mu - \bar{X}}{s\sqrt{1+1/n}}$ is a Student's t random variable with $n-1$ degrees of freedom. Thus,

$$\alpha = Prob\{t \le \frac{l\omega}{\sqrt{1+1/n}}\} = T_{n-1}\left(\frac{l\omega}{\sqrt{1+1/n}}\right).$$

Noting that l is $\Phi^{-1}(\alpha)$ the result follows.

Proof of Lemma 3: Note the following expectation:

$$\begin{aligned} E_{\hat{\gamma}}[E_X[k\omega\hat{\gamma} - X]^+] &= \int_{\hat{\gamma}}\int_{X \le k\omega\hat{\gamma}} (k\omega\hat{\gamma} - x)\, \phi_r(x|\gamma)\, \phi_{nr}(\hat{\gamma}|\frac{\gamma}{nr})\, dx\, d\hat{\gamma} \\ &= k\omega\int_{\hat{\gamma}}\int_{X \le k\omega\hat{\gamma}} \hat{\gamma}\, \phi_r(x|\gamma)\, \phi_{nr}(\hat{\gamma}|\frac{\gamma}{nr})\, dx\, d\hat{\gamma} \\ &\quad - \int_{\hat{\gamma}}\int_{X \le k\omega\hat{\gamma}} x\, \phi_r(x|\gamma)\, \phi_{nr}(\hat{\gamma}|\frac{\gamma}{nr})\, dx\, d\hat{\gamma}. \end{aligned}$$

Using the fact $u\phi_i(u|j) = ij\phi_{i+1}(u|j)$ in both terms (for $\hat{\gamma}$ in the first term and for x in the second) and simplifying, this expectation is equal to

$$\begin{aligned} &= k\omega\int_{\hat{\gamma}}\int_{X \le k\omega\hat{\gamma}} \gamma\, \phi_r(x|\gamma)\, \phi_{nr+1}(\hat{\gamma}|\tfrac{\gamma}{nr})\, dx\, d\hat{\gamma} - \int_{\hat{\gamma}}\int_{X \le k\omega\hat{\gamma}} r\gamma\, \phi_{r+1}(x|\gamma)\, \phi_{nr}(\hat{\gamma}|\tfrac{\gamma}{nr})\, dx\, d\hat{\gamma} \\ &= k\omega\gamma E_{\hat{\gamma}_1}[\Phi_r(k\omega\hat{\gamma}_1|\gamma)] - r\gamma\, E_{\hat{\gamma}}[\Phi_{r+1}(k\omega\hat{\gamma}|\gamma)] \end{aligned}$$

where $\hat{\gamma}_1 \sim Gamma(nr+1, \frac{\gamma}{nr})$ and $\hat{\gamma} \sim Gamma(nr, \frac{\gamma}{nr})$. Then, from Lemma 6 we have

$$E_{\hat{\gamma}}[E_X[k\omega\hat{\gamma} - X]^+] = k\omega\gamma\mathcal{B}_{r,nr+1}\left(\frac{k\omega}{k\omega+nr}\right) - r\gamma\mathcal{B}_{r+1,nr}\left(\frac{k\omega}{k\omega+nr}\right).$$

Given the above, the expectation of $F(y^*(\omega\hat{\gamma});\gamma)$ can easily be found.

References

[1] Biggs R.J., Campion W.M. "The Effect and Cost of Forecast Error Bias for Multiple-Stage Production-Inventory Systems.", *Decision Sciences*, **13**, (1982), 570-584.

[2] Bookbinder J.H., A.E. Lordahl. "Estimation of Inventory Re-ordel Levels Using The Bootstrap Statistical Procedure.", *IIE Transactions*, **21**, 4, (1988).

[3] Ehrhardt R. "The Power Approximation for Computing (S,s) Inventory Policies.", *Management Science*, **25**, (1979), 777-786.

[4] Grundy P.M., Healy M.J.R. and Rees, D.H. "Economic Choice of the Amount of Experimentation.", *J.R. Statistics Soc.*, B. 18, (1956), 32-55.

[5] Hadley G. and Whitin T.M., *Analysis of Inventory Systems*, Prentice-Hall, Inc., N.J. (1963).

[6] Hayes, R.H. "Statistical Estimation Problems in Inventory Control.", *Management Science*, **15**, 11, (1969), 686-701.

[7] Hayes R.H., "Statistical Estimation Models for Optimal Inventory Control.", Ph.D. thesis, Operations Research Program, Stanford University, February (1966).

[8] Jacobs A.R. and Wagner H.M. "Reducing Inventory System Costs by Using Robust Demand Estimators.", *Management Science*, **35**, (1989), 771-787.

[9] Jacobs A.R. and Wagner H.M. "Reducing Inventory System Costs by Using Regression-Derived Estimators of Demand Variability.", *Decision Sciences*, **20**, (1989), 559-574.

[10] Lordahl A.E., J.H. Bookbinder. "Order Statistics Calculation, Costs and Service in an (s,Q) Inventory System.", *Naval Research Logistics*, **41**, (1994), 81-97.

[11] Ritchken H.P. and Sankar R., "The Effect of Estimation Risk in Establishing Safety Stock Levels in an Inventory Model.", *JORS*, **35**, 12, (1984), 1091-1099.

[12] Silver E.A. and Rahmana M.R. "The Cost Effects of Statistical Sampling in Selecting the Reorder Point in a Common Inventory Model.", *JORS*, **37**, 7, (1986), 707-713.

[13] Silver E.A. and Rahmana M.R., 'Biased Selection of the Inventory Reorder Point When Demand Parameters Are Statistically Estimated', *Eng'g. Costs and Prod. Economics*, 12, 283-292, (1987).

[14] Weerahandi S., 'Decision Theoretic Estimation of Optimum Inventory Levels with Natural Invariant Loss Functions', *Statistics and Decisions*, 5, 177-188, (1987).