A Simple and Robust Batch-Ordering Inventory Policy Under Incomplete Demand Knowledge

Liwei Bai*

Christos Alexopoulos†

Mark E. Ferguson‡

Kwok-Leung Tsui§

Georgia Institute of Technology

December 2005

Abstract

Generally, the derivation of an inventory policy requires the knowledge of the underlying demand distribution. Unfortunately, in many settings such as retail, demand is not completely observable in a direct way or inventory records may be inaccurate. A variety of factors, including the potential inaccuracy of inventory records, motivate retailers to seek replenishment policies with a fixed order quantity. We derive estimators of the first two moments of the (periodic) demand by means of renewal theoretical concepts. We then propose a regression-based approximation to improve the quality of the estimators. These estimators are used in conjunction with the Power Approximation (PA) method of Ehrhardt and Mosier (1984) to obtain an (r, Q) replenishment policy. The proposed methodology is robust and easy to code into a spreadsheet application. A series of numerical studies are carried out to evaluate the accuracy and precision of the estimators, and to investigate the impact of the estimation on the optimality of the inventory policies. Our experiments indicate that the proposed (r, Q) policy is very close, with regard to the mean total cost per period, to the (s, S) policy obtained via the PA method when the demand process is fully observable.

Keywords: Inventory Inaccuracies, Power Approximation Method

* Liwei Bai, Carlson Capital, L.P., Dallas, TX 75201, Phone: (214) 932-9648, e-mail: dbai@carlsoncapital.com.

† Christos Alexopoulos, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332. Phone: (404) 894-2361, e-mail: christos@isye.gatech.edu.

‡ Mark E. Ferguson, College of Management, Georgia Institute of Technology, Atlanta, GA 30332. Phone: (404) 894-4330, e-mail: mark.ferguson@mgt.gatech.edu.

§ Kwok-Leung Tsui, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332. Phone: (404) 894-2311, e-mail: ktsui@isye.gatech.edu.

1 Introduction

While modern-day inventory theory has been under development for more than 50 years, the application of the theoretical results from this effort to real-world applications remains a challenge. In particular, many common theoretically optimal inventory policies fail to deliver their promised results in practice because the assumptions in the underlying models often do not hold. One key underlying assumption that is often violated in practice is that the demand distribution and the associated parameter values are known with certainty by the policy maker. In reality, very few firms know the true underlying demand distributions for their products. Instead, firms assume a specific demand distribution and replace the unknown parameters with estimates computed from historical data. Unfortunately, the data used to estimate these parameters is often limited due to frequent product introductions and changing customer preferences. Even in the rare occasions where the customer demand distribution is known with certainty, inventory control is still problematic due to frequent inaccuracies in the system-reported inventory positions.

Recent studies have shown that system-reported inventory positions are often inaccurate. According to DeHoratius and Raman (2004), records were inaccurate for 65% of the SKUs at a publicly traded retailer. Kang and Gershwin (2004) found that cycle counts at the SKU level across the stores of a global retailer matched the system-reported inventory count in only 51% of the cases. The best performing store in their study only achieved a 70-75% accuracy level. Kok and Shang (2005) provide an example of a heavy equipment manufacturing firm where audits uncovered inventory inaccuracies of 1.6% of total inventory value at their distribution centers. Problems arise when firms with system-reported inventory inaccuracies continue to manage their stock with policies developed assuming perfect information about inventory positions and demand distributions. Inventory inaccuracies also affect the demand distribution estimation process, as most common inventory replenishment polices require the estimation (based on historical demand as recorded by the information system) of the first and second moments of the demand distribution. Kok and Shang (2005) estimate the cost penalties at the heavy equipment manufacturer for the misuse of such “perfect-information” replensihment policies at 5%.

Why are system-reported inventory levels so inaccurate? Possible reasons appearing in the literature include misplacement, shrinkage, spoilage, and transaction errors. Misplacement involves the physical stocking of an item in a location where the customer or employee can not locate it. Non-perishable items may be found and replaced during a periodic cycle count and system inventory records can be adjusted, but perishable products often expire before they are noticed or located. Shrinkage typically refers to the theft of items by customers or the firm's employees. Spoilage occurs when a perishable item is not sold before its expiration and transaction errors refer to mistakes in the process of receiving or selling an item. A common type of transaction error in the retail industry involves the mis-scanning of items at the point-of-sale terminal.

The net impact of these reasons on system record inventory counts is almost always an over-reporting of actual inventory levels. While misplacements and transaction errors mostly affect the system record counts and not the physical inventory counts, shrinkage and spoilage directly decrease the physical count, resulting in a system wide over-reporting of physical accounts. This net-loss in actual inventory levels is supported by industry-wide empirical studies: A study of the 200 largest European retailers of consumer goods reports an average stock loss of 1.75% of annual sales (ECR Europe 2003) while a similar study of 118 U.S. retailers reports an average loss of 1.7% (Hollinger, 2003). New tracking technologies such as Radio Frequency Identification (RFID) offer promise of decreasing these numbers, but are currently cost prohibitive for most consumer items. Better security and more frequent cycle counts also help but are also often cost-prohibitive, especially in the consumer goods retail sector. Thus, more robust replenishment policies are needed that account for realities such as inaccurate system-reported inventory levels.

The motivation for our problem comes from a small to mid-sized retailer facing stochastic demand and lead-times, linear holding and backordering cost, fixed order cost, and is restricted to placing orders at the start of each period. If the demands during successive periods are independent and identically distributed (i.i.d.) random variables from a known distribution, the lead-times form an i.i.d. random sequence, and the reported inventory levels are accurate, an (s, S) policy derived from these parameters minimizes the long-run expected undiscounted total cost per period. If only the first two moments of the demand distribution are known, an approximately optimal (s, S) policy can be obtained, e.g., by the Power Approximation (PA) method of Ehrhardt and Mosier (1984). Because of the reasons described above however, recorded inventory is often inaccurate and there is uncertainty about the true demand distribution — even the first two moments of the distribution. Detailed audits of every SKU every period are unrealistic, so the retailer is forced to place orders each period with imperfect information about its inventory position. In addition to an inventory policy that acknowledges the reality of system-reported inventory position inaccuracies, the retailer also desires a policy that reflects other common constraints of the industry such as a stable order quantity and a simple robust policy that does not require dynamic (re-)optimization and can be programmed in a spreadsheet or database software application.

The main objective of this paper is to provide an (r, Q) inventory policy that minimizes the long-run mean (undiscounted) total cost per period subject to the aforementioned constraints. Our choice of an (r, Q) policy over an (s, S) policy comes from the retailer's desire for a stable order quantity in each period. To this end, we first construct estimators for the first two moments of the demand distribution using well-known results from renewal theory and realizations of periods between orders. We proceed with the derivation of approximately optimal inventory policies similar in spirit to the simple and robust PA method first described in Ehrhardt (1979) and revised in Ehrhardt and Mosier (1984). The PA method has been used successfully in a variety of settings and owes its popularity to its simplicity and the surprisingly good fit of the regression model to a variety of demand distributions. We proceed with the development of an (r, Q) policy with $r = s$ that approximates the near-optimal (s, S) policy calculated using the PA method. The performance of the (r, Q) policy is evaluated by means of an experimental grid consisting of 216 design points and is similar to Ehrhardt's experimental design. Since the cost of the (r, Q) policy is sensitive to the variance of the demand per period, we develop an alternative regression-based estimator of the variance of the periodic demand using the same grid. The incorporation of this estimator results in a substantially improved (r, Q) policy with regard to the mean total cost per period. Indeed, a simulation study based on our experimental grid reveals that the relative difference between the average costs per period induced by the proposed PA-based (r, Q) policy and the (s, S) policy under fully observable demands is within $\pm 5\%$ for 94% of the 216 design points.

The rest of the paper is organized as follows: In the next section, we position our research in the context of the relevant literature. §3 contains the key assumptions and notation, and reviews the PA method. In §4, we discuss the estimation of the first two moments of the demand per period and the conversion of (s, S) policies obtained by the PA method to (r, Q) policies. In §5, we evaluate the quality of the moment estimators in §4 and the performance of the (r, Q) policy by means of an extensive experimental study. In §6, we propose a refined estimator for the variance of the demand per period, and use the experimental setting from §5 to show that this estimator yields an (r, Q) policy that is close to the (s, S) policy obtained from the PA method under fully observable demand with regard to the mean total cost per period. §7 contains conclusions and a discussion of managerial implications. The Appendix, available online at <ftp://ftp.isye.gatech.edu/pub/christos/Bai_etal_2005_Appendix.pdf>, contains tables with experimental results.

2 Review of Existing Inventory Inaccuracy Compensation Methods

In a world with perfect information, no system-reported inaccuracies, and no restrictions on the periodic order quantity, an (s, S) policy minimizes the long-run expected total cost per period. Scarf (1960) and Iglehart (1963a) prove the optimality of (s, S) policies, and Veinott and Wagner (1965) consider methods for computing them when demands per period are i.i.d. and normally distributed random variables. Even under these ideal conditions, (s, S) policies are difficult to compute. Thus, several approximations including Ehrhardt's (1979) PA method were developed. Porteus (1985) offers a review of some of the most common approximations. *Our scenario assumes no prior knowledge of the demand distribution type or its moments, and that the total demand in a period is not directly observable*. These assumptions in addition to our pre-stated requirement of a simple robust policy that does not require dynamic (re-)optimization led us to choose the PA method as the basis for our work.

When the demand is observable but its distribution is unknown, a commonly used approach is to fit a statistical model to the observed demand data and use the fitted model to obtain the desired policy. Typically, limited historical data is available to estimate the parameters of the demand distribution. Jacobs and Wagner (1989) investigate how the choices of statistical estimators affect the mean total cost. Specifically, they show that when demand variability is large, exponentially smoothed estimators can substantially outperform the sample mean and sample variance; hence they yield improved (s, S) policies.

If the demand in a period is not observable directly, the problem becomes a lot more challenging. There is little previous research on how a firm should control its inventory under unobservable demand when the demand distribution is also unknown. This situation happens quite frequently in practice, however, due to the system-reported inventory inaccuracies that are common at most retailers. Based on their work at the MIT Auto-ID Center, Kang and Gershwin (2004) discuss five methods that retailers can use in practice to compensate for their inventory inaccuracies: (I) iteratively increasing the safety stock of an inventory policy with a reorder point, (II) cycle counts, (III) manual reset of the inventory record, (IV) periodic negative adjustments of the inventory position record, and (V) using tracking technology such as RFID.

The first method involves a firm incrementally increasing its safety stock (calculated using a policy that assumes perfect inventory information) until acceptable customer service levels are achieved. The second method relies on periodic cycle counts to synchronize the system-recorded inventory level with the physical count. The third method involves monitoring the order frequency of an item and resetting its stock to zero if an unusual number of periods without an order are observed. The fourth method is essentially an attempt to model and subtract the other forms of demand other than demand from purchases. The fifth method simply depends on better inventory tracking.

Of the five methods, the first two are the most common but have major limitations. Method I involves trial and error and often results in excessive safety stocks. A firm using Method II typically ignores the inventory inaccuracy problem during its order replenishment decisions but periodically corrects the inaccuracy through a physical count. A common problem with this method involves determining the frequency of the expensive physical counts. Kok and Shang (2005) include the cost of cycle counts in a joint optimization problem of how much to replenish and how often to count. For a finite horizon with no setup cost, they show that an inspection-adjusted base-stock policy is near-optimal. For the retailer providing our motivation, physical scheduling constraints require that manual cycle counts only occur at certain intervals and a significant amount of time elapses between these intervals. Thus, we do not include the time between cycle counts as a decision variable. In fact, our policy may be used by a firm that never corrects its reported inventory level through cycle counts, making it applicable to situations where cycle counting is impractical.

Method III is becoming more common in the retail industry as it only involves the addition of a simple logic code into the firm's inventory management system. It is only useful if high service levels are not required, as a number of periods may go by before the system's logic resets the inventory level of an item to zero and issues a replenishment request. A second drawback of this method is that it may actually exacerbate the system recorded inventory accuracy problem. A tradeoff involved with setting the threshold for the number of periods with zero demand before the system resets the inventory level of an item equal to zero is that a threshold set too low results in item inventory levels being set equal to zero when units are actually in stock and the number of zero-demand periods actually resulted from a true sequence of low demands. In such cases, the use of this method adds additional system-recorded inventory inaccuracies.

Methods IV and V are discussed in Atali et al. (2005), who measure the value of RFID in the context of a periodic review inventory policy with no fixed ordering cost, multiple replenishments that occur between cycle counts, and imperfect inventory information. They explicitly model four demand streams that affect inventory accuracy: customer demand, misplacements, shrinkage, and transaction errors. They then develop policies based on finite-horizon dynamic programs, which account for the inventory inaccuracy, but are non-stationary in time and vary from period to period based on the time remaining until the next cycle count. In contrast, our policy is a stationary, infinite-horizon policy that accounts for a fixed order cost and does not require an explicit knowledge of the distributions for the multiple drivers of demand. When the fixed ordering cost is sufficiently large, our policy resembles the traditional PA method, which we review in the next section.

3 Review of the Power Approximation Method

We set the length of each period to one day. Replenishment costs consist of a setup cost K and a unit cost c. At the end of each day a cost of h or b is incurred for each unit on hand or backlogged, respectively. We assume that the demands $X_1, X_2, \ldots$ for a specific item during different days are i.i.d. from an unspecified distribution with mean μ and variance $\sigma^2 < \infty$. Lead-times are expressed in days. Let $\mathrm{E}(L)$ and $\mathrm{Var}(L)$ denote the mean and variance of lead-time. The objective is the minimization of the long-run expected undiscounted total cost per day. The use of undiscounted cost is common for retail applications where lead-times are typically expressed in days.

Under the assumptions above, an (s, S) policy is optimal. That is, if during the periodic inspection the inventory position (inventory on hand plus on order), say x, is less than s, an order of $S-x$ units is placed. The computation of the reorder point s and the order-up-to value S requires the complete specification of the demand distribution, and is difficult to carry out in practice. For fixed lead-times and large K and b, Roberts (1962) derives approximations that are easy to compute but still require full knowledge of the demand distribution. Unfortunately, the demand distribution is rarely known in full and lead-times are frequently random; in fact, managers are fortunate if they know the first two moments of the distributions for these random variables. To address these issues, Ehrhardt (1979) proposed the PA method. For fixed lead-times, the PA method assumes the mean total cost per day, T, can be modeled as

$$T/h = cf_1(L, \theta) f_2(\mu, \theta) f_3(\pi, \theta) f_4(\kappa, \theta), \tag{1}$$

where $\theta = \sigma^2/\mu$, $\kappa = K/h$, $\pi = b/h$, $f_i(x, \theta) = x^{\gamma_i(x,\theta)} \exp[\delta_i(x, \theta)]$ $(i = 1, \ldots, 4)$, and

$\gamma_i(x, \theta)$, $\delta_i(x, \theta)$ are linear combinations of variables from the set

$$\{1, x, 1/x, \theta, 1/\theta, x^2, 1/x^2, \theta^2, 1/\theta^2, x\theta, x/\theta, \theta/x, 1/(x\theta)\}.$$

Model (1) is fitted via regression based on optimal values obtained for a grid of 288 test cases using Poisson or negative binomial demand distributions (with variance-to-mean ratios, θ, equal to 3 or 9), three lead-times (0, 2 and 4 days), two values for K (32 and 64), and four values for π (4, 4, 24, and 99). In addition to the parameters s and S, the PA method provides easy-to-compute formulas for the mean holding cost per day, the mean replenishment cost per day, the mean backlog cost per day, and the long-run backlog protection (defined as the probability that a stockout does not occur during a day); see Ehrhardt (1985) for details.

Ehrhardt (1984) shows that for random lead-times, the PA method remains valid when μ is replaced by $\mu_L = [\mathrm{E}(L)+1]\mu$, the mean demand during the time needed to replenish the inventory, and σ^2 is replaced by $\sigma_L^2 = [\mathrm{E}(L)+1]\sigma^2+\mu^2\mathrm{Var}(L)$, the variance of the demand during a lead-time. The policy is determined as follows (the subscript "p" stands for "power"). First calculate

$$D_p = 1.30\mu^{0.494}\kappa^{0.506}(1+\sigma_L^2/\mu^2)^{0.116},$$

$$z = \sqrt{D_p/(\sigma_L\pi)}, \quad \text{and} \tag{2}$$

$$s_p = 0.973\mu_L + \sigma_L(0.183/z + 1.063 - 2.192z).$$

If $D_p/\mu > 1.5$, let $s = s_p$ and $S = s_p + D_p$. Otherwise, set $s = \min\{s_p, S_0\}$ and $S = \min\{s_p + D_p, S_0\}$, where $S_0 = \mu_L + \Phi^{-1}(b/(b+h))$ and $\Phi^{-1}(\cdot)$ is the inverse c.d.f. of the standard normal distribution. With discrete values, s_p, D_p and S_0 are rounded to the nearest integer. If any of the moments μ, σ^2, $\mathrm{E}(L)$ or $\mathrm{Var}(L)$ are not known, they can be replaced by statistical estimates.

In the next section, we address the estimation of the first two moments of the daily total demand (purchase, shrinkage, etc.) by means of a well-known renewal theoretic approach. We then use these parameter value estimates to derive a modified (r, Q) policy and compare it against an optimal (r, Q) policy found via simulations over a grid around $(s, S-s)$.

4 Demand Estimation and Derivation of $(r,\ Q)$ Policies

As described in the previous section, the PA method can use estimates of the mean and variance of the daily demand and lead-time distributions. We now derive good estimations for the parameters of these distributions based on imperfect demand and inventory information. Our estimation procedure is motivated by the retailer whose current inventory system uses a non-optimal "two-bin" periodic review (r, Q) inventory policy, where Q is a non-optimal quantity determined by the retailer's supplier's current case-pack size and $r = Q/2$. The policy works as follows. The inventory status is checked every T days. [1] If the inventory level is less than r, a quantity Q is ordered. Let t_i denote the time that the ith order is placed, which is recorded in a database. When an order of size Q arrives, its arrival time R_i is also recorded in the database. Only these three pieces (t_i, R_i, Q) of historical information are available.

Recall that the mean and variance of the lead-times are also an input to the PA method. Since the historical lead-times can be obtained from $R_i - t_i$, one can easily obtain their sample mean and sample variance. Let $\tau_i = t_i - t_{i-1}$, denote the days between two successive orders (DBO) with $t_0 = 0$. As a random variable, DBO carries some information that we use towards the estimation of the demand moments. Without loss of generality, we assume that system inventory records are reviewed at the start of each day ($T = 1$). The next subsection describes our approach.

4.1 Estimation of Mean and Variance under Imperfect Demand Information

In this section, we provide a method for estimating the mean and variance of the daily demand distribution. We define demand in this case as total item usage, equal to the sum of customer sales, spoilage, shrinkage, permanent misplacements, items sold but not recorded, and any other unknown effects that negatively impact the item's physical count.

Based on our inventory control policy, once the approximate inventory position falls below the

[1] Recall that inventory status is based on system reported inventory levels; actual inventory levels may be lower due to the unobservable demands described in the Introduction. Occasional cycle counts will align these two measures, but it is not required that the cycle counts occur at the time of the inventory level review.

reorder point r, then a order of size Q will be placed. To obtain the estimators for the first two moments of the daily demand, we assume temporarily that delivery is instantaneous and that the total demand during the time interval $[t_{i-1}, t_i)$ elapsed between orders exceeds the order quantity Q. This allows us to count the cycles by means of a renewal process. Further, we use renewal-type approximations for large Q to derive the estimators for μ and σ^2.

Let

$$N(u) \equiv \sup\{n : Z_n = X_1 + X_2 + \cdots + X_n \leq u\}, \quad u \geq 0.$$

Then $\{N(u) : u \geq 0\}$ is a renewal process and the time interval between order replenishments has length

$$\tau = N(Q) + 1 = \inf\{n : Z_n > Q\}.$$

From Heyman and Sobel (1982, Section 5-3) we have

$$\mathrm{E}[N(Q)] = \frac{Q}{\mu} + \frac{\sigma^2 - \mu^2}{2\mu^2} + o(1) \quad \text{as } Q \to \infty, \tag{3}$$

where the notation $o(1)$ denotes a function $h(z)$ such that $h(z)/z \to 0$ as $z \to \infty$. Under the additional assumption $\mathrm{E}(X_i^4) < \infty$, one also has

$$\mathrm{Var}[N(Q)] = \frac{\sigma^2}{\mu^3} Q + c_v + o(1) \quad \text{as } Q \to \infty, \tag{4}$$

where

$$c_v = \frac{2\sigma^2}{\mu^2} + \frac{3}{4} + \frac{5}{4}\frac{\sigma^4}{\mu^4} - \frac{2}{3}\frac{\mathrm{E}(X_1^3)}{\mu^3}.$$

Since $\tau = N(Q) + 1$, we have

$$\mathrm{E}(\tau) = \frac{Q}{\mu} + \frac{\sigma^2 + \mu^2}{2\mu^2} + o(1) \quad \text{as } Q \to \infty$$

and $\mathrm{Var}(\tau) = \mathrm{Var}[N(Q)]$.

Remark 1 In addition to these approximations for the first two moments of τ, one has the following central limit theorem:

$$\frac{\tau - Q/\mu - 1}{\sqrt{\sigma^2 Q/\mu^3}} \xrightarrow{d} N(0, 1) \quad \text{as } Q \to \infty,$$

where $N(\alpha, \beta^2)$ denotes a normal random variable with mean α and variance β^2 and "$\xrightarrow{d}$" denotes convergence in distribution. Hence, we can use the approximation

$$\tau \approx N\left(\frac{Q}{\mu} + 1, \frac{Q\sigma^2}{\mu^3}\right). \quad \triangleleft$$

The following theorem proposes consistent estimators for μ and σ^2.

Theorem 1 *Assume that (a) the daily demands are i.i.d. random variables with mean μ and variance σ^2, (b) deliveries are instantaneous, and (c) the total demand during each time interval between orders exceeds the order quantity Q. Let τ_i be the length of the ith time interval between orders, and let $\bar{\tau}(n)$ and $S_\tau^2(n)$ denote the sample mean and sample variance, respectively, of the τ_i. Consider the following estimators for μ and σ^2:*

$$\hat{\mu} = Q/\bar{\tau}(n) \tag{5}$$

and

$$\widehat{\sigma^2} = \frac{S_\tau^2(n)Q^2}{\bar{\tau}^3(n)}. \tag{6}$$

Then, as $Q \to \infty$ and $n \to \infty$,

$$\hat{\mu} \xrightarrow{p} \mu \quad \textit{and} \quad \widehat{\sigma^2} \xrightarrow{p} \sigma^2, \tag{7}$$

where "$\xrightarrow{p}$" denotes convergence in probability.

Proof Let $M(\cdot) \equiv \mathrm{E}[N(\cdot)]$ be the renewal function. For fixed Q, the (weak) law of large numbers implies

$$\bar{\tau}(n) \xrightarrow{p} M(Q) + 1 \quad \text{as } n \to \infty,$$

hence

$$\bar{\tau}(n)/Q \xrightarrow{p} [M(Q) + 1]/Q \quad \text{as } n \to \infty.$$

Since $M(Q)/Q \to 1/\mu$, as $Q \to \infty$, by the elementary renewal theorem (Heyman and Sobel, 1982), one has

$$\bar{\tau}(n)/Q \xrightarrow{p} 1/\mu \quad \text{as } n \to \infty \text{ and } Q \to \infty. \tag{8}$$

Equation (5) follows from (8) and the continuous mapping theorem (Lehmann, 1990).

To prove (6), write

$$\widehat{\sigma^2} = \frac{S_\tau^2(n)}{\mathrm{Var}(\tau)} \frac{Q^3}{\bar{\tau}^3(n)\mu^3} \frac{\mathrm{Var}(\tau)}{Q} \mu^3. \tag{9}$$

The first term on the r.h.s. of (9) converges to 1 since $S_\tau^2(n) \xrightarrow{p} \mathrm{Var}(\tau)$. The second term of (9) also converges to 1 by $\hat{\mu} \xrightarrow{p} \mu$ and the continuous mapping theorem. On the other hand, equation (4) implies $\mathrm{Var}(\tau)/Q \to \sigma^2/\mu^3$ as $Q \to \infty$. Equation (6) follows from the last three properties and Slutsky's theorem (Lehmann, 1990). ■

4.2 Derivation of (*r*, *Q*) Policies

Once the estimates of the mean and variance of the daily demand over a lead-time are obtained, they can be applied to the PA method to obtain the reorder point s and order-up-to level S. Since the current inventory position is not observable however and the retailer requires a fixed order quantity, such an order-up-to policy cannot be applied directly. Instead, we consider (r, Q) policies with a fixed order quantity Q. In this section, we discuss how to obtain such a policy based on the derived (s, S) policy so that the total cost of an (r, Q) policy is close to that of the optimal (s, S) policy. To test the accuracy of our (r, Q) policies, we use the same numerical examples from Veinott and Wanger (1965) with setup cost $K = 64$, unit holding cost $h = 1$, unit penalty cost $b = 9$, zero lead-time, and Poisson-distributed daily demand. In their paper, they provide optimal values for s and S, and the total cost based on the known information for the demand distribution. The system parameters and optimal values are presented in columns 2-3 of Table 1, with the mean demand listed in column 1.

4.2.1 Direct (*r*, *Q*) Policy

An intuitive and direct transformation of an (s, S) policy to an (r, Q) policy is to use $r = s$ as the reorder point and $Q = S - s$ as the order quantity. The direct (r, Q) policy and the respective cost are shown in the columns 4-5 of Table 1. Unfortunately, due to the inventory record inaccuracies discussed previously, the inventory position prior to the placement of an order will typically be lower than r and the fixed order quantity Q will bring it under S. The poor "fit" of this policy is evident from the entries when the mean demand is 63 or 64.

4.2.2 Adjusted (r, Q) Policy

Suppose momentarily that we use an (s, S) policy (that is, $r = s$ and $Q = S - s$) with instantaneous replenishments and known daily demand distribution. Under this policy, the first replenishment epoch is $N(S-s)+1 = N(Q)+1$ and the inventory position (prior to replenishment) in steady-state is

$$\begin{aligned} S - \mathrm{E}[Z_{N(Q)+1}] &= S - \mu[M(Q)+1] \\ &= S - \mu\left[\frac{Q}{\mu} + \frac{\sigma^2 - \mu^2}{2\mu^2} + 1 + o(1)\right] \\ &= S - Q - \frac{\sigma^2 + \mu^2}{2\mu} + o(\mu) \\ &= s - \frac{\sigma^2 + \mu^2}{2\mu} + o(\mu), \end{aligned} \tag{10}$$

where the second equality is due to equation (3).

For simplicity, one can disregard the terms $\sigma^2/(2\mu)$ and $o(\mu)$, and use the adjusted order quantity $Q = S - s + \mu/2$. (A justification for the elimination of the first term is given in Remark 3 below.) Of course, the approximation in (10) is an asymptotic result when the order quantity Q is large. In the case of a small order quantity, such a modified (r, Q) policy may deviate from the optimal policy significantly. Thus, we consider the following adjustment:

$$Q = \max\left\{S - s + \mu/2, \mathrm{EOQ} \equiv \sqrt{2K\mu/h}\right\}. \tag{11}$$

This empirical adjustment is motivated from the argument that the optimal order quantity should be related to the economic order quantity (EOQ) to allow for the trade-off between the ordering and holding cost. The respective (r, Q) policy and cost are shown in the columns 6-7 of Table 1.

Remark 2 Additional motivation for the use of the EOQ in equation (11) is provided by the case in which deliveries are instantaneous ($L = 0$) and $\sigma^2/\mu = 1$. Then the quantity D_p from equations

Table 1: Comparison of the optimal (s, S) policy and the adjusted (r, Q) policy based on the experiments in Veinott and Wagner (1965)

μ	Optimal (s, S)		Direct (r, Q)		Adjusted (r, Q)		Optimal (r, Q)	
	(s, S)	Cost	(r, Q)	Cost	(r, Q)	Cost	(r, Q)	Cost
21	(15, 65)	50.410	(15, 50)	51.411	(15, 61)	51.157	(15, 56)	50.992
22	(16, 68)	51.630	(16, 52)	52.508	(16, 63)	52.416	(16, 57)	52.193
23	(17, 52)	52.757	(17, 35)	61.532	(17, 55)	53.504	(17, 59)	53.354
24	(18, 54)	53.514	(18, 36)	62.540	(18, 56)	54.659	(18, 60)	54.516
51	(43, 110)	71.612	(43, 67)	82.510	(43, 93)	79.709	(43, 87)	79.495
52	(44, 112)	72.249	(44, 68)	83.180	(44, 94)	80.489	(43, 89)	80.236
55	(47, 118)	74.165	(47, 71)	85.185	(44, 99)	82.995	(44, 90)	82.582
59	(51, 126)	76.679	(51, 75)	87.770	(51, 105)	86.238	(51, 92)	85.566
61	(52, 131)	77.933	(52, 79)	88.655	(52, 110)	87.884	(52, 97)	86.928
63	(54, 73)	78.290	(54, 19)	$> 10^4$	(54, 90)	88.802	(54, 99)	88.376
64	(55, 74)	78.414	(55, 19)	$> 10^4$	(54, 91)	89.263	(54, 97)	89.047

(2) is equal to

$$D_p = 1.30\mu^{0.494}(K/h)^{0.506}2^{0.116} \doteq \text{EOQ}. \quad \triangleleft$$

To evaluate the performance of the adjusted (r, Q) policy, we also searched for an optimal (r, Q) policy and the corresponding cost by running 100 replications of a simulation model over 10 years. The values of r and Q were chosen from a grid around $(s, S - s)$, where s and S are given in column 2 of Table 1. The values of r and Q obtained from the simulation-based search, along with the associated costs, are displayed in columns 8-9. These numerical results show that the adjusted policy based on equation (11) works much better than the direct (r, Q) policy with $r = s$ and $Q = S - s$, and is close to the best (r, Q) policy found by the simulation experiment.

Remark 3 The consideration of the term σ^2/μ towards the derivation of the adjusted order quantity in (11) could result in larger order quantities Q. Since the optimal order quantities in column 8 of Table 1 are smaller than the adjusted order quantities in column 6 (with the optimal reorder points being equal to the respective values for the adjusted policy), such an inclusion would not have a positive effect in this case. $\triangleleft$

5 Numerical Studies for Estimation Accuracy and Policy Performance

In the previous section, we derived asymptotic estimators for the first two moments of the daily demand based on partial demand information from the DBO. In this section we evaluate the "fit" of these estimators by using them to create a policy based on the PA method, and then test the performance of the policy via simulation. We start with an experimental design based on various demand distributions, setup and unit order costs, lead-times, and constant order quantities; this design grid is similar to the grid of Ehrhardt (1979). At each design point we perform a simulation experiment consisting of 100 independent replications. Each replication has two phases. The first phase collects realizations of DBO over a period of 2 years based on a fixed reorder point $r = 15$. This DBO data is used to compute the estimates of μ and σ^2 via equations (5) and (6). The PA method in §3 then uses estimates to derive the approximately optimal (s, S) inventory policy, which is then converted to the adjusted (r, Q) policy using $r = s$ and Q from equation (11). At this instance, the second phase starts with an inventory position of zero and an order of size Q, continues for a time window of 5 years, and returns a realization of the average total cost per day. (The long duration of this phase suffices to alleviate the effects of the initial conditions.) We benchmark the overall performance of our approach in terms of the average total daily cost via comparisons with instances where the demand is observable, the unknown parameters are estimated by their sample moments, and the PA-based (s, S) policy is implemented; we characterize those instances as associated with *fully observable demand*.

Remark 4 Jacobs and Wagner (1989) studied the impact of demand estimation on the total system cost under fully observable demands. In particular, they showed that exponentially smoothed averages and absolute deviations can outperform the sample mean and sample standard deviation, respectively, when the (s, S) policy is computed via the PA method of Ehrhardt (1979). We do not employ the estimators of Jacobs and Wagner (1989) for two reasons: (1) the two-year data collection period is substantially long to allow the computation of sound estimates for μ and σ when the demand is fully observable, and (2) our estimators are not based on smoothing techniques. However, the development of more *robust* estimators for μ and σ (e.g., using exponentially weighted moving averages) when the demand is partially observable is an important problem that is worth future investigation. $\triangleleft$

We use a grid of 216 test cases based on the parameter assignments in Table 2. Three types of demand distributions are used: Poisson, and negative binomial with variance-to-mean ratios of 3 and 5. The appropriateness of these discrete models has been discussed in Chatfield et al. (1966) and Ehrhardt (1979). Each demand distribution is given two mean values, 8 and 16. Two values, 2 and 4, are assigned to lead-time. Since the cost function is linear in the parameters K, b, and h, the value of the unit holding cost is a redundant parameter which is set at unity. The unit penalty costs are 4, 24 or 99, and the setup costs are 32 or 64. The unit replenishment cost is unspecified because it does not effect the computation of an optimal policy for undiscounted and infinite-horizon models. Based on the system parameters above, the reorder point of $r = 15$, and the well-known EOQ model, we calculate the minimal and maximal order quantity for each design point; the minimum is around 25 and the maximum is around 46. Therefore, the constant order quantity is specified at three values: 20, 40, and 80. All combinations of these parameters settings form the experimental grid of 216 cases. Table 11 in the Appendix lists the parameters for each case. Cases 1-72 correspond to Poisson demand with parameter μ, and cases 73-216 correspond to negative binomial demand with parameters r and p.

5.1 Performance Evaluation

In this section, we proceed with a thorough evaluation of the performance of the estimates of the first two moments of the daily demand and the effectiveness of our methodology with regard to the total average daily cost. We first examine the statistical properties of the estimators (5) and (6)

Table 2: System parameters for experimental grid

Factor	Levels	Number of Levels
Demand distribution	Poisson ($\sigma^2 = \mu$) Negative Binomial ($\sigma^2 = 3\mu$) Negative Binomial ($\sigma^2 = 5\mu$)	3
Mean demand (μ)	8, 16	2
Replenishment lead-time (L)	2, 4	2
Replenishment setup cost (K)	32, 64	2
Unit penalty cost (b)	4, 24, 99	3
Unit holding cost (h)	1	1
Order quantity (Q)	20, 40, 80	3

using several standard statistical measurements for their accuracy and precision. Let n_i denote the number of observed DBO within simulation run i, $i = 1, \ldots, 100$. Recall that in the first phase of the experiment, the total number of daily demands is equal to 730 (2 years) but the number of observed DBO is a random variable.

Let $\tau_{i,j}$, $j = 1, \ldots, n_i$, denote the observed DBO during the ith replication. The estimated mean within the ith replication is $\hat{\mu}_i = Q/\bar{\tau}_i$, where $\bar{\tau}_i = \sum_{j=1}^{n_i} \tau_{i,j}/n_i$. The estimate of the mean daily demand over all samples is $\hat{\mu} = Q/\bar{\tau}$, where $\bar{\tau} = \sum_{i=1}^{100} \bar{\tau}_i/100$. Let $S^2_{\tau,i}$ denote the sample variance of the DBO within the ith replication. By equation (6), the estimate of the standard deviation of daily demand based on replication i is

$$\hat{\sigma}_i = \sqrt{S^2_{\tau,i} Q^2/\bar{\tau}_i^3}. \tag{12}$$

The overall estimate of the standard deviation of daily demand is obtained by avaraging the estimates from all replications:

$$\hat{\sigma} = \frac{1}{100} \sum_{i=1}^{100} \hat{\sigma}_i. \tag{13}$$

In addition, we define the overall sample variance of historical DBOs by $S^2_\tau = \sum_{i=1}^{100} S^2_{\tau,i}/100$.

In order to evaluate the accuracy and precision of the estimates, we use the following measures. The first measure is the Relative Root Mean Squared Error (RRMSE) of the mean estimate, denoted by

$$\mathrm{RRMSE}_{\mu} = \sqrt{\frac{\sum_{i=1}^{100} \left(\hat{\mu}_i - \mu\right)^2}{100}} \Big/ \mu. \tag{14}$$

RRMSE is also called the standard error of the estimate and measures the spread of the estimate from the true value. The second measure is called the Relative Standard Deviation (RSD) of the mean estimate, denoted by

$$\mathrm{RSD}_{\mu} = \sqrt{\frac{\sum_{i=1}^{100} \left(\hat{\mu}_i - \hat{\mu}\right)^2}{99}} \Big/ \mu. \tag{15}$$

Moreover, we define the Relative Bias (RBias) by

$$\mathrm{RBias}_{\mu} = \left(\hat{\mu} - \mu\right) / \mu, \tag{16}$$

which measures the relative precision of the estimate. The detailed numerical results are listed in Tables 12 to 17 in the Appendix.

Table 3 contains a summary of the experimental results. The averages of RBias_{μ}, RSD and RRMSE are listed in rows 2-10 based on the level of the coefficient of variation σ/μ and Q, with the remaining parameters being fixed. The last three rows contain the maximum, minimum, and overall average of RBias_{μ}, RSD and RRMSE. Based on these results, we can infer that the estimate (5) of the mean demand works as well as the sample mean of the daily demands as the differences of all three measurements are small. For example, the average of RRMSE is 0.0186 under fully observable demand and 0.0187 under partially observed demand. Furthermore, the difference of the average of RBias_{μ} between partial information and full information is only 0.0004.

Notice that the size of the constant order quantity has negligible effect on the estimate of demand mean, and that the estimate of the mean daily demand performs well throughout the entire experimental grid. As the ratio σ/μ increases, the average of RBias_{μ} and RSD slightly increase as well; this behavior is expected.

Table 3: Summary results for the mean demand estimates under full and partial observation

Parameter		Full Observation			Partial Observation		
		RBias	RSD	RRMSE	RBias	RSD	RRMSE
σ/μ	0.250	0.0002	0.0092	0.0092	0.0004	0.0092	0.0093
	0.354	-0.0004	0.0132	0.0133	-0.0001	0.0132	0.0133
	0.433	0.0002	0.0161	0.0162	0.0006	0.0162	0.0163
	0.559	0.0009	0.0208	0.0209	0.0014	0.0208	0.0210
	0.612	0.0011	0.0227	0.0228	0.0016	0.0227	0.0229
	0.791	0.0003	0.0291	0.0292	0.0013	0.0294	0.0295
Q	20	0.0005	0.0188	0.0188	0.0010	0.0188	0.0189
	40	0.0003	0.0184	0.0185	0.0008	0.0185	0.0186
	80	0.0005	0.0183	0.0184	0.0008	0.0185	0.0186
Max		0.0067	0.0342	0.0342	0.0074	0.0342	0.0344
Min		-0.0041	0.0079	0.0079	-0.0034	0.0080	0.0081
Average		0.0004	0.0185	0.0186	0.0008	0.0186	0.0187

Similarly, we define the error measures for the standard deviation estimate as follows:

$$\mathrm{RRMSE}_{\sigma} = \sqrt{\frac{\sum_{i=1}^{100}\left(\hat{\sigma}_i - \sigma\right)^2}{100}} \Big/ \sigma, \tag{17}$$

$$\mathrm{RSD}_{\sigma} = \sqrt{\frac{\sum_{i=1}^{100}\left(\hat{\sigma}_i - \widehat{\sigma}\right)^2}{99}} \Big/ \sigma, \tag{18}$$

and $\mathrm{RBias}_{\sigma} = \left(\hat{\sigma} - \sigma\right)/\sigma$. The numerical comparisons under fully and partially observed demand are displayed in Tables 18 to 23 in the Appendix. Table 4 contains a summary of these results.

From Table 4 we observe that the estimate of the standard deviation based on partially observed demand performs worse than the estimate obtained under fully observed demand (the last estimator is the sample variance of the daily demands). Since the average of RSD_{σ} under partial demand observation is close to that under full observation, the error mainly comes from RBias_{σ}: the overall average RBias_{σ} under partial observation is almost 50 times larger than that under full observation. Thus, we focus our analysis on RBias_{σ}. Table 4 indicates that both the constant order quantity Q and the coefficient of variation σ/μ significantly affect the performance of the estimate of the standard deviation. As the constant order quantity Q increases, RBias_{σ}, RSD_{σ} and RRMSE_{σ} decrease quickly. Once the order quantity Q becomes large enough, our estimate of standard deviation works well with the average RBias_{σ} being only 0.0784 when $Q = 80$. The impact of the demand coefficient of variation σ/μ on the estimate for σ is more complex. For example, the maximum RBias_{σ} under Poisson demand and order quantity 20 is equal to 0.7476. For the same type of demand distribution but smaller mean demand, the RBias_{σ} decreases: as the mean daily demand drops from 16 to 8 ($\sigma/\mu = 0.354$ and 0.250, respectively), the average RBias_{σ} drops from 0.4764 to 0.1471, respectively. In general, as the ratio σ/μ increases, the RBias_{σ} exhibits a (non-monotone) upward trend. This suggests that improved estimators for σ^2 are needed; this issue is the topic of Section 6.

5.1.1 Evaluation With Regard to Average Total Cost

Next we examine the impact of the demand estimation on the cost estimates. Let C^* denote the average total daily cost for a given case under the (s, S) policy computed via the PA method with the

Table 4: Summary results for the standard deviation estimates for daily demand under full and partial observation

Parameter		Full Observation			Partial Observation		
		RBias	RSD	RRMSE	RBias	RSD	RRMSE
σ/μ	0.250	0.0009	0.0266	0.0267	0.4764	0.0580	0.4817
	0.354	0.0034	0.0277	0.0280	0.1471	0.0709	0.1699
	0.433	0.0005	0.0307	0.0309	0.1834	0.0553	0.1950
	0.559	0.0014	0.0350	0.0352	0.1066	0.0537	0.1243
	0.612	0.0031	0.0348	0.0353	0.0492	0.0671	0.0889
	0.791	0.0028	0.0411	0.0415	0.0253	0.0703	0.0765
Q	20	0.0022	0.0329	0.0332	0.2638	0.0493	0.2734
	40	0.0022	0.0326	0.0329	0.1517	0.0593	0.1723
	80	0.0017	0.0324	0.0327	0.0784	0.0791	0.1224
Max		0.0178	0.0491	0.0498	0.7476	0.0998	0.7490
Min		-0.0067	0.0214	0.0214	0.0035	0.0362	0.0594
Average		0.0020	0.0327	0.0329	0.1647	0.0626	0.1894

actual mean and variance of daily demand. Let $C_{p0,i}$ denote the average total daily cost for a specific case from the ith replication that uses the estimated mean and variance of daily demand under fully observable demand. The average from all replications is denoted by $C^*_{p0} = \sum_{i=1}^{100} C_{p0,i}/100$. Similarly, $C_{p1,i}$ denotes the respective average total daily cost from the ith replication based on the estimated mean and variance under partial information with the overall average denoted by $C^*_{p1} = \sum_{i=1}^{100} C_{p1,i}/100$.

We examine the policy performance based on the known mean and variance of the daily demand. Again we use the measures RRMSE, RSD and RBias to evaluate the policy's performance in terms of average total daily cost by simply replacing $\hat{\mu}_i$, $\hat{\mu}$ and μ by $C_{p0,i}$ ($C_{p1,i}$), C^*_{p0} (C^*_{p1}) and C^*, respectively. The numerical results are presented in Tables 24 to 29 (Appendix), and the summary is presented in Table 5. When the demand is observable, the average total cost is very close to the nearly optimal average total cost; the average RBias over all 216 cases is very small, only 0.0009. Since in most of the 216 cases the demand variance under partial observation is overestimated, most of these cases are associated with higher average total cost than the nearly optimal cost. On the other hand, the average RBias under partial information is 0.047. As the order quantity increases, the RBias under partial demand observation decreases. In the worst case, the RBias is 0.3762.

To reflect practical situations where historical demand data are used to obtain estimates of the mean and variance of the demand, Ehrhart (1979) tested the cost performance of the PA method by substituting estimates of the demand mean and variance in place of the actual parameters by simulating 72 systems, each having a negative binomial demand distribution with $\sigma^2/\mu = 9$. He found that the use of classical estimates with a year's worth of weekly demand history resulted in an aggregate cost that is approximately 6% above the optimal cost for known demand parameters. Based on his results, we propose the following performance measure for each simulation run:

$$\Delta_{p1,i} = \frac{\left(C_{p1,i} - C_{p0,i}\right) \times 100}{C_{p0,i}}, \quad i = 1, \ldots, 100, \tag{19}$$

namely, the percentage by which the average total cost under partially observed demand exceeds that under fully observable demand. The average (denoted by $\bar{\Delta}_{p1}$) and standard deviation of $\Delta_{p1,i}$ over the k replications are displayed in Tables 30 to 32 in the Appendix. These results are summarized

Table 5: Summary results for the average total daily cost under full and partial observation

Parameter		Full Observation			Partial Observation		
		RBias	RSD	RRMSE	RBias	RSD	RRMSE
σ/μ	0.250	-0.0011	0.0144	0.0156	0.1043	0.0213	0.1095
	0.354	-0.0007	0.0167	0.019	0.0275	0.0226	0.0404
	0.433	0.0011	0.0226	0.0236	0.0690	0.0310	0.0780
	0.559	0.0037	0.029	0.0298	0.0523	0.0369	0.0671
	0.612	0.0067	0.0257	0.0278	0.0233	0.0341	0.043
	0.791	-0.0046	0.0326	0.0336	0.0059	0.0429	0.0443
Q	20	0.0009	0.0239	0.0253	0.0766	0.029	0.0873
	40	0.0009	0.0235	0.0248	0.0429	0.0306	0.0584
	80	0.0008	0.0231	0.0246	0.0217	0.0348	0.0454
Max		0.0244	0.05	0.05	0.3762	0.0734	0.378
Min		-0.0177	0.0098	0.0103	-0.0135	0.0084	0.0085
Average		0.0009	0.0235	0.0249	0.0471	0.0315	0.0637

Table 6: Frequencies for $\bar{\Delta}_{p_1}$ under partial demand observation

Range for $\bar{\Delta}_{p_1}$	Number of Cases	Percentage of Cases
$(-\infty, -1.5)$	0	0.0
$[-1.5, 1.5]$	72	33.3
$(1.5, 5)$	76	35.2
$(5, \infty)$	68	31.5

in Table 6, which lists the number of cases in the system having values of $\bar{\Delta}_{p_1}$ in various ranges. 33% of the cases fall in the range $[-1.5, 1.5]$, 35% fall in $(1.5, 5]$, and the remaining 31% fall in $(5, \infty)$. The large percentage of cases with $\bar{\Delta}_{p_1} > 5$ motivates the approach in the next section.

6 Power Approximation for Variance of Daily Demand (PAD)

In this section, we propose an adaptive method to improve the estimate of the variance of the daily demand by using a regression model. Recall that the variance estimator (6) becomes unbiased as $Q \to \infty$ and $n \to \infty$. Since the experiments in §5 use a sufficient amount of data, the significant bias of the variance estimator is most likely due to the small values of the respective order quantities. To address this issue, we consider estimators of the form

$$\sigma^2 = C(S_\tau^2)^\alpha Q^\beta (\bar{\tau})^\gamma, \tag{20}$$

where C, α, β, and γ are constants to be fitted and the index "n" is dropped for simplicity. The model (20) is motivated by the following observations: we see from the previous experiments that the variance estimate degrades with smaller order sizes so the inclusion of Q is obvious. We also see that the coefficient of variation of the demand has an effect, but since we do not know the demand distribution parameters we use S_τ^2 and $\bar{\tau}$ as surrogates for this coefficient. We form a linear model by taking the logarithm of equation (20) and using least squares regression to fit the model based on the experimental grid in Table 2. The independent variables are the sample average of

DBO, the sample variance of DBO, and the fixed order quantity Q. (Of course, we realize that the independent variables of model 21 are statistically dependent, at least in small samples.) The dependent variable is the true variance of the daily demand. The regression model generates the following estimator for σ^2:

$$\sigma_p^2 = 0.7418(S_\tau^2)^{1.2685} Q^{2.0012}/(\bar{\tau})^{3.0060}, \tag{21}$$

where again the subscript "p" stands for "power". For a comparison with the estimator (6), notice that the exponent of the order quantity Q is close to 2 and the exponent of the overall average of DBO is close to -3. However, the exponent of the sample variance of DBO is slightly larger than 1, and the model includes a multiplier that is less than 1. The model in equation (21) has a coefficient of determination $R^2 = 0.9633$.

6.1 Evaluation of the PAD Estimator

In this section, we evaluate the (incremental) effect of the PAD estimate in equation (21). As in Section 5.1, we use the measures, RBias, RSD and RRMSE to evaluate the performance of the variance estimate and the average total daily cost. The numerical results are displayed in the last three columns of Tables 18 to 29 (see Appendix). We summarize the performance of the estimate (21) and the average total cost in Tables 7 and 8, respectively.

Table 7 indicates that the PAD yields an improved estimator compared to (6). (To facilitate the comparison we copy columns 3-5 from Table 5.) For instance, the average RBias decreases by a factor of about 100 (from 0.1647 to 0.0016). Most importantly, the RSD does not vary significantly over the 216 cases. Thus, the regression method "flattens" the bias and makes the estimate less sensitive to the demand parameters and the order quantity. For example, the average absolute RBias is 0.0099, -0.0098 and 0.048 when Q is at 20, 40 and 80, respectively.

Figure 1 plots the differences of the RBias, RSD and RRMSE based on equation (6) and PAD. This figure indicates that the differences have certain patterns based on the demand distribution. The differences of RSD vary slightly around zero. When the mean demand is 16, many RSDs

Table 7: Summary results for the standard deviation estimate under partial demand observation and PAD

Parameter		Equation (6)			PAD		
		RBias	RSD	RRMSE	RBias	RSD	RRMSE
σ/μ	0.250	0.4764	0.0580	0.4817	0.0969	0.0545	0.1388
	0.354	0.1471	0.0709	0.1699	-0.0377	0.0759	0.0858
	0.433	0.1834	0.0553	0.1950	-0.0362	0.0558	0.0686
	0.559	0.1066	0.0537	0.1243	-0.0507	0.0559	0.0781
	0.612	0.0492	0.0671	0.0889	-0.0021	0.0798	0.0911
	0.791	0.0253	0.0703	0.0765	0.0396	0.0874	0.1095
Q	20	0.2638	0.0493	0.2734	0.0099	0.0477	0.0934
	40	0.1517	0.0593	0.1723	-0.0098	0.0636	0.0834
	80	0.0784	0.0791	0.1224	0.0048	0.0933	0.1091
Max		0.7476	0.0998	0.7490	0.2481	0.1351	0.2522
Min		0.0035	0.0362	0.0594	-0.0722	0.0336	0.0370
Average		0.1647	0.0626	0.1894	0.0016	0.0682	0.0953

under equation (6) are less than those by PAD. The differences with regard to RBias and RRMSE is obviously larger when the mean demand is 16.

Since the differences of RRMSE have the same pattern as those of RBias, we plot the differences for RBias under equation (6) and under PAD based on the demand parameter σ/μ in Figure 2. From this figure, we note as the ratio σ/μ increases, the PAD-based estimate becomes larger than the estimate under equation (6).

Based on Table 8, PAD significantly reduces the average total cost to values close to the cost of the nearly optimal (s, S) policy obtained via the PA method. The average RBias over the 216 cases is 0.0002, which is less than the average RBias under fully observable demand. This implies that the total cost induced by PAD is close to the cost under fully observable demand. The average total cost for each case is less sensitive to the demand parameters and the order quantity. The range of RBias on the average total cost is from 0.1170 to -0.0443, which is considerably narrower than

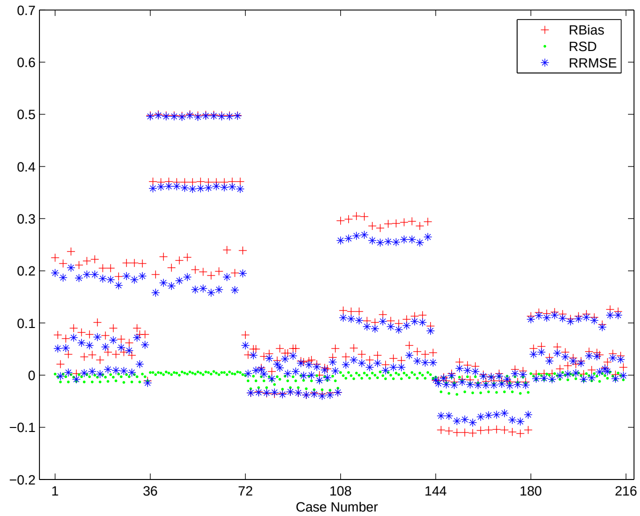

Figure 1: Difference of RBias, RSD and RMSE of the average total cost based on equation (6) and PAD

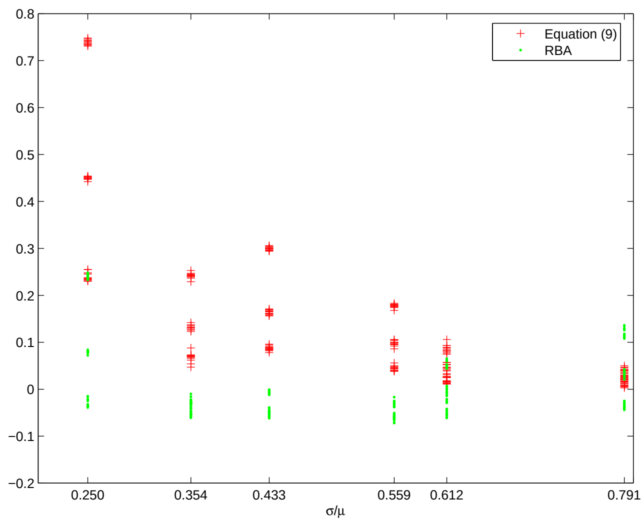

Figure 2: Comparison between RBias for the average total daily cost under partial information and equation (6) and RBias under PAD based on variable coefficient of variation σ/μ

Table 8: Summary results for the average total cost under equation (6) and PAD

Parameter		Equation (6)			RBA		
		RBias	RSD	RRMSE	RBias	RSD	RRMSE
σ/μ	0.250	0.1043	0.0213	0.1095	0.0191	0.0189	0.0347
	0.354	0.0275	0.0226	0.0404	-0.0076	0.0225	0.0264
	0.433	0.069	0.031	0.078	-0.0109	0.0283	0.0319
	0.559	0.0523	0.0369	0.0671	-0.0179	0.0352	0.0408
	0.612	0.0233	0.0341	0.043	0.0057	0.0370	0.0415
	0.791	0.0059	0.0429	0.0443	0.0130	0.049	0.0571
Q	20	0.0766	0.029	0.0873	-0.0019	0.0266	0.0368
	40	0.0429	0.0306	0.0584	-0.0032	0.0302	0.0351
	80	0.0217	0.0348	0.0454	0.0059	0.0387	0.0443
Max		0.3762	0.0734	0.3780	0.1170	0.0994	0.1297
Min		-0.0135	0.0084	0.0085	-0.0443	0.0102	0.0111
Average		0.0471	0.0315	0.0637	0.0002	0.0318	0.0387

the range based on equation (6).

As in §5.1.1, let $C_{p_2,i}$ denote the average total cost from the ith replication based on the demand estimate from PAD. We define the percentage differences

$$\Delta_{p_2,i} = \frac{\left(C_{p_2,i} - C_{p_0,i}\right) \times 100}{C_{p_0,i}}, \quad i = 1, \ldots, 100. \tag{22}$$

The average ($\bar{\Delta}_{p_2}$) and standard deviation of $\Delta_{p_2,i}$ over the 100 replications are listed in Tables 30 to 32 (see Appendix). Similarly, we list the number of cases in various ranges in Table 9. Nearly 60% of the cases fall in the range ± 1.5, a remarkable 94% falls in the range ± 5, while only 6% induce $\bar{\Delta}_{p_2} > 5$. Therefore, it seems that the combination of the PA and PAD methods works as well as the PA method alone under fully observable demand in terms of the mean total cost per day.

Table 9: Frequencies for $\bar{\Delta}_{p_2}$ under PAD

Range for $\bar{\Delta}_{p_2}$	Number of Cases	Percentage of Cases
$(-\infty, -5)$	0	0.0
$[-5, -1.5)$	60	27.8
$[-1.5, 1.5]$	112	51.9
$(1.5, 5]$	31	14.4
$(5, \infty)$	13	6.0

6.2 Extrapolation

We end the experimental evaluation of PAD with a few cases that are outside the experimental grid in Table 2. A single case with interpolated parameter settings is used as a base case: negative binomial daily demand with $\sigma^2 = 4\mu$, $\mu = 12$, $L = 3$, $h = 1$, $b = 49$, $K = 48$, and $Q = 60$. We create each test case by changing the value of one parameter while holding the other parameters at their base values. The performance measures relative to the mean total cost per day are presented in Table 10.

Based on Table 10, the average RBias induced by PAD and equation (6) are -0.001 and 0.06, respectively. This suggests that PAD works well with regard to the aggregated average total cost even though the parameter values are beyond the original range. Based on the average RBias of 0.006 from column 3, the PAD variance estimator and equation (11) yield an (r, Q) policy that rivals the near optimal (s, S) policy resulting from the PA method when the daily demand is observed.

7 Conclusions

In this paper, we provide an inventory policy for a retailer facing stochastic demand and lead-times, linear holding and backordering costs, fixed order cost, and who is restricted to placing orders at predetermined time intervals. The policy must be robust to system reported inventory inaccuracies

Table 10: Single parameter extrapolations

Parameter	Value	Full Observation			Partial Observation			RBA		
		RBias	RSD	RRMSE	RBias	RSD	RRMSE	RBias	RSD	RRMSE
σ^2/μ	7	0.003	0.039	0.039	0.014	0.049	0.051	0.015	0.055	0.057
μ	24	0.014	0.026	0.030	0.112	0.036	0.118	-0.013	0.033	0.035
	30	0.002	0.026	0.027	0.151	0.038	0.156	-0.007	0.034	0.034
	36	-0.002	0.021	0.021	0.213	0.033	0.215	0.019	0.028	0.033
K	16	-0.004	0.038	0.039	0.031	0.059	0.067	-0.018	0.065	0.067
	72	0.008	0.025	0.026	0.026	0.034	0.043	-0.002	0.038	0.038
p	2	0.005	0.017	0.018	0.003	0.017	0.018	0.008	0.018	0.020
	132	0.008	0.030	0.031	0.036	0.045	0.058	-0.006	0.049	0.049
	199	0.006	0.033	0.034	0.043	0.051	0.066	-0.003	0.055	0.055
L	0	0.000	0.016	0.016	0.014	0.023	0.027	-0.004	0.025	0.025
	5	-0.003	0.034	0.034	0.028	0.045	0.053	-0.013	0.048	0.050
Q	10	0.011	0.027	0.029	0.119	0.031	0.123	-0.027	0.026	0.037
	15	0.009	0.030	0.032	0.097	0.037	0.104	-0.023	0.032	0.039
	85	0.013	0.028	0.031	0.029	0.041	0.050	0.014	0.047	0.050
	90	0.012	0.028	0.030	0.032	0.048	0.058	0.025	0.056	0.062
	100	0.007	0.032	0.033	0.021	0.046	0.050	0.019	0.055	0.058

and reflect other common constraints of the retailing industry, such as a stable order quantity and a simple policy that does not require dynamic (re-)optimization and can be programmed in a spreadsheet or database software application. To meet these requirements, we propose an (r, Q) inventory policy that minimizes the long-run mean (undiscounted) total cost per period. Our choice of an (r, Q) policy over an (s, S) policy comes from the retailer's desire for a stable order quantity in each period. To this end, we first construct estimators for the first two moments of the demand distribution using well-known results from renewal theory and realizations of periods between orders. We proceed with the derivation of approximately optimal inventory policies similar in spirit to the simple and robust Power Approximation (PA) method first described in Ehrhardt (1979). The PA method has been used successfully in a variety of settings and owns its popularity to its simplicity and the surprisingly good fit of the regression model to a variety of demand distributions. Since the suppliers of the retailer prefer to ship fixed quantities that do not vary from one period to the next, we propose an (r, Q) policy with $r = s$ that approximates the near optimal (s, S) policy calculated using the PA method.

The performance of the (r, Q) policy is evaluated using a simulation study over a range of distributions and parameter values from an experimental grid that consists of 216 design points and is similar to Ehrhardt's experimental design. Since the cost of the (r, Q) policy is sensitive to the variance of the demand per period, we develop an alternative regression-based estimator of the variance of the periodic demand using the same grid. The incorporation of this estimator results in substantially improved (r, Q) policies with regard to the mean total cost per period. The relative difference between the average cost resulting from the PA method with the mean estimator (5) and the variance estimator (21) for the demand and the average cost resulting from the PA method under fully observable demand is within $\pm 5\%$ for 94% of the 216 cases in our experimental grid.

We end with a few problems worth future research. First, one could incorporate service level constraints; this metric is frequently used by retailers as penalty costs are hard to obtain. An accurate PA method for the case in which the backorder cost is replaced by a service level constraint was proposed by Schneider and Ringuest (1990). Second, one could investigate the application of smoothing techniques to the estimation of the demand moments μ and σ (see Remark 4). For instance, the sample mean $\bar{\tau}(n)$ in equations (5)-(6) can be replaced by an exponentially smoothed average and the sample standard deviation $S_\tau(n)$ in equation (6) can be replaced by a modified absolute deviation (cf. Jacobs and Wagner 1989). Third, one could look at modeling the tradeoff between improving the accuracy of system reported inventory and the associated costs. Fourth, one could study the effect of order size on the degree of inventory inaccuracy. For instance, one would expect that large order sizes would induce more unobservable demand of the types mentioned in the Introduction. Finally, our model can be modified to address the possibility of lost sales.

References

[1] Atali, A., Lee, H., and Ozer, O. (2004) If the inventory manager knew: Value of RFID under imperfect inventory information. Working Paper. Stanford Graduate School of Business, Stanford University, Stanford, CA.

[2] Chatfield, C., Ehrenberg, A., and Goodhardt, G. (1966) Progress on a simplified model of stationary purchasing behaviour. *Journal of the Royal Statistical Society A,* **129**, 317-360.

[3] DeHoratius, N. and Raman, A. (2004) Inventory record inaccuracy: An empirical analysis. Working Paper. University of Chicago, Chicago, IL.

[4] ECR Europe (2003) Shrinkage: Introducing a collaborative approach to reducing stock loss in the supply chain. ECR Europe, Brussels, Belgium.

[5] Ehrhardt, R. (1979) The power approximation for computing (s, S) policies. *Management Science*, **25**, 777-786.

[6] Ehrhardt, R. (1984) (s, S) policies for a dynamic inventory model with stochastic lead-times. *Operations Research*, **32**, 121-132.

[7] Ehrhardt, R. (1985) Easily computed approximations for (s, S) inventory system characteristics. *Naval Research Logistics Quarterly*, **32**, 347-359.

[8] Ehrhardt, R. and Mosier, C. (1984) A revision of the power approximation for computing (s, S) policies. *Management Science*, **30**, 618-622.

[9] Feller, W. (1949) Fluctuation theory of recurrent events. *Transactions of the American Mathematical Society*, **67**, 98-119.

[10] Heyman, D.P., and Sobel, M.J. (1982) *Stochastic Models in Operations Research, Volume I*. McGraw-Hill, New York.

[11] Hollinger, R.C. (2003) 2002 National retail security survey final report. Working Paper. University of Florida, Gainsville, FL.

[12] Iglehart, D.L. (1963a) Optimality of (s, S) policies in the infinite horizon dynamic inventory problem. *Management Science*, **9**, 259-267.

[13] Iglehart, D.L. (1963b) Dynamic programming and stationary analysis of inventory problems. In *Multistage Inventory Models and Techniques*, eds. H.E. Scarf, D.M. Gilford, and M.W. Shelly. Stanford University Press, Stanford, CA.

[14] Jacobs, R. and Wagner, H.M. (1989) Reducing inventory system costs by using robust demand estimators. *Management Science*, **435**, 771-787.

[15] Kang, Y. and Gershwin, S.B. (2004) Information inaccuracy in inventory systems — stock loss and stockout. Working Paper. Department of Mechanical Engineering, Massachusetts Institute of Technology, Boston, MA.

[16] Kok, A.G. and Shang, K.H. (2005) Inspection and replenishment policies for systems with inventory record inaccuracy. Working Paper. Fuqua School of Business, Duke University, Durham, NC.

[17] Lehmann, E. (1990) *Elements of Large Sample Theory*. Springer-Verlag, New York.

[18] Porteus, E.L. (1985) Numerical comparisons of inventory policies for periodic review systems. *Operations Research*, **33**, 134-152.

[19] Roberts, D.M. (1962) Approximations to optimal policies in a dynamic inventory model. Chapter 13 in *Studies in Applied Probability and Management Science*, eds. K. Arrow, S. Karlin, and H. Scarf. Stanford University Press, Stanford, CA.

[20] Scarf, H.E. (1960) The optimality of (s, S) policies in the dynamic inventory problem. In *Mathematical Methods in the Social Sciences*. Stanford University Press, Stanford, CA.

[21] Schneider, H. and J.L. Ringuest (1990) Power approximation for computing (s, S) policies using service level. *Management Science*, **36**, 822-834.

[22] Veinott, A.F. and Wagner, H.M. (1965) Computing optimal (s, S) inventory policies. *Management Science*, **11**, 525-552.