OPERATIONS RESEARCH
Vol. 54, No. 2, March–April 2006, pp. 311–323
ISSN 0030-364X | EISSN 1526-5463 | 06 | 5402 | 0311

DOI 10.1287/opre.1050.0258

An Optimal Lot-Sizing and Offline Inspection Policy in the Case of Nonrigid Demand

Shoshana Anily
Faculty of Management, Tel Aviv University, Tel Aviv 69978, Israel, anily@post.tau.ac.il

Abraham Grosfeld-Nir
Academic College of Tel-Aviv-Yaffo, Tel Aviv, Israel, agn@mta.ac.il

A batch production process that is initially in the in-control state can fail with constant failure rate to the out-of-control state. The probability that a unit is conforming if produced while the process is in control is constant and higher than the respective constant conformance probability while the process is out of control. When production ends, the units are inspected in the order they have been produced. The objective is to design a production and inspection policy that guarantees a zero defective delivery in minimum expected total cost.

The inspection problem is formulated as a partially observable Markov decision process (POMDP): Given the observations about the quality of the items that have already been inspected, the inspector should determine whether to inspect the next unit or stop inspection and possibly pay shortage costs. We show that the optimal policy is of the control limit threshold (CLT) type: The observations are used to update the probability that the production process was still in control while producing the candidate unit for inspection. The optimal policy is to continue inspection if and only if this probability exceeds a CLT value that may depend on the outstanding demand and the number of uninspected items. Structural properties satisfied by the various CLT values are presented.

Subject classifications: dynamic programming; optimal control; Markov: applications; probability: stochastic model application.
Area of review: Manufacturing, Service, and Supply Chain Operations.
History: Received June 2002; revised October 2002, July 2003, December 2004, January 2005; accepted January 2005.

1. Introduction

In a world of growing, complexity targeting zero defectives becomes an absolute necessity if the possibility of complex and expensive systems becoming inactive due to the failure of quite simple components is to be avoided. This reality may lead to meticulous inspection procedures, which sometimes are more costly than manufacturing. Juran (1993) expressed a view about the increasing role of quality control in supporting the firm's long-term competitive strategy as follows "While the twentieth century has been the Century of Productivity, the twenty first century will be the Century of Quality" (p. 47). The need for elaborate inspection is particularly important when products are custom made, in small quantities, because often in such situations neither the manufacturer nor the customer has genuine experience concerning the product's quality.

We consider a batch production process as described in Montgomery (2001): Items in a batch are processed sequentially. A production process in (statistical) control is operating with only chance causes of variations, i.e., a background noise that results in an inherent (acceptable) amount of natural variability. Other types of variability may arise as a result of improperly adjusted or controlled machines, operator errors, or defective raw material. These are called *assignable causes*. A process operating in the presence of such causes is said to be *out of control*. The production process switches to the out-of-control state at the first occurrence of an assignable cause. While in control, the production process may get out of control with a constant failure rate while producing the next unit in the batch. Once the process becomes out of control, it remains in this state until the batch production is completed. Ross (1983, p. 25) referred to continuous-time constant failure rate processes whose failure time is exponentially distributed. We consider the analogous discrete-time case of such processes where the failure time is geometrically distributed.

Set-up operations take place before initiating the production of a batch to restore the optimal operating conditions. Manufactured units are classified by attributes as either conforming or nonconforming to the customer's requirements. More specifically, while the process is in control, the probability that a produced unit is conforming is constant and higher than the constant respective conformance probability while the process is out of control. We assume that it is impossible to carry out a direct inspection of the production process as long as it works due to the inaccessibility of some of the process's components, but the process deterioration can be measured indirectly by inspecting the output.

A major objective of statistical process control (SPC) is to perform online inspection to detect as early as possible the occurrence of assignable causes, which enables the operator to undertake corrective actions. The literature distinguishes between two types of inspection rules: (1) static SPC rules with fixed parameters (sampling interval, sample size, and control chart limits) and (2) dynamic SPC (DSPC) rules, where one or more parameters are variable; see Tagaras (1998). Because the Bayesian approach in developing SPC rules for online inspection bears similarities with our formulation of offline inspection, we mention briefly some relevant SPC-oriented works. Calabrese (1995) analyzed, based on the Bayesian approach, SPC schemes with fixed sample sizes and sampling intervals in the finite time horizon: The samples' observations are used to derive the optimal decision rule on whether to continue production for the next interval or to stop and inspect the process for assignable causes. Porteus and Angelus (1997) analyzed similar infinite- and finite-horizon practical problems where inspection time and restoration time are not necessarily negligible. They questioned the effectiveness of static SPC rules that are widely used in practice, in comparison to DSPC rules. Tagaras and Nikolaidis (2002) compared the economic performance of various Bayesian dynamic $\overline{X}$ charts in finite productions runs when the sample interval and the sample size are allowed to be determined dynamically. They find that adaptive sample intervals have the most positive impact on the economic performance. The main criticism about DSPC rules, which were shown to be optimal, is their complexity; see Parkhideh and Case (1989). Still, DSPC rules have been used in practice for many years (see Western Electric Inc. 1958), and the practical interest in them keeps growing. Tagaras (1994, 1996) showed by a numerical study that at least 10% quality cost savings can be achieved by using DSPC policies rather than static rules.

Sometimes online inspection is impossible, usually because the inspection is performed in an exogenous facility where special equipment is installed. Offline inspection, which is the subject of this paper, is performed after completing the production of the batch, and its main purpose is to test the quality of the manufactured units. The use of offline inspection is very popular, for example, in the printed circuit assembly industry: Offline manual X-ray inspection allows detection of solder defects in specific areas of the board. The technology employed by these systems has improved over time and many provide high-quality X-ray images. Some systems also provide software tools that help in detecting defects—especially for area array packages such as ball grid array (BGA). The importance of such nondestructive testing by X-ray continuously increases in the inspection of small to medium-size lots.

In accordance with the just-in-time philosophy, we assume that it is the producer's responsibility to ensure a delivery free from defects. In the yield distribution described above, the only way to guarantee zero defects is by employing a 100% inspection procedure. We assume that the items in a lot are inspected sequentially in the order they have been produced by an offline nondestructive and error-free procedure. Note that our model is closely related to lot-by-lot acceptance sampling for attributes and item-by-item sequential sampling, which sentence lots to acceptance or rejection based on a random sample; see Montgomery (2001). These two methods are usually not exhaustive sampling schemes, except for the case of very high cost of defectives, where 100% inspection may prove to be optimal; see Hald (1981) and Tagaras and Lee (1987). Many authors considered the inspection of goods produced by a process with constant failure rate, e.g., Grosfeld-Nir et al. (2000), Hassin (1984), He et al. (1996), Porteus (1986, 1990), and Raz et al. (2000). In contrast to our model, these papers assume that all units produced while the process is in control (out of control) are conforming (defective), and their main objective was to minimize the expected number of inspections required to detect the failure point. The detection of the failure point in such processes completely reveals the quality of all the items, as only the ones produced before the failure point are conforming; thus, it is equivalent to 100% inspection.

Our model assumes that the manufacturer faces a certain demand of a made-to-order product, and is restricted to producing a single batch whose size is a decision variable. This last restriction may be a result of a rush order, lead time considerations, and/or limited availability of production and inspection resources. Zipkin (2000, §9.4.8) considers such single-attempt production models. Defective and conforming units in excess of the demand are assumed to have no value. Thus, once the inspector detects a sufficient number of conforming units to fulfill the demand, inspection is stopped and all uninspected and defective units are scrapped. If inspection terminates while the number of conforming units is short of the demand, a per-unit shortage cost is incurred.

An optimal policy minimizes the expected total of production, inspection, and shortage costs. A policy is defined by the lot size to be produced and a dynamic rule that determines when to stop inspection. Referring to the second decision, note that the true state of the core process (production process) can only be inferred from the quality of the inspected units: Many defectives make it plausible that the process is out of control, and therefore it is economical to cease inspection even though the demand has not yet been satisfied. The decision on whether to cease inspection is taken with incomplete information about the core process. Thus, we formulate the inspection problem as a finite-time partially observable Markov decision process (POMDP). As will be demonstrated later, the resulting formulation is unusual in the sense that it consists of two "time-parameters:" the outstanding demand and the number of as yet uninspected units. The optimal batch size is determined based on the optimal solution to the inspection problem. Lot-sizing issues in the context of inspection and restoration were also addressed by Porteus (1986), Rosenblatt and Lee (1986a, b), Lee and Rosenblatt (1987), and Porteus (1990).

Sondik (1971) and Bertsekas (1976) showed that, using Bayes' rule, a POMDP can be converted into a Markov decision process (MDP): With each new observation, the probability distribution of the core process can be evaluated. These updated probabilities serve as state variables of the MDP. The main problem in solving the resulting MDP is its continuous state space. Sondik (1971) was the first to address and solve finite-time POMDPs by an exact "one-pass" algorithm. He showed that the optimal expected total reward as a function of the state variables is piecewise linear and concave. One of the main results when two actions are possible (Albright 1979, Bertsekas 1976, Lovejoy 1987, and White 1979) is that the optimal policy has a control limit threshold (CLT) structure, i.e., performs a certain action if and only if the state variable exceeds the CLT. In this paper, we show that all these properties prevail in our two "time-parameter" formulation. More precisely, for any number of yet uninspected units and any outstanding demand, the optimal inspection rule is of the CLT type, i.e., inspect the next unit if and only if the probability that the process was still in control exceeds the CLT. We also introduce a procedure for computing the various CLT values and prove structural properties that they satisfy.

In contrast to fully observed MDPs, analytically, POMDPs are usually hard to solve because of the prohibitively large size of the state space. Apart from a method developed by Grosfeld-Nir (1996) for a two-state POMDP with uniformly distributed observations, no analytical formula has yet been proposed to resolve the practical issue of computing the CLT. For algorithms, solution techniques, and bounds, see Lovejoy (1991a, b), White (1991), White and Scherer (1989), and Lauritzen and Nilsson (2001). For applications of POMDPs in a variety of areas, see Lane (1989), Monahan (1982), and references therein. Still, the major application is in machine maintenance/replacement and quality control. Calabrese (1995) used the POMDP methodology to solve an SPC problem.

In the context of random yield, our model is categorized as *nonrigid demand*. Yano and Lee (1995) reviewed the literature on random-yield production models with constant and random nonrigid and rigid demand. White (1970) considers a production setting similar to ours with a single run and a nonrigid constant demand, but without inspection costs. A recent paper by Scarf (2005) considers an infinite-horizon nonrigid demand model where the inventory manager is allowed to meet only a fraction of the demand in each period: Exercising this option may be profitable even at the risk of potential loss of good will if the sales price in a period is low relatively to the restocking cost.

This paper is organized as follows: In §2, we introduce notation, formulation, and preliminary results. In §3, we analyze the resulting POMDP and the structure of the optimal policy. Section 4 provides some numerical results. Section 5 concludes the paper with a discussion on the contribution of the results and possible extensions.

2. Problem Description

A manufacturer faces a nonrigid demand for D_0 units. A single batch is produced. Production involves a fixed setup cost α and a variable cost β, i.e., the cost of producing a batch of size N is $(\alpha+\beta N)$. Units are inspected, at a cost γ per unit, according to a first-come first-served (FCFS) policy, and only conforming units are delivered to the customer. When stopping inspection, if the number of good units is short of the demand, a cost s per unit shortage is incurred. Noninspected as well as defective units have no value. Clearly, $s>\gamma$ because otherwise inspection is never profitable. The objective is to minimize the expected total of production, inspection, and shortage costs. A policy is defined by the lot size and the stopping inspection rule.

2.1. Preliminaries and Notation: The Production and Inspection Processes

We consider a production process that can be in control or out of control. The true state of the process is unobservable and can only be inferred from the quality of the products. A unit produced while the process is in control (out of control) ends up conforming with probability θ_0 (θ_1); presumably, $\theta_0>\theta_1$. Let Z be the state of the production process while producing a specific unit, where $Z=0$ if the process is in control and $Z=1$ if it is out of control. Also, let $Y=0$ ($Y=1$) denote that this unit is conforming (defective). Then,

$$P(Y=0\mid Z=0)=\theta_0 \quad\text{and}\quad P(Y=0\mid Z=1)=\theta_1. \tag{1}$$

The process is probabilistically deteriorating with constant failure rate equal to $(1-r)$. That is, let Z denote the state while producing a specific unit, and let $\widehat{Z}$ be the state while producing the next unit in the batch. Then,

$$P(\widehat{Z}=0\mid Z=0)=r \quad\text{and}\quad P(\widehat{Z}=0\mid Z=1)=0. \tag{2}$$

Moreover, we assume that $\widehat{Z}$ is conditionally independent of Y given Z, i.e.,

$$P(\widehat{Z}=0\mid Y=y, Z=0)=P(\widehat{Z}=0\mid Z=0)=r. \tag{3}$$

After terminating the production of a batch, inspection begins. The inspection process perfectly reveals the quality of the inspected units. Based on these observations, we calculate the probability that the production process was in control while producing the candidate unit for inspection. This conditional probability is the key information necessary for solving the inspection problem.

Let x be the probability that the process was in control while producing the unit that is currently a candidate for inspection, i.e., $x=P(Z=0)$. Initially, when considering the inspection of the first unit in the batch $x=r$ because there is a chance of $1-r$ that the process gets out of control while producing the first unit. In general, the value of x is based on the quality of the inspected units. We refer to x

as the *information state*. It will become apparent that x is the state variable based on which decisions are made.

Let $P_x(A)$ represent the probability of event A given the information state x, and note that by definition $P_x(Z=0)=x$ and $P_x(A\mid Z=z)=P(A\mid Z=z)$ because, given the current state of the production process (i.e., $Z=0$ or $Z=1$), the information state x that is based on the quality of the inspected items is irrelevant. Let Y denote the quality of the unit that is currently a candidate for inspection. We define $p(x)=P_x(Y=0)$ and $q(x)=1-p(x)=P_x(Y=1)$. Note that x is the probability that the process was in control while producing the current unit, while $p(x)$ is the probability that this unit will end up conforming. By using the rules of conditional probability, we obtain $p(x)=P_x(Y=0)=\sum_{z=0,1}P_x(Y=0\mid Z=z)P_x(Z=z)=\sum_{z=0,1}P(Y=0\mid Z=z)P_x(Z=z)=x\theta_0+(1-x)\theta_1$. Thus, we get that $p(x)=\theta_1+(\theta_0-\theta_1)x$. The function $p(x)$ is linear and strictly increasing as $\theta_0>\theta_1$.

Note that Y is observable, whereas Z and $\widehat{Z}$ are unobservable. The next information state, i.e., the one attached to $\widehat{Z}$, is updated based on x and the observation Y as follows: $h_y(x)=P_x(\widehat{Z}=0\mid Y=y)$. The next lemma investigates the functions $h_0(x)$ and $h_1(x)$.

Lemma 1. (a)

$$h_0(x)=\frac{rx\theta_0}{x\theta_0+(1-x)\theta_1}=\frac{rx\theta_0}{p(x)},\qquad h_1(x)=\frac{rx(1-\theta_0)}{x(1-\theta_0)+(1-x)(1-\theta_1)}=\frac{rx(1-\theta_0)}{q(x)}. \tag{4}$$

(b) *Both $h_0(x)$ and $h_1(x)$ are continuous and strictly increasing functions, $h_0(x)$ is strictly concave, and $h_1(x)$ is strictly convex for $0<x<1$. $h_0(0)=h_1(0)=0$, $h_0(1)=h_1(1)=r$, and $h_1(x)<rx<h_0(x)$ for $0<x<1$. Moreover, the inverse functions $h_0^{-1}(\cdot)$ and $h_1^{-1}(\cdot)$ of $h_0(x)$ and $h_1(x)$, respectively, are well defined and are strictly increasing as well.*

(c) *If $r\theta_0\leqslant\theta_1$, then $0<h_0(x)<x$ for $x\in(0,1]$. If $r\theta_0>\theta_1$, let $x^*=(r\theta_0-\theta_1)/(\theta_0-\theta_1)$. Then, $h_0(x^*)=x^*$, $x^*<h_0(x)<x$ for $x>x^*$, and $h_0(x)>x$ for $0<x<x^*$. (See Figure 1.)*

Figure 1. The functions $h_0(x)$ and $h_1(x)$ for $r\theta_0>\theta_1$.

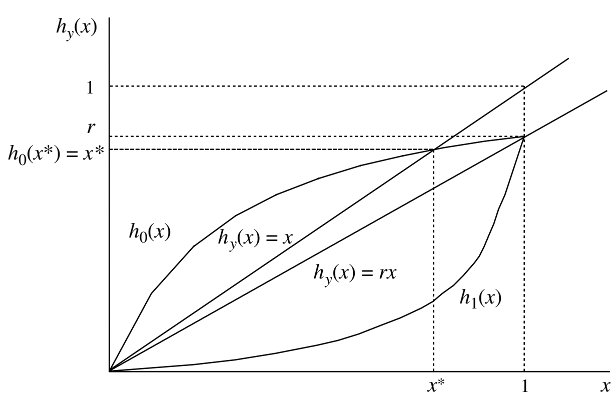

Proof. (a) (2) and (3) together with $P_x(\widehat{Z}=0\mid Y=y, Z=0)=P_x(\widehat{Z}=0\mid Z=0)=r$ imply that

$$\begin{aligned}h_y(x)&=P_x(\widehat{Z}=0\mid Y=y)\\&=P_x(\widehat{Z}=0\mid Y=y, Z=0)P_x(Z=0\mid Y=y)\\&=rP_x(Z=0\mid Y=y).\end{aligned}$$

Recall that $x=P_x(Z=0)$ and $P_x(Y=y\mid Z=z)=P(Y=y\mid Z=z)=\theta_z$ for $z=0,1$ and apply Bayes' formula to obtain

$$h_y(x)=\frac{rxP(Y=y\mid Z=0)}{xP(Y=y\mid Z=0)+(1-x)P(Y=y\mid Z=1)}.$$

(b) The proof follows directly from the definitions of $h_0(x)$ and $h_1(x)$ and verification of their first and second derivatives. Note also that the linear function rx coincides with the functions $h_0(x)$ and $h_1(x)$ at $x=0$ and $x=1$. The fact that $h_0(x)$ and $h_1(x)$ are strictly increasing implies that their inverse functions $h_0^{-1}(\cdot)$ and $h_1^{-1}(\cdot)$ are well defined, and they are strictly increasing.

(c) This part follows by comparing $h_0(x)$ to the function x, and by the fact that both functions are strictly increasing and $h_0(x)$ is strictly concave. □

According to Lemma 1(c), the information state can improve (increase) by observing the quality of a new unit only if (i) the unit is conforming; (ii) $r\theta_0>\theta_1$; and (iii) the information state x is sufficiently small, i.e., $x<x^*$. Otherwise, the information state deteriorates. This means that the information state can deteriorate even after revealing a good unit. This is in contrast to Porteus and Angelus (1997), which stated in the context of DSPC that the probability that the process is out of control jumps up if the inspected unit is found to be bad, and it dips down if it is found conforming. The reason for this counterintuitive result is the hazard rate $(1-r)$: If the hazard rate is relatively large, the identification of a conforming unit may be insufficient to counterbalance the deterioration caused by assignable causes.

2.2. The Expected Cost Functions

Given a candidate unit for inspection, the manufacturer should determine whether to continue or stop inspection, based on the quality of the inspected units. We denote these two actions by INSP (inspect) and STOP. We refer to this subproblem as the *inspection problem*. In the following, we formulate the inspection problem as a POMDP that depends on the information state x. For that sake, we define the following cost functions:

$V_{D_0}(n)$ for $n\geqslant 1$ is the expected total cost if a batch of size n is produced while the demand is D_0, and an optimal (FCFS) inspection policy is followed thereafter.

$G_{D,K}(x)$ is the remaining expected inspection and shortage costs for outstanding demand D, if the K, $K\geqslant 0$, last units in the batch are available for inspection, the information state is x, and an optimal inspection policy is followed.

$G^{\mathrm{INSP}}_{D,K}(x)$ is the remaining expected inspection and shortage costs for outstanding demand D if the K, $K \geqslant 1$, last units in the batch are available for inspection, the information state is x, the current action is INSP, and thereafter an optimal inspection policy is followed. We also define $G^{\mathrm{INSP}}_{D,0}(x) = sD$.

Note that $G^{\mathrm{INSP}}_{D,K}(x)$ and $G_{D,K}(x)$ are nonincreasing in K, because any additional uninspected unit can be discarded at no extra cost. Thus,

$$G^{\mathrm{INSP}}_{D,K}(x) \leqslant G^{\mathrm{INSP}}_{D,K-1}(x) \quad \text{for } K \geqslant 1, \tag{5a}$$

$$G_{D,K}(x) \leqslant G_{D,K-1}(x) \quad \text{for } K \geqslant 1. \tag{5b}$$

$G^{\mathrm{STOP}}_{D,K}(x) = sD$ is the shortage cost incurred when stopping inspection while the outstanding demand is D. Therefore, $G^{\mathrm{STOP}}_{D,K}(x)$ is independent of K, and $G_{D,K}(x) = \min\{G^{\mathrm{INSP}}_{D,K}(x), G^{\mathrm{STOP}}_{D,K}(x)\}$.

If $K \leqslant D$, then at most K units of demand can be supplied and a shortage cost for at least $(D-K)$ units must be incurred by any policy. Thus,

$$G^{\mathrm{INSP}}_{D,K}(x) = G^{\mathrm{INSP}}_{K,K}(x) + s(D-K) \quad \text{for } 0 \leqslant K \leqslant D. \tag{6a}$$

Recall that $s > \gamma$, which implies the following inequality:

$$G^{\mathrm{INSP}}_{D+1,K}(x) \leqslant G^{\mathrm{INSP}}_{D,K}(x) + s \quad \text{for } K > D \geqslant 0, \tag{6b}$$

because the left-hand side represents the expected remaining cost of an optimal policy for demand $D+1$ and K uninspected units given that the first unit is inspected. The right-hand side represents the expected remaining cost of a feasible policy in which the K uninspected units are used to satisfy at most D units of demand, and a shortage cost for at least one unit is incurred. In the sequel, we assume that if $G_{D,K}(x) = G^{\mathrm{INSP}}_{D,K}(x) = sD$, i.e., both the STOP and the INSP policies have the same expected cost, then STOP is used. Technically, we need this assumption to be able to assign to each information state a single optimal policy.

2.3. The Production Problem

The production problem is to determine the elements of the set

$$N_{D_0} = \arg\min\{V_{D_0}(n)\colon n \geqslant 1\},$$
$$\text{where } V_{D_0}(n) = \alpha + \beta n + G_{D_0,n}(r). \tag{7}$$

That is, the set N_{D_0} is the set of optimal production lots for demand D_0. Thus, given $G_{D_0,n}(r)$ for $n \geqslant 1$, the optimal lot size can be calculated via a search over n in (7). Note that the set-up cost α has no effect on the optimal policy. Its effect is on the strategic decision on whether the contract is profitable. In this paper, we design the optimal policy, assuming that a batch of size $n \geqslant 1$ is produced. Each inspected unit costs at least $\beta+\gamma$. If $s \leqslant \beta+\gamma$, then it is optimal to produce the minimum lot size, i.e., $N_{D_0} = \{1\}$. Thus, a necessary condition for producing a batch of size $n > 1$ is $s > \beta+\gamma$. In the following, we develop two upper bounds on the optimal lot size by considering the expected cost of two specific policies, namely, producing a batch of size 1, and producing a batch of size D_0. The existence of the upper bounds shows that theoretically the search for the optimal lot size is finite in the problem's parameters. These bounds were helpful in searching for the optimal lot sizes in §4. Further properties on $V_{D_0}(n)$ as a function of the lot size n should be explored in future research to enable a more efficient search on the optimal lot size.

The derivation of the upper bounds is based on the fact that for $n \in N_{D_0}$, $V_{D_0}(n) \leqslant \min\{V_{D_0}(1), V_{D_0}(D_0)\}$. Note also that for any lot size, $V_{D_0}(n) \geqslant \alpha + \beta n + \gamma D_0$ because γD_0 is a lower bound on the total of inspection and shortage costs. We consider each of the two upper bounds mentioned above separately: If a single unit is produced, then $V_{D_0}(1) \leqslant \alpha + \beta + sD_0$. Thus, any optimal lot size $n \in N_{D_0}$ satisfies $n \leqslant 1 + ((s-\gamma)D_0)/\beta$.

The other upper bound on N_{D_0} is obtained by considering $V_{D_0}(D_0)$: The minimal expected inspection and shortage costs when a lot of size D_0 is produced are bounded from above by the expected cost of the policy of inspecting the whole batch. Let D_0^G denote the random variable of the number of good units in a batch of size D_0. Then, $V_{D_0}(D_0) \leqslant E[\alpha + \beta D_0 + \gamma D_0 + s(D_0 - D_0^G)] = \alpha + (\beta + \gamma + s)D_0 - sE(D_0^G)$. To compute $E(D_0^G)$, we define an additional random variable $B(D_0)$ that represents the last unit in a batch of size D_0 where the process was still in control. Clearly, $0 \leqslant B(D_0) \leqslant D_0$. The random variable $B(D_0)$ is distributed according to the truncated geometric probability distribution with a hazard rate of $(1-r)$, thus $P(B(D_0) \geqslant m) = r^m$ for $m = 0, \ldots, D_0$, $E(B(D_0)) = \sum_{m=1}^{D_0} r^m = r(1-r^{D_0})/(1-r)$, and

$$\begin{aligned} E(D_0^G) &= E_{B(D_0)}[E_{D_0^G}(D_0^G/B(D_0))] \\ &= E_{B(D_0)}(\theta_0 B(D_0) + \theta_1(D_0 - B(D_0)) \\ &= \frac{\theta_1 D_0 + (\theta_0-\theta_1) r(1-r^{D_0})}{1-r}. \end{aligned}$$

Thus, for $n \in N_{D_0}$, $\alpha + \beta n + \gamma D_0 \leqslant V_{D_0}(n) \leqslant V_{D_0}(D_0)$, and therefore,

$$n \leqslant \frac{[\beta + s(1-\theta_1)]D_0}{\beta} - \frac{s(\theta_0-\theta_1)r(1-r^{D_0})}{\beta(1-r)}.$$

The minimum of the two bounds provides an upper bound on the optimal batch size.

In the rest of the paper, we solve the inspection problem.

2.4. The Inspection Problem

The objective here is to find an optimal stopping-inspection rule given the outstanding demand D and $K \geqslant 1$ uninspected units. The action INSP is optimal if and only if

$G^{\text{INSP}}_{D,K}(x) \leqslant G^{\text{STOP}}_{D,K}(x) = s\cdot D$. To simplify the analysis, we define the following:

$$\Delta_{D,K}(x) = G^{\text{INSP}}_{D,K}(x) - G^{\text{STOP}}_{D,K}(x) = G^{\text{INSP}}_{D,K}(x) - s\cdot D \quad \text{for } K \geqslant 1 \text{ and } D \geqslant 1. \tag{8}$$

In particular, $\Delta_{D,0}(x) = 0$ and $\Delta_{0,K}(x) = 0$ for $K \geqslant 0$ and $D \geqslant 0$.

Note that $\Delta_{D,K}(x)$ represents the expected savings if inspection is stopped. Thus, for given D and K, STOP is an optimal policy if and only if $\Delta_{D,K}(x) \geqslant 0$. A recursive system of the Δ functions enables us to use a constant reference value of 0 instead of $s\cdot D$ if the recursive system of the G functions was used. We rewrite (5a), (6a), and (6b) as follows:

$$\Delta_{D,K}(x) \leqslant \Delta_{D,K-1}(x) \quad \text{for } D \geqslant 0 \text{ and } K \geqslant 1, \tag{9}$$

$$\Delta_{D,K}(x) = \Delta_{K,K}(x) \quad \text{for } 0 \leqslant K \leqslant D, \tag{10a}$$

$$\Delta_{D+1,K}(x) \leqslant \Delta_{D,K}(x) \quad \text{for } K > D \geqslant 0. \tag{10b}$$

We also let $\Delta^{\text{OPT}}_{D,K}(x) = G_{D,K}(x) - s\cdot D$. Note that $G_{D,K}(x) \leqslant s\cdot D$ implies that $\Delta^{\text{OPT}}_{D,K}(x) \leqslant 0$. Thus, $\Delta^{\text{OPT}}_{D,K}(x) = 0$ if and only if the policy STOP is optimal. Otherwise, $\Delta^{\text{OPT}}_{D,K}(x) < 0$. If $\Delta^{\text{OPT}}_{D,K}(x) < 0$, then $|\Delta^{\text{OPT}}_{D,K}(x)|$ represents the expected savings that result from continuing inspection relative to the policy that stops inspection.

The following recursive equations define the inspection problem:

$$\begin{aligned}
\Delta^{\text{OPT}}_{D,K}(x) &\equiv G_{D,K}(x) - s\cdot D \\
&= \min\{0, \Delta_{D,K}(x)\} \quad \text{for } K \geqslant 0 \text{ and } D \geqslant 0, \\
\Delta_{D,K}(x) &= G^{\text{INSP}}_{D,K}(x) - s\cdot D \\
&= \gamma + p(x)G_{D-1,K-1}(h_0(x)) \\
&\quad + q(x)G_{D,K-1}(h_1(x)) - s\cdot D \\
&= \gamma - s\cdot p(x) + p(x)[G_{D-1,K-1}(h_0(x)) - s(D-1)] \\
&\quad + q(x)[G_{D,K-1}(h_1(x)) - s\cdot D] \\
&= \gamma - s\cdot p(x) + p(x)\Delta^{\text{OPT}}_{D-1,K-1}(h_0(x)) \\
&\quad + q(x)\Delta^{\text{OPT}}_{D,K-1}(h_1(x)),
\end{aligned} \tag{11}$$

where the last equation holds for $D \geqslant 1$ and $K \geqslant 1$. The boundary conditions are

$$\Delta_{0,K}(x) = 0 \quad \text{and} \quad \Delta_{D,0}(x) = 0. \tag{12}$$

In the context of POMDPs, K—the number of yet uninspected units—is usually interpreted as a time parameter—the number of remaining periods to the end of the horizon. This interpretation is also valid here: Assume that the duration of inspecting a unit is one time period—entailing a cost γ. However, in contrast to traditional POMDPs, we also need to consider the outstanding demand D. It is as if, in the context of POMDP's, there are two "time parameters" (K and D) with the problem terminating as soon as we run out of one of the two "time resources." Furthermore, one of the "time parameters," the number of uninspected items, is policy dependent. To the best of our knowledge, this is the first time that such a POMDP is explored. In the next subsection, we discuss a certain type of functions that turns out to be helpful in our analysis.

2.5. PLDC Functions

A continuous function is said to be PLDC if it is piecewise linear, strictly decreasing, and concave. If $\gamma/s > \theta_1$, $\Delta_{D,K}(x)$ for $\min\{D,K\} \geqslant 1$ is PLDC, and therefore is associated with a single root $L_{D,K}$ that satisfies $\Delta_{D,K}(L_{D,K}) = 0$. If $L_{D,K} < r$, $L_{D,K}$ is said to be CLT. In this case, the optimal action is INSP if and only if the information state x is greater than $L_{D,K}$. The next lemma, which is proved in the appendix, is helpful in the proof of Theorem 1.

Lemma 2. *Suppose that the continuous function $f(x)$ is piecewise linear decreasing and concave for $x > 0$. Moreover, there exists $\tau_1 > 0$ such that $f(x) \equiv 0$ for $x \in [0, \tau_1]$. Then, the function*

(i) *$p(x)f(h_0(x))$ is piecewise linear decreasing and concave,*

(ii) *$q(x)f(h_1(x))$ is piecewise linear decreasing and concave, and*

(iii) *$k(x) = g(x) + p(x)f(h_0(x)) + q(x)f(h_1(x))$ is PLDC for any function $g(x)$ that is PLDC.*

Theorem 1. *If $\gamma/s > \theta_1$, then $\Delta_{D,K}(x)$ for $D \geqslant 1$, $K \geqslant 1$, is PLDC with a unique positive root.*

Proof. The proof is by induction on K. For any $D \geqslant 1$ and $K = 1$, according to (11) and (12), $\Delta_{D,1}(x) = \gamma - s\cdot p(x) = \gamma - s\cdot\theta_1 - s(\theta_0 - \theta_1)x$, which is PLDC. The condition $\gamma/s > \theta_1$ implies that the root of $\Delta_{D,1}(x)$, namely, $L_{D,1} = (\gamma/s - \theta_1)/(\theta_0 - \theta_1) > 0$.

Suppose that for any D, $\Delta_{D,k}(x)$, $k \leqslant K-1$, is PLDC with a unique positive root. We show that also $\Delta_{D,K}(x)$ is PLDC. By hypothesis, $\Delta_{D-1,K-1}(x)$ and $\Delta_{D,K-1}(x)$ are PLDC, each associated with a positive root; therefore, $\Delta^{\text{OPT}}_{D-1,K-1}(x)$ and $\Delta^{\text{OPT}}_{D,K-1}(x)$ are piecewise linear decreasing and concave and are equal to 0 near 0. Consequently, there exists $\tau > 0$ such that in $[0,\tau)$, $\Delta^{\text{OPT}}_{D-1,K-1}(h_0(x)) \equiv \Delta^{\text{OPT}}_{D,K-1}(h_1(x)) \equiv 0$. Thus, in this interval $\Delta_{D,K}(x) = \gamma - s\cdot p(x)$ (see (11)). By Lemma 2, the function $\Delta_{D,K}(x)$ is PLDC. Because $\Delta_{D,K}(0) = \gamma - s\cdot p(0) = \gamma - s\theta_1 > 0$, $\Delta_{D,K}(x)$ has a unique positive root. □

Figures 2a and 2b demonstrate the form of the functions $\Delta_{D,K}(x)$ and $\Delta^{\text{OPT}}_{D,K}(x)$ if $\gamma/s > \theta_1$ and $r = 0.9$; because $L_{D,K}$—the root of $\Delta_{D,K}(x)$—is less than r, it serves as the CLT.

Figure 2. (a) The function $\Delta_{D,K}(x)$. (b) The respective function $\Delta^{\mathrm{OPT}}_{D,K}(x)$.

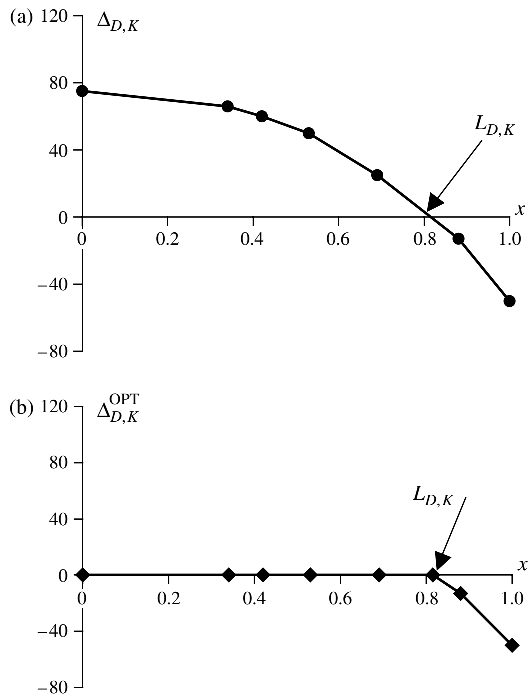

3. The Solution to the Inspection Problem

In §3.1, we identify sets of parameters for which the optimal solution is degenerate (always STOP or always INSP), independently of D or K or the information state x. In §3.2, we analyze the inspection problem for the nondegenerate cases. We will use the following notation: $\mathscr{L}_1 = (\gamma/s-\theta_1)/(\theta_0-\theta_1)$ and $\Omega_1(x) = \gamma - s\cdot p(x)$. Note that $\mathscr{L}_1$ is the unique root of $\Omega_1(x)$, i.e., $\Omega_1(\mathscr{L}_1) = 0$.

3.1. When Is the Optimal Inspection Rule Static?

Here we identify combinations of the parameters for which the optimal inspection policy is static, i.e., it does not depend on D and K. As is demonstrated below, the critical factor is the ratio between the inspection cost and the shortage cost, that is, the size of γ/s.

Theorem 2. *Consider the inspection problem with $D\geqslant 1$, $K\geqslant 1$. Then,*

(i) *If $\gamma/s \geqslant p(r) = \theta_1 + (\theta_0-\theta_1)r$, the optimal policy is STOP, i.e., $\Delta^{\mathrm{OPT}}_{D,K}(x)\equiv 0$ for any $x\in[0,r]$.*

(ii) *If $\gamma/s\leqslant\theta_1$, then the optimal policy is INSP, i.e., $\Delta^{\mathrm{OPT}}_{D,K}(x) = \Delta_{D,K}(x) < 0$ for $x>0$.*

(iii) *If $\max\{\theta_1, r\cdot\theta_0\}\leqslant \gamma/s\leqslant p(r)$, the optimal policy is INSP if and only if $x>\mathscr{L}_1$, i.e., $L_{D,K}=\mathscr{L}_1$, independently of D and K. Moreover, $\Delta_{D,K}(x) = \Omega_1(x)$ for $x\leqslant h_0^{-1}(\mathscr{L}_1)$.*

Proof. (i) In view of (11) and (12), $\gamma/s\geqslant p(r)$ implies that $\Delta_{D,1}(x) = \gamma - s\cdot p(x)\geqslant 0$, and therefore $\Delta^{\mathrm{OPT}}_{D,1}(x)=0$ for $x\in[0,r]$. Assume by induction that $\Delta^{\mathrm{OPT}}_{D,k}(x)=0$ for $k\leqslant K-1$, $D\geqslant 1$, and any $x\in[0,r]$. We will prove that $\Delta^{\mathrm{OPT}}_{D,K}(x)=0$ for $x\in[0,r]$. According to the inductive assumption and (11),

$$\Delta_{D,K}(x) = \gamma - s\cdot p(x)\geqslant \gamma - s\cdot p(r)\geqslant 0.$$

Thus, $\Delta^{\mathrm{OPT}}_{D,K}(x)=0$ and STOP is optimal for any triplet $D\geqslant 1$, $K\geqslant 1$, and $x\in[0,r]$.

(ii) Note that $\gamma/s\leqslant\theta_1$ implies $\Delta_{D,1}(x) = \gamma - s\cdot p(x) < 0$ (because $p(x)>\theta_1$), and therefore $\Delta^{\mathrm{OPT}}_{D,1}(x)<0$ for $x>0$. Assume by induction that $\Delta_{D,k}(x)<0$ for $x>0$, $k\leqslant K-1$, and $D\geqslant 1$. We will prove that $\Delta^{\mathrm{OPT}}_{D,K}(x)<0$ for $x>0$. According to the inductive assumption and (11),

$$\Delta_{D,K}(x) < \gamma - sp(x) < \gamma - s\theta_1\leqslant 0.$$

Thus, $\Delta_{D,K}(x) = \Delta^{\mathrm{OPT}}_{D,K}(x)<0$ for $x>0$, i.e., the optimal policy is INSP.

(iii) First observe that $\max\{\theta_1, r\theta_0\}\leqslant\gamma/s\leqslant p(r)$ implies $h_0(\mathscr{L}_1)<\mathscr{L}_1$, or equivalently, $\mathscr{L}_1<h_0^{-1}(\mathscr{L}_1)$, because

$$\mathscr{L}_1 = \frac{\gamma/s-\theta_1}{\theta_0-\theta_1} > \frac{r\theta_0-\theta_1}{\theta_0-\theta_1} = x^*;$$

see Lemma 1(c).

For $K=1$, $\Delta_{D,1}(x) = \gamma - s\cdot p(x) = \Omega_1(x)$ with $L_{D,1}=\mathscr{L}_1$. Note that $\Delta_{D,1}(0) = \gamma - s\cdot\theta_1 > 0$ and $\Delta_{D,1}(r) = \gamma - s\cdot p(r) < 0$; thus $0<\mathscr{L}_1<r$. Assume that $L_{D,k}=\mathscr{L}_1$ for $k\leqslant K-1$. We will prove that $L_{D,K}=\mathscr{L}_1$. Assume that $x\leqslant h_0^{-1}(\mathscr{L}_1)$. Then, $h_1(x)<h_0(x)\leqslant\mathscr{L}_1$. Therefore, $\Delta^{\mathrm{OPT}}_{D-1,K-1}(h_0(x)) = \Delta^{\mathrm{OPT}}_{D,K-1}(h_1(x)) = 0$ and $\Delta_{D,K}(x) = \gamma - s\cdot p(x) = \Omega_1(x)$. Thus, $\Delta_{D,K}(\mathscr{L}_1)=0$ and $L_{D,K}=\mathscr{L}_1$. □

Note that $\mathscr{L}_1$ satisfies $p(\mathscr{L}_1)=\gamma/s$, and therefore $x>\mathscr{L}_1$ is equivalent to $p(x)>\gamma/s$. Thus, under the condition of Theorem 2(iii), inspection continues if and only if the probability that the first uninspected unit is conforming is higher than γ/s.

3.2. The Case of the Dynamic Optimal Inspection Rule

In view of Theorem 2, the interesting set of parameters to focus on is $\theta_1<\gamma/s<r\theta_0$. Theorem 1 implies that the functions $\Delta_{D,K}(x)$ for $D\geqslant 1$ and $K\geqslant 1$ are PLDC, and therefore their roots $L_{D,K}$ are well defined. The following observation is helpful in the analysis.

Observation 1. Note that $\theta_1<\gamma/s<r\theta_0$ implies $0<\mathscr{L}_1<x^*$ to conclude that $h_0^{-1}(\mathscr{L}_1)<\mathscr{L}_1$.

In light of Observation 1, and the fact that $h_1(x)<rx<x$ for any $x\in(0,1)$, we obtain that for problems with $\gamma/s<r\theta_0$, the following inequalities hold:

$$h_0^{-1}(\mathscr{L}_1)<\mathscr{L}_1<h_1^{-1}(\mathscr{L}_1).$$

Theorem 3 analyzes the case where $D \geqslant K$. Note that if at some point the demand was at least as large as the number of uninspected units, i.e., $D \geqslant K$, then also in the future $D \geqslant K$; because the number of uninspected units decreases by a faster rate than the outstanding demand. In Theorem 3 we prove that if $D \geqslant K$, then the optimal inspection rule depends only on K and the information state x, but it is independent of the demand D. The proofs of the next two theorems are in the appendix.

Theorem 3. *Suppose that $\theta_1 < \gamma/s < r\theta_0$ and $K \geqslant 1$.*

(i) *For $D \geqslant K$, the functions $\Delta_{D,K}(x)$ are independent of D, so we denote $\Omega_K(x) = \Delta_{D,K}(x)$ and let $\mathscr{L}_K = L_{D,K}$ be its unique root. For $D \geqslant K$, the INSP policy is optimal if and only if $x > \mathscr{L}_K$.*

(ii) *The functions $\Omega_K(x)$ consist of at most $2^K - 1$ linear segments.*

(iii) *Assume that $K \geqslant 2$. For $x \leqslant h_0^{-1}(\mathscr{L}_{K-1})$, $\Omega_K(x) = \cdots = \Omega_1(x) = \gamma - s \cdot p(x)$, and for $x > h_0^{-1}(\mathscr{L}_{K-1})$, $\Omega_K(x) < \Omega_{K-1}(x)$.*

(iv) $0 < \mathscr{L}_{K+1} < \mathscr{L}_K < \cdots < \mathscr{L}_1 < x^*$. *Moreover, $\mathscr{L}_{K+1} \in (h_0^{-1}(\mathscr{L}_K), \mathscr{L}_K)$.*

(v) *There exists a constant $\mathscr{L}^*$, $\mathscr{L}^* \geqslant 0$, such that $\lim_{K\to\infty} \mathscr{L}_K = \mathscr{L}^*$.*

In the sequel, we let $\mathscr{L}_K$ be the single root of the function $\Omega_K(x) = \Delta_{K,K}(x)$ for $K \geqslant 1$. This definition is equivalent to the one in Theorem 3(i). We also let $\mathscr{L}^* = \lim_{K\to\infty} \mathscr{L}_K$ as defined in Theorem 3(v). Theorem 4 completes Theorem 3 by considering the case $D \leqslant K$. In this range, the optimal inspection rule depends on D and K.

Theorem 4. *Suppose that $\theta_1 < \gamma/s < r\theta_0$ and $K \geqslant D$.*

(i) *For $D = 1$, $L_{1,K} = \mathscr{L}_1$. Moreover, $\Delta_{1,K}(x) = \Omega_1(x)$ for $x \in [0, h_1^{-1}(\mathscr{L}_1)]$.*

(ii) *For a fixed $D(K)$, the sequence $L_{D,K}$ is nonincreasing in $K(D)$ and $\mathscr{L}_K \leqslant L_{D,K} \leqslant \mathscr{L}_D$. Moreover, $\Delta_{D,K}(x) = \Omega_1(x)$ for $x \leqslant h_0^{-1}(\mathscr{L}_{K-1})$.*

4. Numerical Results

To solve the inspection problem numerically, we computed recursively in D and K approximations for the functions $\Delta_{D,K}(x)$, $K \leqslant 14$, $D \leqslant 10$, $x \in \{0, 0.001, 0.002, \ldots, r\}$. The running time for each problem (including the search for the optimal production lot) was less than two seconds. We also tried a finer grid for the x-values, but this proved unnecessary.

4.1. The Inspection Problem

As follows from Theorem 2, the interesting set of parameters for the inspection problem satisfies $0 < \theta_1 < \gamma/s < r\theta_0$. The CLTs of the inspection problems depend on the cost parameters only via the ratio γ/s, and they do not depend on the production cost parameters, namely, α and β. The variable production cost β is used only in determining the optimal lot size. In our numerical test, we consider the following basic set of parameters, which we believe to be realistic:

r	θ_0	θ_1	γ	s
0.98	0.90	0.40	2	3

For this set of parameters, we obtain the CLT values $L_{D,K}$ that are shown in Table 1.

Table 1 reveals, as proved in Theorems 3(iv) and 4(ii), that the CLTs are nonincreasing in $K(D)$ for fixed $D(K)$. Table 1 shows that for fixed K, the CLT values remain constant for $D \geqslant K$, as proved in Theorem 3(i). Moreover, by Table 1, for fixed D, the CLT values are strictly decreasing in K for $K \leqslant D$, see Theorem 3(iv). Table 1, as well as additional numerical examples that we have explored, suggest that for fixed D, the CLT values remain constant as long as $K > D$. We leave the further investigation of this property for future research.

Table 2 explores the dependency of the CLT values on γ/s. Note that if $\gamma/s < \theta_1 = 0.4$, the optimal policy is always to inspect, i.e., CLT $= 0$; see Theorem 2(ii). If $\gamma/s > p(r) = 0.89$, the optimal policy is always to stop. That is, CLT $= r$; see Theorem 2(i). Also, if $\max\{\theta_1, r \cdot \theta_0\} \leqslant \gamma/s \leqslant p(r)$, i.e., $0.882 \leqslant \gamma/s \leqslant 0.89$, the optimal policy is INSP if and only if $x > \mathscr{L}_1 = (\gamma/s - \theta_1)/(\theta_0 - \theta_1)$; see Theorem 2(iii). Thus, in this last interval, the CLT is strictly increasing in γ/s. Our numerical study shows that also for $0.4 \leqslant \gamma/s \leqslant 0.882$, the CLT values increase in γ/s.

Table 1. The CLTs as a function of D and K for the basic set of parameters.

	K													
D	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	**0.53**	0.53	0.53	0.53	0.53	0.53	0.53	0.53	0.53	0.53	0.53	0.53	0.53	0.53
2	0.53	**0.46**	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46
3	0.53	0.46	**0.41**	0.41	0.41	0.41	0.41	0.41	0.41	0.41	0.41	0.41	0.41	0.41
4	0.53	0.46	0.41	**0.36**	0.36	0.36	0.36	0.36	0.36	0.36	0.36	0.36	0.36	0.36
5	0.53	0.46	0.41	0.36	**0.33**	0.33	0.33	0.33	0.33	0.33	0.33	0.33	0.33	0.33
6	0.53	0.46	0.41	0.36	0.33	**0.30**	0.30	0.30	0.30	0.30	0.30	0.30	0.30	0.30
7	0.53	0.46	0.41	0.36	0.33	0.30	**0.28**	0.28	0.28	0.28	0.28	0.28	0.28	0.28
8	0.53	0.46	0.41	0.36	0.33	0.30	0.28	**0.26**	0.26	0.26	0.26	0.26	0.26	0.26
9	0.53	0.46	0.41	0.36	0.33	0.30	0.28	0.26	**0.25**	0.24	0.24	0.24	0.24	0.24
10	0.53	0.46	0.41	0.36	0.33	0.30	0.28	0.26	0.25	**0.23**	0.23	0.23	0.23	0.23

Table 2. The CLTs for several values of γ/s for our basic set of parameters, $K = 9$ and $1 \leqslant D \leqslant 10$.

K	9	9	9	9	9
D	$\gamma/s = 0.35$	$\gamma/s = 0.50$	$\gamma/s = 0.67$	$\gamma/s = 0.80$	$\gamma/s = 0.89$
1	0	0.201	0.534	0.801	0.98
2	0	0.158	0.463	0.756	0.98
3	0	0.129	0.407	0.715	0.98
4	0	0.109	0.363	0.679	0.98
5	0	0.095	0.329	0.648	0.98
6	0	0.085	0.301	0.622	0.98
7	0	0.076	0.279	0.600	0.98
8	0	0.069	0.259	0.580	0.98
9	0	0.066	0.248	0.565	0.98
10	0	0.066	0.248	0.565	0.98

4.2. The Production Problem

Next, we explore in our numerical test the dependency of the optimal production lot size on β (all other parameters of the basic set remain fixed). We assume in our runs that $\alpha = 0$, and we let the optimal lot size be 0 if production is not profitable, that is, if it is cheaper to pay the shortage cost for the whole demand than producing a lot. A necessary condition for production to be profitable is $\beta + \gamma < s$. Thus, it implies that for the choice of our basic parameters, $\beta < 1$. We should note that some offline inspection systems are very costly relative to the production cost in view of the special equipment that is needed. For example, X-ray inspection of assembly printed circuits mentioned in the introduction may necessitate the use of complex machine vision systems, which are expensive in comparison to the low unit production cost. Therefore, in our computational test we allow γ to be 2.5, and even 10 times higher than the unit production cost.

In Table 3 we search for the optimal lot size over the range defined by the minimum of the upper bounds developed in §2.3. These bounds get tighter as β is larger or D_0 is smaller. For example, for $\beta = 0.8$ (0.2), the upper bound on the optimal lot size varies between 2 to 13 (6 to 51) as D_0 is increased from 1 to 10. Our limited computational study seems to support the conjecture that the function $V_{D_0}(n)$ is unimodal (with a single local minimum) and possibly even convex in n. As mentioned in §2.3, we believe that further research is needed to explore properties of the function $V_{D_0}(n)$ to enhance the computation of the optimal lot size.

Table 3. Optimal expected costs and optimal production lots for the problem with basic parameters, various values of β, and $\alpha = 0$.

	$\beta = 0.2$		$\beta = 0.4$		$\beta = 0.6$		$\beta = 0.8$	
D_0	V_{D_0}	$n^*_{D_0}$	V_{D_0}	$n^*_{D_0}$	V_{D_0}	$n^*_{D_0}$	V_{D_0}	$n^*_{D_0}$
1	2.53	1	2.73	1	2.93	1	3	0
2	5.09	2	5.49	2	5.89	2	6	0
3	7.68	3	8.28	3	8.88	3	9	0
4	10.29	4	11.09	4	11.88	3	12	0
5	12.92	5	13.92	5	14.88	3	15	0
6	15.55	7	16.77	6	17.88	3	18	0
7	18.2	8	19.64	7	20.88	3	21	0
8	20.86	9	22.53	8	23.88	3	24	0
9	23.55	10	25.42	9	26.88	3	27	0
10	26.24	11	28.33	10	29.88	3	30	0

We let $n^*_{D_0} \geqslant 0$ be an optimal lot size for demand D_0. The optimal expected cost for demand D_0 presented in Table 3 is the minimum between the shortage cost sD_0 and the total of the production cost $(\alpha + \beta n^*_{D_0})$ plus the expected inspection cost, given that initially $n^*_{D_0}$ units are available for inspection.

Table 3 reveals that for high values of β, $\beta = 0.8$, it is still cheaper (as for $\beta = 1$) to pay the shortage cost for the whole demand rather than producing a lot. This is intuitive, as the strategy of producing a lot necessitates a cost of at least $\beta + \gamma \geqslant 2.8$ for each unit that is inspected, in addition to the extra cost incurred if the inspected unit turns out to be defective. The policy of not producing, and paying the shortage cost of $s = 3$ per unit, is risk free. When β is reduced to 0.6, the optimal strategy is to produce lots of a size that do not exceed 3. Indeed, for $D_0 \leqslant 3$, the optimal lot size equals the demand, but for $D_0 > 3$ the optimal lot is of size three units. In this example, the variable production cost is still too large, so the producer is better off not taking the risk of producing a lot of a size that is larger than the demand. Moreover, the relatively high unit production cost does not make it profitable to produce lots of size greater than 3 as the chance that the production system is still in control while producing the fourth unit is not high enough to offset the cost incurred if this unit is not inspected, or if it is inspected and found nonconforming. When β is decreased to $\beta = 0.4$, the optimal lot size increases so that for $D_0 \leqslant 10$, we produce a lot of a size that is as large as the demand. When β is further reduced to $\beta = 0.2$ and demands are in the range $D_0 \leqslant 5$, the optimal lot size is equal to the demand. However, for $5 < D_0 \leqslant 10$, the optimal lot size is slightly above the demand. The explanation for this behavior is that for sufficiently small demands, the problem parameters are such that the chance of getting a lot where all of its first five units are conforming is reasonably high, to make the expected cost of producing units beyond the demand not profitable. However, for $5 < D_0 \leqslant 10$, the optimal lot size is greater than the demand by one unit, meaning that the risk of obtaining at least one defective unit among the first 6–10 units in the lot is high enough to justify the production of one extra unit. From additional runs that we performed, it is clear that as β gets smaller the marginal cost of producing units beyond the demand that may not be inspected is low enough in comparison to the marginal expected savings on the shortage cost if these units are inspected.

The focus of this paper is on the inspection problem. However, Table 3 seems to suggest two additional properties regarding the production problem: (i) the optimal lot size never decreases in the demand given that all other parameters are fixed; and (ii) the expected total cost is not only increasing in the demand, but also the marginal cost (per unit demand) is nondecreasing in the demand. We leave these properties as open questions.

5. Conclusions

We consider a production and inspection problem that deals with a lot-sizing and inspection policy. The manufacturer is contracted to a zero-defect, nonrigid delivery. The objective is to identify a policy that minimizes the expected production, inspection, and shortage costs. The optimal lot size can be found by a bounded search once the inspection problem is solved to optimality. The inspection problem is formulated as an unusual POMDP with two "time parameters:" the outstanding demand and the number of uninspected items. The optimal policy for the inspection problem is proved to be a simple CLT structure where the decision whether to stop or continue inspection depends on a single information state that is updated dynamically with the acquisition of new observations. The CLT divides the range of possible information states into two disjoint intervals so that in one of the intervals the optimal policy is to stop inspection, and in its complement it is preferable to continue inspection. The CLT values can be computed recursively in the demand and the number of uninspected items.

It is interesting to note that Porteus and Angelus (1997) observed that in online inspection the remaining lot size affects the optimal DSPC rule. More specifically, towards the end of the run more negative evidence is needed to warrant restoration of the production process; see also Crowder (1992). Our results regarding offline 100% inspection support these findings in the sense that, in general, the optimal inspection rule is not static: The operator should be more conservative and careful (i.e., more negative evidence is required) when stopping inspection and paying the shortage cost the larger is the number of yet uninspected items or the larger is the remaining demand.

In particular, Theorems 3(iv) and 4(ii) can be interpreted as "monotonicity of the finite horizon control limits." We note that this is in contrast to Grosfeld-Nir (1996), which showed that the finite-horizon control limits are not necessarily monotone. This result accentuates the fact that "two time parameter" POMDPs may differ from the traditional case.

Our problem is meaningful only for a finite demand D and a finite number K of uninspected items; thus, it corresponds to the finite-horizon case. Furthermore, the horizon length is policy dependent. The structural properties of the various CLT values that we prove reduce the computational effort involved in the explicit derivation of the expected cost functions. However, in general, the calculation of the optimal policy is possible only for small values of K because the expected cost PLDC functions consist of a number of segments that grows exponentially with K. These computational aspects suggest the need for application of numerical approximation methods that reduce the uncountably infinite state space that the information state can assume to a finite grid of points. Some methods adjust the grid dynamically as the algorithm progresses, whereas others keep it fixed. Lovejoy (1991b) proposes an easily manageable fixed-grid method to develop upper and lower value function bounds, which help in generating policies that are nearly optimal. That paper also provides references for a number of alternative approximation techniques that may be helpful in approximating the PLDC expected cost function for sufficiently large D and K by a function that is differentiable almost everywhere. The root of such a function may serve as an approximation to L^*—the limiting CLT value when D and K increase. In our numerical test we show that for reasonable demand levels, a simple approximation of the expected cost functions on a grid provides good quality solutions in a few seconds.

An interesting question for future research involves the generalization of the production process to one with increasing failure rate. Banerjee and Rahim (1988) and Rahim and Banerjee (1993) considered such DSPC models. In the offline 100%-inspection case the optimal inspection policy will not only depend on D and K, but it will also depend on the lot size for the production problem. This problem is more complex because not only is the optimal lot size dependent on the solution to the inspection problem (see (7)), but also the solution to the inspection problem is dependent on the optimal lot size.

Another interesting question deals with multiple lots under rigid and nonrigid demand. Under rigid demand the production problem has to be solved for all possible outstanding demand levels. Under nonrigid demand, probably a more involved version, the manufacturer has the right to stop the production and inspection processes and pay a shortage cost if necessary. In the first version the manufacturer faces two possible actions for the inspection problem, namely, "inspect" or "produce," where in the second, a third option exists, which is STOP production and inspection and pay shortage costs if necessary.

Appendix

Proof of Lemma 2. $f(x)$ is assumed to be piecewise linear. Therefore, there exists a collection of intervals $[\tau_0 = 0, \tau_1), [\tau_1, \tau_2), \ldots, [\tau_m, \tau_{m+1}), [\tau_{m+1}, \tau_{m+2}), \ldots$ such that $f(x)$ is linear in each of them. That is, $f(x) = \varphi_m + \omega_m x$ for $x \in [\tau_m, \tau_{m+1})$ and $m \geqslant 0$. The fact that $f(x) = 0$ near zero, and is decreasing, and concave implies that (1) $\varphi_0 = 0$; (2) $\varphi_m < \varphi_{m+1}$ for $m \geqslant 0$; (3) $\omega_0 = 0$; and (4) $\omega_m > \omega_{m+1}$ for $m \geqslant 0$.

Let $g_1(x) = p(x)f(h_0(x))$ and $g_2(x) = q(x)f(h_1(x))$. Both functions are well defined for $x > 0$ because $h_0(x) >$

$h_1(x) > h_1(0) = 0$, as follows from Lemma 1(b). The continuity of g_1 and g_2 follows from the continuity of the functions p, q, f, h_0, and h_1.

(i) We start by proving that $g_1(x)$ is linear in each of the following intervals $[0, h_0^{-1}(\tau_1))$, $[h_0^{-1}(\tau_1), h_0^{-1}(\tau_2)), \ldots$. Consider a specific interval $x \in [h_0^{-1}(\tau_m), h_0^{-1}(\tau_{m+1}))$, i.e., $h_0(x) \in [\tau_m, \tau_{m+1})$. By definition, in this interval $g_1(x) = p(x)(\varphi_m + \omega_m h_0(x)) = \varphi_m p(x) + \omega_m r x \theta_0$, where the last equality follows from the definition of $h_0(x)$; see (4). As can be seen, g_1 is linear in x in this interval, which proves that g_1 is piecewise linear. To prove that g_1 is concave, we need to show that it is concave at the break points, i.e., it is sufficient to show that the slope of g_1 to the left of τ_{m+1} is strictly greater than to the right of τ_{m+1}. This is equivalent to showing that $\varphi_m(\theta_0 - \theta_1) + \omega_m r\theta_0 > \varphi_{m+1}(\theta_0 - \theta_1) + \omega_{m+1} r\theta_0$. Using the fact that g_1 is continuous at τ_{m+1} implies that $(\omega_m - \omega_{m+1}) r\theta_0 \tau_{m+1} = (\varphi_{m+1} - \varphi_m)\theta_1 + (\varphi_{m+1} - \varphi_m)(\theta_0 - \theta_1)\tau_{m+1}$. Given that $\varphi_{m+1} > \varphi_m$, we obtain that the continuity implies the following inequality: $(\omega_m - \omega_{m+1}) r\theta_0 > (\varphi_{m+1} - \varphi_m)(\theta_0 - \theta_1)$, which proves that $g_1(x)$ is concave. In view of the concavity of g_1, it is sufficient to show that it is flat in the first interval to prove that it is decreasing. However, in the interval $[0, h_0^{-1}(\tau_1))$, $g_1(x) = p(x)(\varphi_0 + \omega_0 h_0(x)) \equiv 0$ because $\varphi_0 = \omega_0 = 0$.

(ii) To complete the proof of part (ii), consider the function $g_2(x)$ in the collection of intervals $[0, h_1^{-1}(\tau_1))$, $[h_1^{-1}(\tau_1), h_1^{-1}(\tau_2)), \ldots$. Let $x \in [h_1^{-1}(\tau_m), h_1^{-1}(\tau_{m+1}))$, i.e., $h_1(x) \in [\tau_m, \tau_{m+1})$. By definition, in this interval $g_2(x) = q(x)(\varphi_m + \omega_m h_1(x)) = \varphi_m q(x) + \omega_m r x(1 - \theta_0)$, where the last equality follows from the definition of $h_1(x)$; see (4). Thus, g_2 is linear in the interval, which proves that g_2 is piecewise linear. We prove that g_2 is concave at the break point τ_{m+1} in the same way as we did for g_1. That is, we show that $-\varphi_m(\theta_0 - \theta_1) + \omega_m r(1 - \theta_0) > -\varphi_{m+1}(\theta_0 - \theta_1) + \omega_{m+1} r(1 - \theta_0)$. Using the continuity of g_2 at τ_{m+1} implies that $(\omega_m - \omega_{m+1}) r(1 - \theta_0)\tau_{m+1} = (\theta_0 - \theta_1)(\varphi_m - \varphi_{m+1})\tau_{m+1} + (\varphi_{m+1} - \varphi_m)(1 - \theta_1)$. Using the fact that $\varphi_{m+1} > \varphi_m$, we conclude the concavity proof of $g_2(x)$. As in the proof of part (i), we show that g_2 is decreasing by proving that it is flat near 0. Consider the interval $[0, h_1^{-1}(\tau_1)]$, where $g_2(x) = q(x)(\varphi_0 + \omega_0 h_1(x)) \equiv 0$ because $\varphi_0 = \omega_0 = 0$.

(iii) Note that the function $k(x)$ is the sum of three piecewise linear decreasing and concave functions. The fact that $g(x)$ is strictly decreasing implies, therefore, that $k(x)$ is PLDC. □

Proof of Theorem 3. First note that part (v) is a direct consequence of part (iv). According to Theorem 1, the functions $\Delta_{D,K}(x)$ are PLDC. According to the definitions of x^* and $\mathscr{L}_1$, $\theta_1 < \gamma/s < r\theta_0$ implies that $0 < \mathscr{L}_1 < x^*$. We prove parts (i)–(iv) by induction on K. For part (i), note that for any outstanding demand $D \geqslant 1$ and $K = 1$, the function $\Delta_{D,1}(x) = \gamma - sp(x) = \Omega_1(x)$ is independent of D, linear, and strictly decreasing in x. The optimal policy is to inspect the uninspected unit if and only if $x > \mathscr{L}_1$. Thus, $L_{D,1} = \mathscr{L}_1$ for $D \geqslant 1$. We assume by induction that the PLDC functions $\Omega_k(x)$ for $1 \leqslant k \leqslant K-1$, are well defined. The inequalities $\Omega_1(x) \geqslant \Omega_2(x) \geqslant \cdots \geqslant \Omega_{K-1}(x)$ follow from (9). Their PLDC property implies the existence of a unique CLT value for each $\Omega_k(x)$, named $\mathscr{L}_k$ for $k \leqslant K-1$. The fact that $\Omega_1(x) \geqslant \Omega_2(x) \geqslant \cdots \geqslant \Omega_{K-1}(x)$ implies that $\mathscr{L}_{K-1} \leqslant \cdots \leqslant \mathscr{L}_2 \leqslant \mathscr{L}_1 < x^*$ and therefore, by using Lemma 1(b) and 1(c) regarding the functions $h_0(\cdot)$, $h_1(\cdot)$ and the definition of x^*, we obtain that $\mathscr{L}_k \in (h_0^{-1}(\mathscr{L}_k), h_1^{-1}(\mathscr{L}_k))$ for $k \leqslant K-1$.

We start by proving parts (i), (iii), and (iv) for K, and any $D \geqslant K \geqslant 2$, namely: (i) the PLDC function $\Delta_{D,K}(x)$ is independent of D; which implies that $\Omega_K(x)$ and its root $\mathscr{L}_K$ are well defined. (iii) For $x \leqslant h_0^{-1}(\mathscr{L}_{K-1})$, $\Omega_K(x) = \gamma - sp(x) = \Omega_1(x)$, and for $x > h_0^{-1}(\mathscr{L}_{K-1})$, $\Omega_K(x) < \Omega_{K-1}(x)$; and (iv) $\mathscr{L}_K \in (h_0^{-1}(\mathscr{L}_{K-1}), \mathscr{L}_{K-1})$, which implies $\mathscr{L}_{K-1} > \mathscr{L}_K > 0$. As is shown below, $h_0^{-1}(\mathscr{L}_{K-1})$ and $h_1^{-1}(\mathscr{L}_{K-1})$ are break points of $\Delta_{D,K}(x)$. We distinguish between three cases regarding the value of the information state x:

(a) If $x \leqslant h_0^{-1}(\mathscr{L}_{K-1})$, then $h_1(x) < h_0(x) \leqslant \mathscr{L}_{K-1}$. By the inductive assumptions, $\mathscr{L}_{K-1} \leqslant \mathscr{L}_1$ for $K \geqslant 2$. Lemma 1(c) and $\mathscr{L}_1 < x^*$ imply that $h_0^{-1}(\mathscr{L}_{K-1}) < \mathscr{L}_{K-1}$. Thus, after inspecting the first uninspected item, we obtain a new problem with $K-1$ uninspected items, and a new information state that is either $h_0(x)$ or $h_1(x)$; independently of the quality of the first item, the new information state for $K-1$ uninspected items does not exceed $\mathscr{L}_{K-1}$. According to the inductive assumption for $K-1$, the best policy after inspecting the first item is STOP. Thus, $\Delta^{\text{OPT}}_{D-1,K-1}(h_0(x)) = \Delta^{\text{OPT}}_{D,K-1}(h_1(x)) = 0$ and by (11), $\Delta_{D,K}(x) = \gamma - sp(x) = \Omega_1(x)$. Let $\Omega_K(x) \equiv \Omega_1(x)$ in this interval. Observe that $\Omega_1(x) > 0$ in $[0, \mathscr{L}_1)$ and because $[0, h_0^{-1}(\mathscr{L}_{K-1})] \subset [0, \mathscr{L}_1)$, we obtain that $\Omega_K(x) > 0$ for $x \leqslant h_0^{-1}(\mathscr{L}_{K-1})$, which means that $\mathscr{L}_K > h_0^{-1}(\mathscr{L}_{K-1})$.

(b) If $h_0^{-1}(\mathscr{L}_{K-1}) < x \leqslant h_1^{-1}(\mathscr{L}_{K-1})$, then $h_1(x) \leqslant \mathscr{L}_{K-1} < h_0(x)$. Thus, if the first item is found conforming, the new information state is $h_0(x)$, which is strictly greater than $\mathscr{L}_{K-1}$, and according to the inductive assumption, inspection continues. Otherwise, if the item is defective, the new information state $h_1(x)$ is smaller than $\mathscr{L}_{K-1}$ and inspection stops. Thus, $\Delta^{\text{OPT}}_{D-1,K-1}(h_0(x)) = \Delta_{D-1,K-1}(h_0(x)) = \Omega_{K-1}(h_0(x)) < \Omega_{K-1}(\mathscr{L}_{K-1}) = 0$ and $\Delta^{\text{OPT}}_{D,K-1}(h_1(x)) = 0$. Substituting into (11) results in $\Delta_{D,K}(x) = \gamma - sp(x) + p(x)\Omega_{K-1}(h_0(x)) < \Omega_1(x)$.

(c) If $x > h_1^{-1}(\mathscr{L}_{K-1})$, then $h_0(x) > h_1(x) > \mathscr{L}_{K-1}$, which implies by the inductive assumption that after inspecting the first item, inspection continues independently of its quality. That is, $\Delta^{\text{OPT}}_{D-1,K-1}(h_0(x)) = \Delta_{D-1,K-1}(h_0(x)) = \Omega_{K-1}(h_0(x)) < 0$ and $\Delta^{\text{OPT}}_{D,K-1}(h_1(x)) = \Delta_{D,K-1}(h_1(x)) = \Omega_{K-1}(h_1(x)) < 0$. By substituting into (11), we get that $\Delta_{D,K}(x) = \gamma - sp(x) + p(x)\Omega_{K-1}(h_0(x)) + q(x)\Omega_{K-1}(h_1(x)) < \Omega_1(x)$.

In all three intervals, $\Delta_{D,K}(x)$ is independent of D, so $\Omega_K(x) \equiv \Delta_{D,K}(x)$ for $D \geqslant K$ and its unique root $\mathscr{L}_K$ is well defined. For $K \geqslant 2$, because $\Omega_K(x) \leqslant \Omega_{K-1}(x)$, we obtain that the unique root of $\Omega_K(x)$, namely, $\mathscr{L}_K$, satisfies $\mathscr{L}_K \leqslant \mathscr{L}_{K-1}$. In general the form of $\Omega_K(x)$ is as follows:

$$\Omega_K(x) = \begin{cases} \gamma - sp(x), & x \leqslant h_0^{-1}(\mathscr{L}_{K-1}), \\ \gamma - sp(x) + p(x)\Omega_{K-1}(h_0(x)), & h_0^{-1}(\mathscr{L}_{K-1}) < x \leqslant h_1^{-1}(\mathscr{L}_{K-1}), \\ \gamma - sp(x) + p(x)\Omega_{K-1}(h_0(x)) + q(x)\Omega_{K-1}(h_1(x)), & x > h_1^{-1}(\mathscr{L}_{K-1}). \end{cases} \tag{13}$$

This completes the proof of part (i).

To prove part (iii), note that we have already shown that $\Omega_K(x) = \Omega_1(x)$ for $x \leqslant h_0^{-1}(\mathscr{L}_{K-1})$. Moreover, for $K = 2$, we have also proved that $\Omega_2(x) < \Omega_1(x)$ whenever $x > h_0^{-1}(\mathscr{L}_1)$. Assume by induction that for $x > h_0^{-1}(\mathscr{L}_{k-1})$, $\Omega_k(x) < \Omega_{k-1}(x)$, which implies that $\mathscr{L}_k < \mathscr{L}_{k-1}$ for $k \leqslant K-1$. It remains to prove that for $x > h_0^{-1}(\mathscr{L}_{K-1})$, $\Omega_K(x) < \Omega_{K-1}(x)$. By the inductive assumption, the sequence $\mathscr{L}_k$, $k \leqslant K-1$, is strictly decreasing; thus, $h_0^{-1}(\mathscr{L}_{K-1}) < h_0^{-1}(\mathscr{L}_{K-2})$. Moreover, to the left of $h_0^{-1}(\mathscr{L}_{K-1})$, the functions $\Omega_K(x)$ and $\Omega_{K-1}(x)$ coincide. We distinguish between three cases:

Case 1. $x \in (h_0^{-1}(\mathscr{L}_{K-1}), h_0^{-1}(\mathscr{L}_{K-2})]$: Note that according to the inductive assumption (iv) and according to Lemma 1(b), $h_0^{-1}(\mathscr{L}_{K-2}) < \mathscr{L}_{K-1} < h_1^{-1}(\mathscr{L}_{K-1})$. $h_0(x) > \mathscr{L}_{K-1}$ implies that $\Omega_{K-1}(h_0(x)) < 0$, and $h_1(x) < \mathscr{L}_{K-1}$ implies that $\Omega_{K-1}(h_1(x)) = 0$. Thus, $\Omega_K(x) = \gamma - sp(x) + p(x)\Omega_{K-1}(h_0(x))$. Therefore, in this interval $\Omega_K(x) < \Omega_{K-1}(x) = \gamma - sp(x)$, where the form of $\Omega_{K-1}(x)$ follows from the inductive assumption (iii).

Case 2. If $x \in (h_0^{-1}(\mathscr{L}_{K-2}), h_1^{-1}(\mathscr{L}_{K-2})]$, then $h_0(x) > \mathscr{L}_{K-2} > \mathscr{L}_{K-1}$. According to (i), inspection continues if the first unit is found conforming, thus $\Delta^{\text{OPT}}_{D-1, K-1}(h_0(x)) = \Delta_{D-1, K-1}(h_0(x)) = \Omega_{K-1}(h_0(x)) < 0$. (11) and $\Delta^{\text{OPT}}_{D, K-1}(h_1(x)) \leqslant 0$ imply that $\Omega_K(x) \leqslant \gamma - sp(x) + p(x)\Omega_{K-1}(h_0(x)) < \gamma - sp(x) + p(x)\Omega_{K-2}(h_0(x)) = \Omega_{K-1}(x)$ where the strict inequality follows from inductive assumption (iii) for Ω_{K-1} and Ω_{K-2} and the equality follows from the general form of the function $\Omega_{K-1}(x)$; see (13).

Case 3. If $x > h_1^{-1}(\mathscr{L}_{K-2})$, then the following inequalities hold: $h_0(x) > h_1(x) > \mathscr{L}_{K-2} > \mathscr{L}_{K-1}$. That is, independently of the quality of the first item, inspection continues. Thus, $\Omega_K(x) = \gamma - sp(x) + p(x)\Omega_{K-1}(h_0(x)) + q(x)\Omega_{K-1}(h_1(x)) < \gamma - sp(x) + p(x)\Omega_{K-2}(h_0(x)) + q(x)\Omega_{K-2}(h_1(x)) = \Omega_{K-1}(x)$, where the strict inequality follows from inductive assumption (iii) for $\Omega_{K-1}(x)$ and $\Omega_{K-2}(x)$, and the equality from the form of $\Omega_{K-1}(x)$; see (13). This concludes the proof of part (iii). Part (iv) then follows directly from (iii).

It remains to prove part (ii), which is equivalent to proving that it has at most $2^K - 2$ break points. The statement holds for $K = 1$ because $\Omega_1(x)$ is linear. Assume by induction that the statement holds for $\Omega_k(x)$, for $k \leqslant K-1$, and we will prove it for $\Omega_K(x)$. For that sake, recall the structure of $\Omega_K(x)$; see (13). $\Omega_K(x)$ has break points at $h_0^{-1}(\mathscr{L}_{K-1})$ and $h_1^{-1}(\mathscr{L}_{K-1})$. The function is linear to the left of $h_0^{-1}(\mathscr{L}_{K-1})$ but to the right of this point we should consider all possible break points of both $\Omega_{K-1}(h_0(x))$ and $\Omega_{K-1}(h_1(x))$. The inductive assumption implies that each of the functions $\Omega_{K-1}(h_0(x))$ and $\Omega_{K-1}(h_1(x))$ is associated with at most $2^{K-1} - 2$ break points. Therefore, the total number of break points of $\Omega_K(x)$ is bounded from above by $2 + 2(2^{K-1} - 2) = 2^K - 2$. □

Proof of Theorem 4. (i) The proof is by induction on K. For $K = 1$, see Theorem 3(i). Assume that (i) holds for any $k \leqslant K-1$; we will prove that it holds for $k = K$. Recall from Lemma 1(b) that $h_1(\mathscr{L}_1) < \mathscr{L}_1$ or, equivalently, $h_1^{-1}(\mathscr{L}_1) > \mathscr{L}_1$. We will show that for $x \leqslant h_1^{-1}(\mathscr{L}_1)$, $\Delta_{1, K}(x) = \Omega_1(x)$, and for $x > h_1^{-1}(\mathscr{L}_1)$, $\Delta_{1, K}(x) < \Omega_1(x)$. This will imply that $L_{1, K}$ coincides with the root of $\Omega_1(x)$. Thus, the optimal policy for $D = 1$ is INSP if and only if $x > \mathscr{L}_1$.

Let $x \leqslant h_1^{-1}(\mathscr{L}_1)$. In this interval, if the first item is inspected and is found defective, then the new information state is $h_1(x) \leqslant \mathscr{L}_1$. According to the induction assumption for $D = 1$ and $K-1$, the optimal policy is STOP, i.e., $\Delta^{\text{OPT}}_{1, K-1}(h_1(x)) = 0$. Thus, (11) and (12) imply that $\Delta_{1, K}(x) = \gamma - s \cdot p(x) = \Omega_1(x)$. Let $x > h_1^{-1}(\mathscr{L}_1)$. In this interval $h_1(x) > \mathscr{L}_1$, thus if the first unit is found to be defective, then the new information state $h_1(x)$ is larger than $\mathscr{L}_1$. According to the inductive assumption for $D = 1$ and $K-1$, the optimal policy is INSP. Thus, $\Delta^{\text{OPT}}_{1, K-1}(h_1(x)) < 0$ and $\Delta_{1, K}(x) < \gamma - s \cdot p(x) = \Omega_1(x)$. This proves that the two PLDC functions $\Delta_{1, K}(x)$ and $\Omega_1(x)$ coincide on the first interval $x \leqslant h_1^{-1}(\mathscr{L}_1)$. Because $\mathscr{L}_1$ is within this interval, $\mathscr{L}_1$ is also the root of $\Delta_{1, K}(x)$ for $K \geqslant 1$, i.e., $L_{1, K} = \mathscr{L}_1$ for $K \geqslant 1$.

(ii) For $K > D$, according to (9) and Theorem 3(i), $\Delta_{D, K}(x) \leqslant \Delta_{D, K-1}(x) \leqslant \Delta_{D, D}(x) = \Omega_D(x)$, which implies that for $K \geqslant D$, the sequence $L_{D, K}$ is nonincreasing in K for fixed D. (10b) implies that for $K \geqslant D$, the sequence $L_{D, K}$ is nonincreasing in D for fixed K. Thus, $\mathscr{L}_K = L_{K, K} \leqslant L_{D, K} \leqslant L_{D, D} = \mathscr{L}_D$. In view of Theorem 3(iii), it remains to show that for $K > D$ and $x \leqslant h_0^{-1}(\mathscr{L}_{K-1})$, $\Delta_{D, K}(x) = \Omega_1(x)$: In this interval $h_1(x) < h_0(x) \leqslant \mathscr{L}_{K-1} \leqslant L_{D, K-1} \leqslant L_{D-1, K-1}$, i.e., independently of the quality of the first item, the optimal policy is STOP inspection. Thus, $\Delta^{\text{OPT}}_{D-1, K-1}(h_0(x)) = \Delta^{\text{OPT}}_{D, K-1}(h_1(x)) = 0$ and therefore $\Delta_{D, K}(x) = \gamma - s \cdot p(x) = \Omega_1(x)$. □

Acknowledgments

The authors are grateful to Yigal Gerchak and Rafi Hassin for their insightful comments and suggestions that helped improve the presentation of earlier drafts of this paper. They also thank the anonymous associate editor and referees for their comments, which helped them to improve this paper significantly. This research was supported by the Georges Leven High-Tech Management School, Faculty of Management, Tel Aviv University.

References

Albright, S. C. 1979. Structural results for partially observable Markov decision processes. *Oper. Res.* **27**(5) 1041–1053.

Banerjee, P. K., M. A. Rahim. 1988. Economic design of X-bar control charts under Weibull shock models. *Technometrics* **30** 407–414.

Bertsekas, D. 1976. *Dynamic Programming and Stochastic Control*. Academic Press, New York.

Calabrese, J. M. 1995. Bayesian process control for attributes. *Management Sci.* **41**(4) 637–645.

Crowder, S. V. 1992. An SPC model for short production runs: Minimizing expected cost. *Technometrics* **34** 64–73.

Grosfeld-Nir, A. 1996. A two state partially observable Markov decision process with uniformly distributed observations. *Oper. Res.* **44**(3) 458–463.

Grosfeld-Nir, A., Y. Gerchak, Q.-M. He. 2000. Manufacturing to order with random yield and costly inspection. *Oper. Res.* **48**(5) 761–767.

Hald, A. 1981. *Statistical Theory of Sampling Inspection by Attributes*. Academic Press, New York.

Hassin, R. 1984. A dichotomous search for a geometric random variable. *Oper. Res.* **32** 423–439.

He, Q.-M., Y. Gerchak, A. Grosfeld-Nir. 1996. Optimal inspection order when process' failure rate is constant. *Internat. J. Reliability, Quality Safety* **3**(1) 25–41.

Juran, J. M. 1993. Made in U.S.A.: A renaissance in quality. *Harvard Bus. Rev.* **71**(July/August) 42–50.

Lane, D. E. 1989. A partially observable model of decision making by fishermen. *Oper. Res.* **37**(2) 240–254.

Lauritzen, S. L., N. Nilsson. 2001. Representing and solving decision problems with limited information. *Management Sci.* **47**(9) 1235–1251.

Lee, H. L., M. J. Rosenblatt. 1987. Simultaneous determination of production cycle and inspection schedules in a production system. *Management Sci.* **33** 1125–1136.

Lovejoy, W. S. 1987. Ordered solutions for dynamic programs. *Math. Oper. Res.* **12**(2) 269–276.

Lovejoy, W. S. 1991a. A survey of algorithmic methods for POMDP's. *Ann. Oper. Res.* **28** 47–65.

Lovejoy, W. S. 1991b. Computationally feasible bounds for partially observed Markov decision processes. *Oper. Res.* **39**(1) 162–175.

Monahan, G. E. 1982. A survey of partially observable Markov decision processes: Theory, models and algorithms. *Management Sci.* **28**(1) 1–16.

Montgomery, D. C. 2001. *Introduction to Statistical Quality Control*, 4th ed. John Wiley and Sons, New York.

Parkhideh, B., K. Case. 1989. The economic design of a dynamic X-bar-control chart. *IIE Trans.* **21** 313–323.

Porteus, E. L. 1986. Optimal lot-sizing, process quality improvement and setup cost reduction. *Oper. Res.* **34**(1) 137–144.

Porteus, E. L. 1990. The impact of inspection delay on process and inspection lot-sizing. *Management Sci.* **36**(8) 999–1007.

Porteus, E. L., A. Angelus. 1997. Opportunities for improved statistical process control. *Management Sci.* **43**(9) 1214–1228.

Rahim, M. A., P. K. Banerjee. 1993. A generalized model for the economic design of X-bar control charts for production systems with increasing failure rate and early replacement. *Naval Res. Logist.* **40** 787–809.

Raz, T., Y. T. Herer, A. Grosfeld-Nir. 2000. Economic optimization of off-line inspection in production processes with constant failure rates. *IIE Trans.* **32** 205–217.

Rosenblatt, M. J., H. L. Lee. 1986a. Economic production cycles with imperfect production processes. *IIE Trans.* **18** 48–55.

Rosenblatt, M. J., H. L. Lee. 1986b. A comparative study of continuous and periodic inspection policies in deteriorating production systems. *IIE Trans.* **18** 2–9.

Ross, S. M. 1983. *Stochastic Processes*. John Wiley and Sons, New York.

Scarf, H. E. 2005. Optimal inventory policies when sales are discretionary. *Internat. J. Production Econom.* **93–94** 111–119.

Sondik, E. J. 1971. The optimal control of a partially observable Markov processes. Ph.D. thesis, Department of Engineering-Economic Systems, Stanford University, Stanford, CA.

Tagaras, G. 1994. A dynamic programming approach to the economic design of X-bar charts. *IIE Trans.* **26** 48–56.

Tagaras, G. 1996. Dynamic control charts for finite production runs. *Eur. J. Oper. Res.* **91** 38–55.

Tagaras, G. 1998. A survey of recent developments in the design of adaptive control charts. *J. Quality Tech.* **30**(3) 212–231.

Tagaras, G., H. L. Lee. 1987. Optimal Bayesian single sampling plans with modified beta prior distribution. *Naval Res. Logist.* **34**(6) 789–801.

Tagaras, G., Y. Nikolaidis. 2002. Comparing the effectiveness of various Bayesian $\bar{X}$ control charts. *Oper. Res.* **50**(5) 878–888.

Western Electric Company, Inc. 1958. *Statistical Quality Control Handbook*, 2nd ed. Mack Printing Co., Easton, PA.

White, C. C. 1979. Optimal control-limit strategies for a partially observed replacement problem. *Internat. J. Systems Sci.* **10** 321–331.

White, C. C. 1991. A survey of solution techniques for POMDP. *Ann. Oper. Res.* **32** 215–230.

White, C. C., W. T. Scherer. 1989. Solution procedures for partially observed Markov decision processes. *Oper. Res.* **37**(5) 791–797.

White, J. A. 1970. On absorbing Markov chains and optimum batch production quantities. *AIIE Trans.* **2**(1) 82–87.

Yano, C. A., H. L. Lee. 1995. Lot-sizing and random yields: A review. *Oper. Res.* **43**(2) 311–334.

Zipkin, P. H. 2000. *Foundations of Inventory Management.* Irwin/McGraw-Hill.