IIE Transactions (2005) **37**, 995–1009

ISSN: 0740-817X print / 1545-8830 online
DOI: 10.1080/07408170500232149

Economic optimization of off-line inspection in a process subject to failure and recovery

ALEXANDER FINKELSHTEIN[1], YALE T. HERER[2,*], TZVI RAZ[3] and IRAD BEN-GAL[4]

[1]*Economic Department, Cellcom, Israel*
[2]*Faculty of Industrial Engineering and Management, Technion—Israel Institute of Technology, 32000 Haifa, Israel*
E-mail: yale@technion.ac.il
[3]*Faculty of Management and* [4]*Department of Industrial Engineering, Tel Aviv University, Tel Aviv, Israel*

Received and accepted May 2004

In certain types of processes, verification of the quality of the output units is possible only after the entire batch has been processed. We develop a model that prescribes which units should be inspected and how the units that were not inspected should be disposed of, in order to minimize the expected sum of inspection costs and disposition error costs, for processes that are subject to random failure and recovery. The model is based on a dynamic programming algorithm that has a low computational complexity. The study also includes a sensitivity analysis under a variety of cost and probability scenarios, supplemented by an analysis of the smallest batch that requires inspection, the expected number of inspections, and the performance of an easy to implement heuristic.

1. Introduction

Most researchers and practitioners in the field of production management agree that process improvement contributes more to product quality than the inspection of process output. However, quality inspection, the focus of this study, remains the main mechanism for preventing defects introduced during the production process from reaching the customer.

We distinguish between two inspection modalities: on-line and off-line. On-line inspection is performed during or immediately after the production operations and allows not only the detection of defects but also the adjustment of the production process, if warranted. On-line inspection provides, in effect, the basis for statistical process control and is the more effective of the two. However, in some cases, due to technological or operational constraints, it is infeasible or impractical to perform inspection during the production process. This may be true for cases where the physical environment of the production process does not permit carrying out the inspection in a reliable manner, where the cost of a setup adjustment in mid-batch is high, or where the time for the inspection operation is significant when compared to the time required for the production process itself. In such cases, the accepted procedure is to preserve the production order of the units and to carry out off-line inspection after the entire batch has been processed. Then, by inspecting the output it is possible to identify units for which the process deviated from the original setup. One of the main drawback with off-line inspection is that due to the fact that the units have already been produced before inspection begins, it is not possible for the inspection process to affect the number of units produced.

Even though economic optimization of off-line inspection problems has not drawn as much attention as on-line inspection problems, there are several papers on this important subject in the literature. Most of these papers relate to an environment where there is a finite ordered batch of units produced by a machine subject to random breakdowns.

Hassin (1984) investigated the optimal full inspection policy that minimizes the expected number of inspections needed to locate the exact transition point for the case where the last unit of the batch is known to be non-conforming and the transition probability for every step is time-independent. He *et al.* (1996) extended the model to the case where the actual quality of the last unit of the batch is unknown. Herer and Raz (2000) investigated the same problem setting as Hassin (1984), but allowed for parallel inspections, i.e., inspecting more than one unit at the same time. They developed a policy which is based on information theory.

Raz *et al.* (2000) investigated, as we do here, the economic optimization of the search for the transition unit. They formulated and solved the problem of determining the optimal inspection/disposition policy. The objective function was to minimize expected costs, including the inspection cost per unit, the penalty for accepting a non-conforming

*Corresponding author

unit, and the penalty for rejecting a conforming unit. They solved the problem under the assumption of constant and variable transition probabilities.

In this paper we extend the work of Raz *et al.* (2000) by taking into account the ability of the production process to recover after a failure. Fine (1983) also examined a recovering process. Although his model of recovery is identical to ours, the problem he studies ("effect of quality-based learning on optimal inspection policies") is very different. Fine (1983) gives several examples of where recovery may occur. Recovery may occur in the steel production or metal processing industries. Consider a batch which is passed through an oven on a conveyor belt. The oven's temperature may fluctuate during the processing of the batch. Since it is infeasible to stop the process and inspect and remove the non-conforming units when the process is running, the batch is inspected after the last unit is produced. In this case, rather than searching for the single transition unit, we wish to identify all units that were produced when the process was out of control. If we know the identity of these units with certainty, we will reject all these units that were produced when the process was out of control, and accept all the others. Knowledge of all transition points can be acquired by inspecting all the units in the batch, resulting in high inspection costs. Reducing the number of units inspected would reduce the inspection costs, but would also introduce the probability of acceptance/rejection errors, along with the associated penalties. The inspection facility is assumed to be error-free, meaning that only conforming units will be classified as conforming and only non-conforming units will be classified as non-conforming.

The status of the system during the production of a batch is modeled as a discrete-time two-state Markov process. The objective is to define the inspection/disposition procedure that minimizes the sum of the inspection cost and the penalty costs for incorrect disposition decisions. In addition to an optimal $O(N^2)$-time dynamic programming algorithm, a simple, easy to manage, heuristic rule that provides local optimization and opportunities for parallel inspection, is presented and investigated. The performance of the heuristic policy and the behavior of the optimal policy are explored under various combinations of cost parameters and transition probabilities. The results are used to derive managerial insights for operational and design issues.

This paper is organized as follows. Section 2 describes the production process and defines the problem of finding the optimal inspection/disposition policy, and Section 3 formulates the mathematical model and presents the dynamic programming algorithm used to find the optimal policy. Section 4 explores the behavior of the optimal inspection/disposition policy under various process parameters. Some operational and managerial implications are addressed, in particular the need for inspection and the expected number of inspections. Finally, Section 5 presents a heuristic approach to the problem and compares its performance and that of other heuristic approaches to the optimal policy. Section 6 summarizes the findings and suggests some further research directions.

2. Problem definition and research objectives

2.1. *The process*

We consider a production process that can be in one of two states (see Fig. 1):

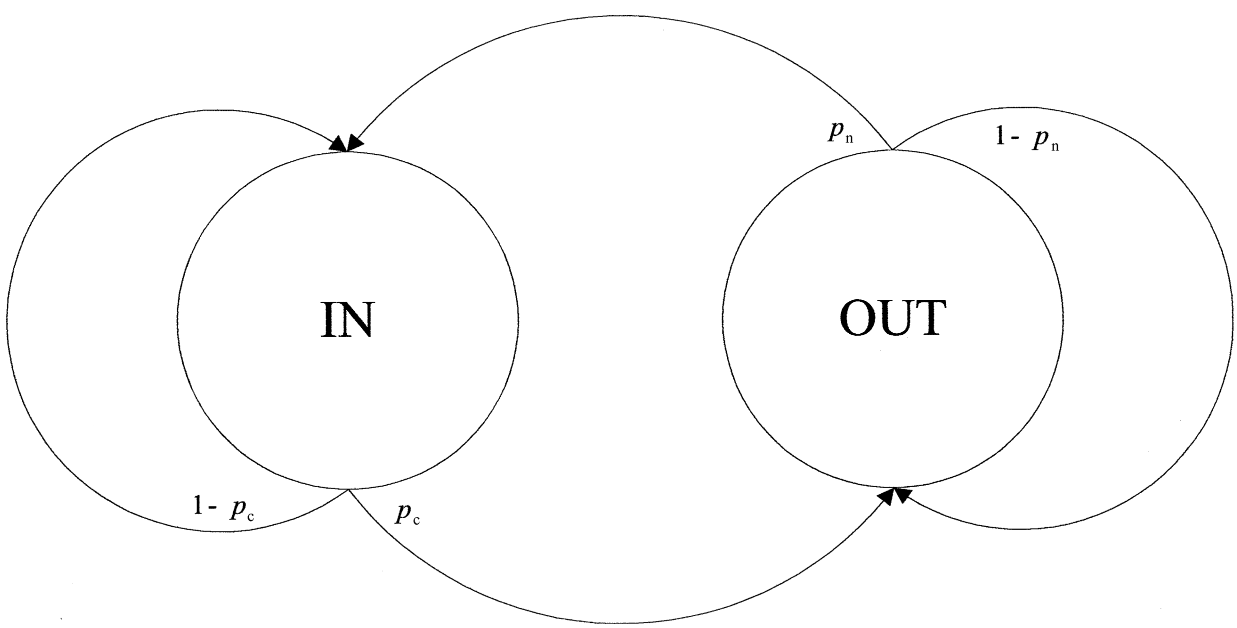

Fig. 1. Markov chain representation of the production system.

1. The process is properly set up and adjusted, such that all the units produced conform to specifications and are acceptable to the customer. This will be referred to as the IN state.
2. The process is incorrectly adjusted, such that all the units fail to meet specifications and are not acceptable to the customer. This will be referred to as the OUT state.

The process operates in a batch mode, and the order in which the units were produced within the batch is preserved. The units in each batch have to be classified according to whether they conform (and should be accepted) or do not conform (and should be rejected) to the quality specifications.

The states of the process before production starts and at the end of production may be known (either IN or OUT) or unknown. Typically a process starts in the IN state. The probability of a transition from the IN state to the OUT state is known and assumed to be constant. Once the transition occurs and the process shifts into the OUT state, all the units produced are non-conforming. There is a constant and known probability (generally different from the first one) that the process will recover and return to the IN state during the production of each unit. The unit produced when a transition occurs is by definition in the new state. Once the process returns to the IN state it again produces conforming units. The process can again shift to the OUT state and so on until completion of the entire batch. Because of the one-to-one relationship between "process state" and "unit condition" these terms will be used interchangeably.

A batch of units consists of a series of conforming and non-conforming units that follow each other. A batch normally starts with a series of conforming units (possibly empty) which ends at the first IN-OUT transition unit. The production of non-conforming units continues until the first OUT-IN transition unit and so on until the end of the batch. Due to the stochastic nature of the process, the batch can contain only conforming units, which means that an IN-OUT transition never occurred. Similarly, the batch can contain only non-conforming units, which means that a transition to the OUT state occurred during the production of the first unit and an OUT-IN transition never occurred.

The classification of any given unit can be done in one of two ways: (i) either the unit is inspected and classified according to the inspection result; or (ii) the unit is classified based on the inspection results of other units in the same batch. This latter option carries with it the risk of committing one of two possible errors: (i) classifying a conforming unit as non-conforming; and (ii) classifying a non-conforming unit as conforming. The only way to avoid these classification errors with certainty is to inspect all the units, but this will probably be too costly. Note that inspecting a unit, not only determines its status with certainty, but also provides information on the probable status of the units around the unit inspected. In fact due to the Markovian nature of the transitions, inspecting a batch divides the batch into two parts: (i) the units produced before (and including) the inspected unit; and (ii) the units produced after the inspected unit. The question of optimally classifying the units in these two sub-batches is identical to the original problem of how to classify the entire batch, albeit the size of the batch and the status of the process (IN, OUT or unknown) at the beginning and/or at the end of the batch may be different. This last observation will form the basis of our dynamic programming solution procedure.

2.2. *Cost function and optimal policy*

If we had perfect *a priori* knowledge of all the transition points of the process, then the optimal policy would be straightforward: accept all conforming units and reject all non-conforming units. The cost of this policy is unavoidable for a given production batch and is the absolute minimum achievable, and will thus serve as our point of reference. Production costs include the unit production costs multiplied by the number of units in the batch, and may also include any applicable setup costs for the batch as a whole. These production costs were all incurred prior to the inspection, and are in fact sunk costs in our model. Consequently, they are not considered separately in the costs in the objective function.

We will address the problem of finding the inspection/classification policy that minimizes the sum of the expected costs above this reference point. Three costs are considered by our model:

1. **The cost of inspection.** All the costs involved in determining whether a unit is conforming or non-conforming. This cost is the same for all units regardless of whether the unit is conforming or non-conforming and is calculated as the sum of direct labor and materials required to inspect one unit, plus an allocation to reflect fixed costs (equipment depreciation, facilities, etc.) plus an allocation for overhead charges, if appropriate.
2. **The cost of classifying conforming units as non-conforming.** If the demand is essentially unlimited, then this cost includes any lost profits that could have been obtained if it were known that the unit was indeed conforming and was eventually sold. Thus, it is the revenue that would have been received for the part minus the value of the additional inputs (labor, materials, machine run time, etc.) that would have been invested in the unit.

   If the units are being produced for a contract that specifies a certain quantity, then we must produce a replacement unit. This cost includes production and inspection costs incurred by the unit, as well as any other costs accumulated so far by the unit (inventory carrying charges, internal transportation, etc.).

   In both these cases we must add the cost of actually scrapping the unit. If there is a salvage value for non-conforming items then this part of the cost may be negative. Although the actual costs may vary among units, it is conceivable that average cost values are available as part of the managerial cost accounting system in place.

3. **The cost of mistakenly classifying non-conforming units as conforming.** This cost is the most difficult to estimate, as it includes the costs of repairing the damage that the non-conforming unit may cause downstream in the value chain.

   From a practical perspective, we suggest the following approach to estimate this parameter. If the process in question is the last one before the product reaches an external customer, we will assume that any non-conforming unit will eventually be detected as such and will have to be repaired, replaced, or dealt with otherwise according to the warranty conditions. Thus, the cost of mistakenly allowing a non-conforming unit to reach the customer can be estimated as the average warranty cost per product unit.

   On the other hand, if the output of the process goes on to the next step in the production sequence, then the cost of missing a non-conforming unit may be estimated as the value of the additional inputs (labor, materials, machine run time, etc.) that will be invested in the unit up to the next inspection where it will be discovered to have been non-conforming.

Note that all these costs are constant and independent of the placement of the unit in the batch and of when the unit is inspected.

Consider, for example, a commercial printing operation that produces packaging materials for consumer goods (wrappers for fast food; labels for jars, etc.). The company has a long-term contract with a fast food chain to supply wrappers printed with the chain logo and the name of the food product, in four colors. The wrappers are printed, cut and bound in batches, with all the processing being done automatically. There is a possibility that the paper sheets will be misfed into the printing machines. The probability of such an event is estimated at 0.0001. Misfeeding is a result of the paper not being properly aligned, and lasts for a relatively small number of cycles. The average run length of misfed sheets is estimated at 100. Assuming an exponential distribution, this run length is equivalent to a probability of recovery of 0.01.

Before shipping the batch to the customer, the wrappers are inspected for defects. Inspecting a single wrapper (pulling it out of the batch, going over it, and replacing it) takes about 1 minute, at an estimated cost (labor + overhead) of \$0.80. The cost of producing a wrapper is estimated at \$0.04, which would be the cost of discarding an acceptable wrapper. The contract with the fast food chain stipulates that if a batch of wrappers contains more than two defectively printed wrappers, then the printing supplier is charged \$250 to cover the costs of screening the entire batch for defects, by the food chain personnel. This figure will be used as the estimated cost of allowing a defective wrapper to reach the customer.

Before starting the development of our optimal solution procedure, let us examine the spectrum of possible solutions. At one end of the spectrum stands the "inspect-all" policy. This policy may be appealing when the penalties for erroneous acceptance/rejection are very high. However, in most situations it is unappealing due to the inherently high inspection costs. At the other extreme stands the "no-inspection" policy. Given the transition (failure and recovery) probabilities of the process and penalties for erroneous acceptance/rejection, the expected penalties for disposition errors can be calculated for each unit of the batch and the appropriate action taken. This policy becomes more appropriate when the production process is stable and when the inspection is costly relative to the penalty costs. The inspect-all policy minimizes the penalty costs (they are zero) by paying for inspecting all of the units, whereas the no-inspection policy minimizes the inspection costs (they are zero) by incurring the risk of paying penalties for the erroneous classification of units. Since our cost is the sum of inspection and penalty costs, we develop a solution method that will minimize the expected value of this sum and will most likely fall between these two extremes.

3. Model formulation

In this section we develop a mathematical formulation of our model and develop an optimal $O(N^2)$-time dynamic program to find the optimal inspection/disposition policy. Before we go into the details of the model, we first present our notation starting with the parameters of the model:

N = number of units in the process batch;
K = number of units in the portion of the process batch being considered, $1 \leq K \leq N$;
p_c = IN-OUT transition probability. This is the probability of a transition from the IN (conforming) state to the OUT state while producing a unit;
p_n = OUT-IN transition probability. This is the probability of a transition from the OUT (non-conforming) state to the IN state while producing a unit, i.e., the process recovers;
C_I = cost of inspection per unit;
C_P = cost of incorrect acceptance, i.e., the penalty resulting from allowing a non-conforming unit to reach the process customer, above and beyond the rework or replacement costs that are normally incurred for any non-conforming unit;
C_S = cost of incorrect rejection, i.e., the loss resulting from scrapping a conforming unit.

To introduce our solution procedure we will need the following notation.

S_b = the status of the system before the start of the batch;
S_e = The status at the end of the batch (i.e., the status of the last unit);

S_m = the status of the unit from the middle of the batch which was inspected.

The variables S_b, S_e, S_m can have one of three possible values, c, n, or u representing, respectively, conforming (IN), non-conforming (OUT), and unknown.

$P_i^{S_b}$ = the probability that unit i will be conforming, given that before the batch started the process was in the S_b state;

$\alpha_i^{S_b S_e}(K)$ = the probability that unit i will be conforming in a batch of size K, given that before (after) the batch started (completed) the process was in the S_b (S_e) state;

$W^{S_b S_e}(K)$ = minimal expected cost (above the reference point) of classifying all the units in a batch of size K without inspection, given that before (after) the batch started (completed) the process was in the S_b (S_e) state;

$G^{S_b S_e}(K)$ = minimal expected cost (above the reference point) of classifying all the units in a batch of size K, given that before (after) the batch started (completed) the process was in the S_b (S_e) state;

$G_j^{S_b S_e}(K)$ = minimal expected cost (above the reference point) of classifying all the units in a batch of size K, given that before (after) the batch started (completed) the process was in the S_b (S_e) state, given that unit j is to be inspected.

The recursive nature of the problem is illustrated in Fig. 2. We begin with a batch of N units. We inspect one of the units, thus dividing the whole batch into two sub-batches. For each of these sub-batches we again inspect a unit, dividing each of these sub-batches into sub-sub-batches. This process continues until it is no longer economically advisable to continue inspecting; at this time we determine the disposition of each of the units using the information that we have about the batch, i.e., S_b and S_e. The recursive equations (i.e., the dynamic program) for determining when and which unit to inspect are the subject of Sections 3.3 and 3.4.

3.1. *Probabilistic aspects*

Before moving to our dynamic programming formulation we first develop a mathematical expression for $\alpha_i^{S_b S_e}(K)$, i.e., the probability that a particular unit will be conforming. First we consider the case where we know the status of the process both before the batch was started and when the batch was completed (i.e., the status of the last unit).

We start by examining the Markov chain representation of our system (see Fig. 1). The case when either p_c or p_n are equal to zero reduces to the problem studied by Raz *et al.* (2000). The transition matrix associated with this Markov chain is as follows:

$$\begin{array}{ccc} & \text{IN} & \text{OUT} \\ \text{IN} & 1-p_c & p_c \\ \text{OUT} & p_n & 1-p_n \end{array}.$$

Using an induction argument on i one can easily show that for all $i \geq 0$ (for notational convenience we define $\beta = 1 - p_n - p_c$):

$$P_i^c = \frac{\beta^i p_c + p_n}{p_n + p_c} \quad \text{and} \quad 1 - P_i^n = \frac{\beta^i p_n + p_c}{p_n + p_c}. \tag{1}$$

Note the symmetry of the problem: the IN and OUT states are symmetric.

Applying Bayes' theorem to the definition of $\alpha_i^{S_b S_e}(K)$ we obtain:

$$\begin{aligned}\alpha_i^{S_b S_e}(K) &= \text{prob}\{\text{unit } i \text{ is conforming} \mid \text{the batch started in the } S_b \text{ state and ended in the } S_e \text{ state}\},\\ &= \frac{\text{prob}\{\text{unit } i \text{ is conforming and the batch ended in the } S_e \text{ state} \mid \text{it started in the } S_b \text{ state}\}}{\text{prob}\{\text{the batch ended in the } S_e \text{ state} \mid \text{the batch started in the } S_b \text{ state}\}}.\end{aligned}$$

Thus,

$$\alpha_i^{cc}(K) = \frac{P_i^c P_{K-i}^c}{P_K^c} = \frac{(\beta^{K-i} p_c + p_n)(\beta^i p_c + p_n)}{(p_n + p_c)(\beta^N p_c + p_n)},$$

and similarly,

$$\alpha_i^{cn}(K) = \frac{P_i^c(1 - P_{K-i}^c)}{1 - P_K^c}, \quad \alpha_i^{nc}(K) = \frac{P_i^n P_{K-i}^c}{P_K^n}, \quad \text{and}$$

$$\alpha_i^{nn}(K) = \frac{P_i^n(1 - P_{K-i}^c)}{1 - P_K^n}.$$

If the process starts in the unknown state and finishes in the unknown state then we have no information about the units produced and thus we use the steady-state probabilities, that is:

$$\alpha_i^{uu}(K) = \frac{p_n}{p_n + p_c}. \tag{2}$$

If the process starts in a known state (either c or n) and ends in the unknown state, then we have information of how the process started and no information of how the process ended. Thus:

$$\alpha_i^{S_b u}(K) = P_i^{S_b}.$$

In the uncommon situation where the status of the process is unknown at the beginning of the batch and known at the end of the batch we can use the reversibility of the Markov chain to obtain:

$$\alpha_i^{u S_e}(K) = P_{K-i}^{S_e}.$$

Now, having demonstrated how to calculate the probabilities $\alpha_i^{S_b S_e}$ we proceed to our dynamic programming

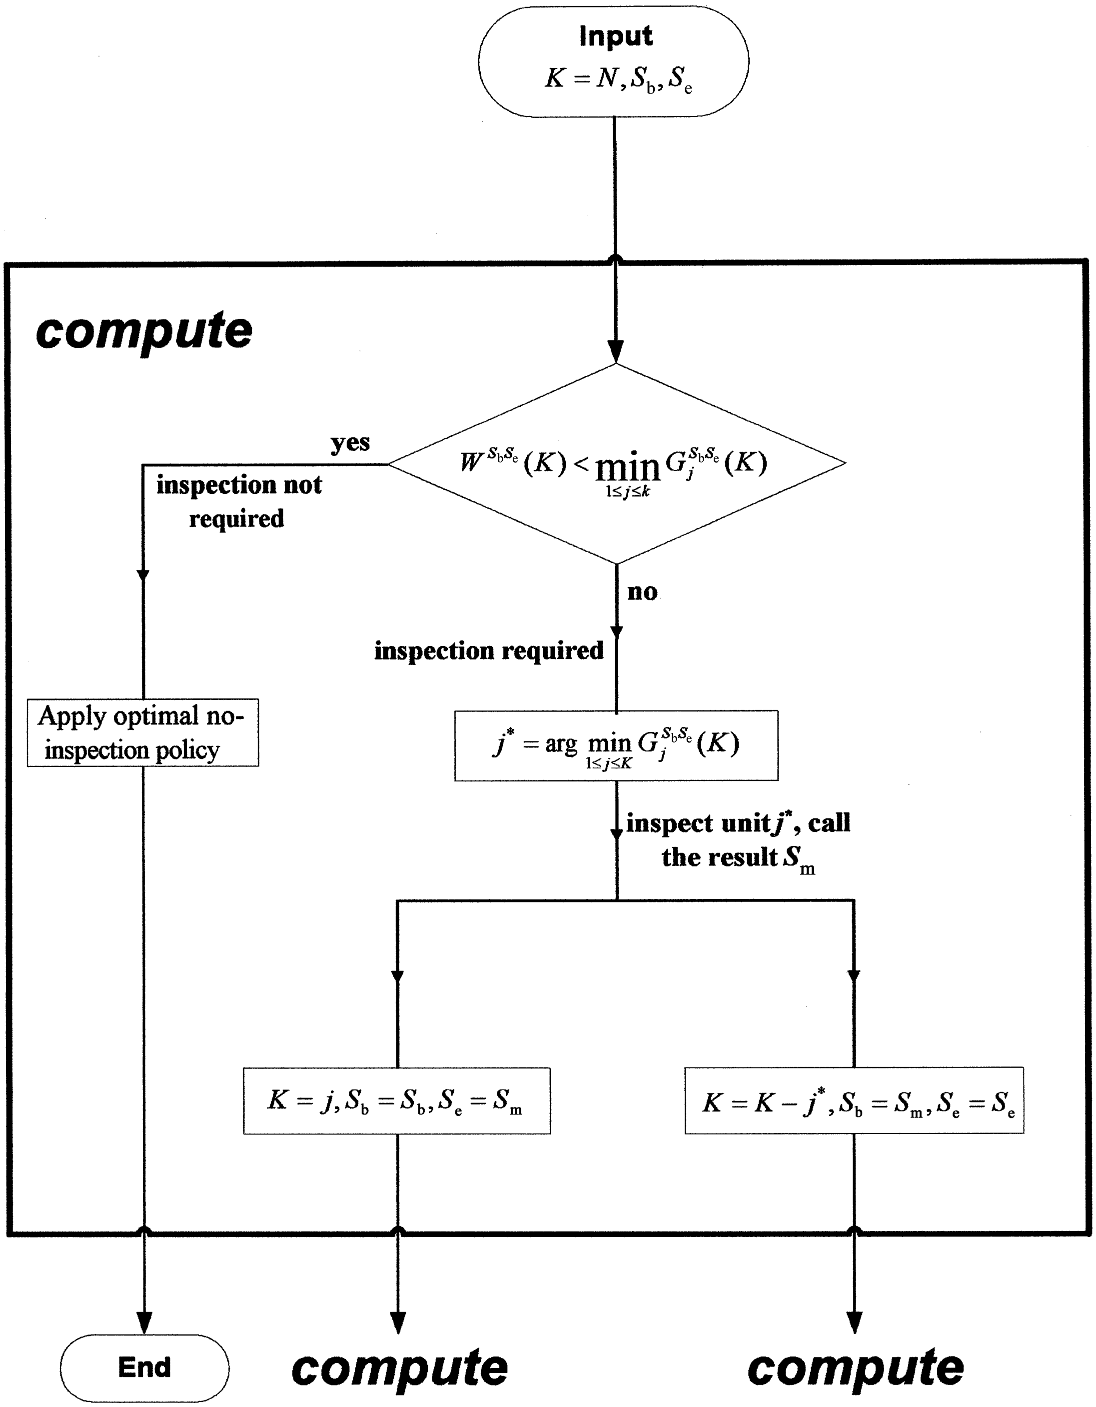

Fig. 2. Calculation of the optimal inspection/disposition policy.

formulation for finding the optimal inspection/disposition policy.

3.2. *Optimal inspection/disposition policy*

The basic problem is to find the first unit to be inspected in a batch of K units (initially $K = N$) when the status of the process before producing the first unit is S_b and the status of the process after producing the Kth unit is S_e. Suppose that we inspect a unit from the middle of the batch, call it unit j, in order to obtain a value for S_m. This action effectively creates two batches: one of size j, including units $1, \ldots, j$ for which the status of the process before producing the first unit (i.e., unit 1) is S_b and the status of the process after producing the last unit (i.e., unit j) is S_m; and the other of size $K - j$, including units $j + 1, \ldots, K$ for which the status of the process before producing the first unit (i.e., unit $j + 1$) is S_m and the status of the process after producing the last unit (i.e., unit K) is S_e. Due to the memoryless property of Markov processes, the probabilities for the status of the units of the second batch are independent of the probabilities of the status of the first batch, i.e., the two batches, given the results of the inspections, are independent of each other. Moreover, the two problems associated with the two parts of the batch (after inspecting unit j) are identical to the initial one, except that the batch size is reduced.

The discussion above assumed that some unit will be inspected. However, this will not always be economically desirable. Thus, we must also consider whether it is cheaper (lower expected cost) to classify all the units without further inspection.

3.3. *No-inspection policy*

In order to determine whether inspection is economically justified we must consider the optimal policy if no inspections are performed. Once we decide not to inspect, then we have to determine the disposition of each unit. To do this we look at the expected cost if we decide to reject the unit, $\alpha_i^{S_bS_e}(K)C_S$ (the probability that the unit is conforming multiplied by the penalty of rejecting a conforming unit), and the expected cost if we decide to accept the unit, $(1-\alpha_i^{S_bS_e}(K))C_P$ (the probability that the unit is nonconforming multiplied by the penalty of accepting a nonconforming unit), and choose the lesser of the two. Summing over all units in the batch we have that:

$$W^{S_bS_e}(K) = \sum_{i=1}^{K} \min\left(\alpha_i^{S_bS_e}(K)C_S, \left(1-\alpha_i^{S_bS_e}(K)\right)C_P\right).$$

3.4. *Recursive equations for determining the optimal policy*

The minimum cost, $G^{S_bS_e}(K)$, can be calculated using the following recursion:

$$G^{S_bS_e}(K) = \min\left[W^{S_bS_e}(K), \min_{1\le j\le K} G_j^{S_bS_e}(K)\right], \tag{3}$$

where

$$\begin{aligned} G_j^{S_bS_e}(K) = C_I &+ \alpha_j^{S_bS_e}(K)\left(G^{S_bc}(j) + G^{cS_e}(K-j)\right) \\ &+ \left(1-\alpha_j^{S_bS_e}(K)\right)\left(G^{S_bn}(j) + G^{nS_e}(K-j)\right). \end{aligned} \tag{4}$$

The outer minimum in Equation (3) represents the choice between not inspecting, $W^{S_bS_e}(K)$, and inspecting one of the units from 1 to K, $\min_{1\le j\le K} G_j^{S_bS_e}(K)$. Equation (4) represents the fact that if we choose to inspect, then we must pay the unit inspection cost C_I. Furthermore, the inspection of unit j could reveal it to be conforming ($S_m = c$), with probability $\alpha_j^{S_bS_e}(K)$, or non-conforming ($S_m = n$), with probability $1-\alpha_j^{S_bS_e}(K)$. Recall that inspection divides the batch into two parts and we must account for the cost of inspecting/disposing of the units in both parts. The first part of the batch ends with the unit being inspected (unit j) and thus has size j, and the second part of the batch starts with unit $j+1$ and thus has size $K-j$. The boundary conditions for the recursion in Equation (3) are: $G^{S_bc}(1) = G^{S_bn}(1) = G^{S_bS_e}(0) = 0$. Note that the boundary condition $G^{S_bS_e}(0) = 0$ comes into play only when we inspect the last unit in the batch, which in turn is only reasonable when the status of the last unit is unknown.

The computational effort necessary to arrive at the optimal inspection/disposition policy is modest. Examining the straightforward implementation of the computations we see that the calculation of each $P_i^{S_b}$ requires $O(1)$ time and the calculation of each $\alpha_i^{S_bS_e}(K)$ also requires $O(1)$ time. Since there are $O(N)$ of the former quantities and $O(N^2)$ of the latter, the total computational effort to calculate all these probabilities is $O(N^2)$ time. After calculating these probabilities the calculation of each $W^{S_bS_e}(K)$ requires $O(N)$ time and since these calculations need to be performed for $K = 1, \ldots, N$, the total calculation time for these quantities is $O(N^2)$. Finally, after calculating the probabilities and the cost of not inspecting, the calculation of each $G^{S_bS_e}(K)$ requires $O(N)$ time and since these calculations need to be performed $K = 1, \ldots, N$, the total calculation time for these quantities is $O(N^2)$. Thus, the time required for a batch of N units is $O(N^2)$.

4. Behavior of the optimal solution

In this section we investigate, through a set of computational experiments, the behavior and sensitivity of the optimal solution to different situations. Additionally, we investigate some operational aspects of the problem that might be of interest to managers, such as the need for inspection and the expected workload for the inspection facility.

4.1. *Parameters for the computational study*

In order to investigate the behavior of the optimal inspection/disposition policy under different situations, 10 sets of cost parameters and 12 different probability combinations were used, making a total of 120 combinations. Additionally, the number of units in the batch varied from one to 500. For the sake of simplicity, we only examine the most common situation, namely that the process starts in the IN state, $S_b = c$, and the status of the process at the end of the batch is unknown, $S_e = u$. The analysis for the other situations would be similar.

The relevant cost parameters for the model are the unit inspection cost, C_I, and the penalties for incorrect acceptance, C_P, and incorrect rejection, C_S. We follow Raz *et al.* (2000) and check 10 combinations of these parameters, termed cost scenarios, as listed in Table 1. These cost scenarios have the following attributes. In scenario A both penalty costs are infinite indicating that we require "perfect information", thus complete inspection is required. In scenario B we require "zero-defects", i.e., no non-conforming part can be

Table 1. Cost scenarios used for the computational study

Cost scenario	C_I	C_P	C_S	*Represented case*
A	1	∞	∞	Perfect information policy
B	1	∞	1	Zero-defects policy
C	1	50	10	Sensitivity analysis with respect to C_P
D	1	10	10	
E	1	1	10	
F	1	10	50	Sensitivity analysis with respect to C_S
G	1	10	1	—together with scenario D
H	50	1	1	Sensitivity analysis with respect to C_I
I	10	1	1	
J	1	1	1	

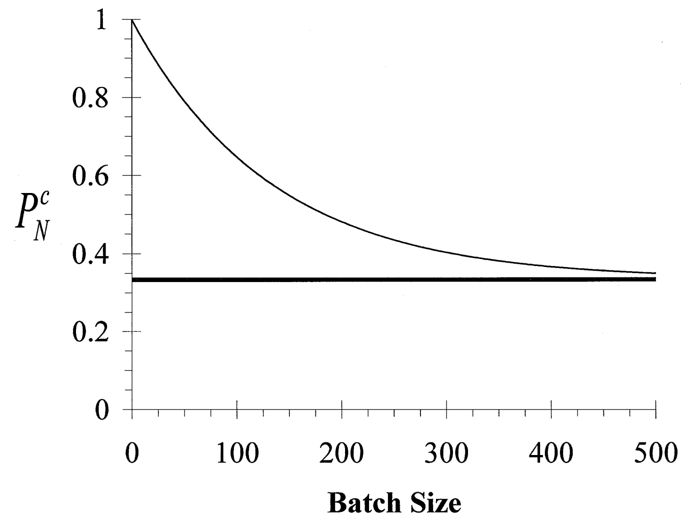

Fig. 3. Illustration of convergence of probabilities for probability scenario I.

accepted, but conforming parts can be rejected. Scenarios C, D and E allow us to perform a sensitivity analysis with respect to C_P. Scenarios F, D and G allow us to perform a sensitivity analysis with respect to C_S. Scenarios H, I and J allow us to perform a sensitivity analysis with respect to C_I.

For each of the 10 cost scenarios, we considered 12 pairs of transition probabilities (p_c, p_n), termed probability scenarios, see Table 2. We again follow Raz *et al.* (2000) in our choice of IN-OUT transition probabilities. To examine the model's sensitivity to the recovery probability, three different values of p_n were used for each p_c, namely, $p_n = 0.5p_c$, p_c, $2p_c$.

The optimal inspection/disposition policy was determined for batch sizes ranging from one unit to 500 units. The reason we did not evaluate the policy for larger batch sizes stems from the Markovian property of the process. Our system converges rather quickly to its steady-state values. Recall that for our computational study $S_b = c$ and $S_e = u$. Thus, we can know how close the system is to reaching its steady-state value by examining the closeness of $P_n^{S_b}$ to its steady-state value $(p_n/(p_n + p_c)$, see Equations (1) and (2)). This is illustrated in Fig. 3 for scenario I. We choose to examine scenario I because it has the slowest convergence of all the probability scenarios. The x-axis is the batch size N and the y-axis is $P_N^{S_b}$. Because the system with the slowest convergence has reached steady-state for a batch size of 500 items, there is no need to examine larger batches.

4.2. *Results of the computation study*

4.2.1. *Expected average cost per unit of the optimal inspection/disposition policy*

Table 3 presents the expected average cost per unit of the optimal inspection/disposition policy, $G^{S_bS_e}(N)/N$, under scenario V for the 10 different cost scenarios and for batch sizes ranging from one to 500. By examining Table 3 we note that as the batch size grows the expected average cost per unit asymptotically converges to a constant value. This indicates that the expected average cost per unit of the optimal inspection/disposition policy for "large" batches is independent of the batch size. The intuitive explanation is straightforward. Similarly to what we saw in Section 4.1 with respect to the last item in the batch, as the batch size grows, the probability that a "middle" unit is conforming converges to its steady-state value rather quickly. As a result the marginal cost of treating each additional "middle" unit converges to some constant value. Therefore, for "large" batches the expected average cost per unit is independent of the batch size. This phenomenon was also observed for all the other probability scenarios.

This observation is important because it allows us to normalize the optimal policy costs in large batches by the batch size. For this reason, for the rest of the computational study, we will present our results in terms of the expected average

Table 2. Probability scenarios used for the computational study

Probability scenario	*I*	*II*	*III*	*IV*	*V*	*VI*	*VII*	*VIII*	*IX*	*X*	*XI*	*XII*
p_c	0.005	0.005	0.005	0.01	0.01	0.01	0.05	0.05	0.05	0.1	0.1	0.1
p_n	0.0025	0.005	0.01	0.005	0.01	0.02	0.025	0.05	0.1	0.05	0.1	0.2

Table 3. Expected average cost per unit under probability scenario V

C_I	C_P	C_S	*Scenario*	*Batch size* 1	50	100	150	200	250	300	350	400	450	500
1	∞	∞	A	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
1	∞	1	B	0.990	0.812	0.713	0.655	0.620	0.597	0.581	0.570	0.561	0.554	0.549
1	50	10	C	0.500	0.192	0.184	0.180	0.177	0.175	0.174	0.173	0.173	0.172	0.172
1	10	10	D	0.100	0.135	0.134	0.134	0.134	0.134	0.134	0.134	0.134	0.134	0.134
1	1	10	E	0.010	0.080	0.088	0.092	0.094	0.096	0.097	0.097	0.098	0.098	0.099
1	10	50	F	0.100	0.145	0.152	0.156	0.159	0.161	0.162	0.163	0.163	0.164	0.164
1	10	1	G	0.100	0.122	0.115	0.112	0.109	0.108	0.107	0.106	0.106	0.105	0.105
50	1	1	H	0.010	0.188	0.287	0.345	0.380	0.403	0.419	0.430	0.439	0.446	0.451
10	1	1	I	0.010	0.188	0.237	0.245	0.256	0.257	0.262	0.262	0.265	0.265	0.266
1	1	1	J	0.010	0.069	0.071	0.072	0.072	0.072	0.073	0.073	0.073	0.073	0.073

cost per unit. Note that for a given batch size the two measures (expected cost and expected average cost per unit) differ by a constant. Thus, the same inspection/dispostion policy minimizes both measures.

4.2.2. *Effect of model parameters on the expected average cost per unit*

Clearly as the cost parameters increase, the expected average cost per unit of the optimal inspection/disposition policy will also increase. Likewise, we would expect the values of the transition probabilities to have an effect on the expected average cost per unit. In order to investigate and quantify these points we evaluated the expected average cost per unit of the optimal inspection/disposition policy for all 120 combinations of the cost and probability scenarios for a batch size of 500. Examining the results in Table 4 the following observations can be made:

- When the penalty costs are infinitely high (scenario A), the expected average cost per unit is insensitive to probability values and is equal to the cost of inspection. The cost of a disposition error is so high that we have to inspect every unit in the batch. That is, we require perfect information about each unit in the batch.
- When the value of either of the penalties increases, the expected average cost per unit increases as well, albeit by a smaller percentage (compare scenarios E, D, C and G, D, F). Since for these sets of parameters not all of the units will be inspected, some disposition errors are inevitable, and the higher the penalty values, the higher the expected average cost per unit.
- When the stability of the production process decreases, as identified by higher values of p_c and p_n, the expected average cost per unit generally increases. Intuitively, this is because when the process is unstable, inspection is less useful in that it gives us less information. When we inspect a unit we find out the status of the process at the time it finished producing that unit. One might infer that other units (produced both before and after this unit) are likely to have the same status. When the process is stable this inference is reasonable, but when the process in unstable, the inference is unreliable.

Table 4. Expected average cost per unit for a large batch (i.e., $N = 500$)

			p_c	0.005			0.01			0.05			0.1		
C_I	C_P	C_S	p_n scenario	0.0025 I	0.005 II	0.01 III	0.005 IV	0.01 V	0.02 VI	0.025 VII	0.05 VIII	0.1 IX	0.05 X	0.1 XI	0.2 XII
1	∞	∞	A	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
1	∞	1	B	0.506	0.598	0.710	0.421	0.549	0.688	0.350	0.509	0.670	0.341	0.504	0.668
1	50	10	C	0.085	0.110	0.144	0.129	0.172	0.226	0.353	0.476	0.620	0.529	0.711	0.915
1	10	10	D	0.067	0.084	0.104	0.106	0.134	0.166	0.305	0.387	0.466	0.467	0.592	0.709
1	1	10	E	0.052	0.060	0.068	0.085	0.099	0.108	0.246	0.276	0.272	0.376	0.400	0.332
1	10	50	F	0.084	0.101	0.118	0.138	0.164	0.190	0.405	0.472	0.525	0.617	0.708	0.768
1	10	1	G	0.052	0.068	0.089	0.077	0.105	0.139	0.195	0.280	0.378	0.277	0.403	0.561
50	1	1	H	0.327	0.358	0.290	0.351	0.451	0.312	0.337	0.491	0.330	0.335	0.496	0.332
10	1	1	I	0.158	0.185	0.203	0.222	0.266	0.282	0.337	0.491	0.330	0.335	0.496	0.332
1	1	1	J	0.037	0.046	0.055	0.058	0.073	0.087	0.154	0.196	0.223	0.224	0.285	0.310

Table 5. Threshold batch size (i.e., minimal batch size for which an inspection is performed)

				p_c											
				0.005			*0.01*			*0.05*			*0.1*		
C_I	C_P	C_S	p_n *scenario*	*0.0025* *I*	*0.005* *II*	*0.01* *III*	*0.005* *IV*	*0.01* *V*	*0.02* *VI*	*0.025* *VII*	*0.05* *VIII*	*0.1* *IX*	*0.05* *X*	*0.1* *XI*	*0.2* *XII*
1	∞	∞	A	1	1	1	1	1	1	1	1	1	1	1	1
1	∞	1	B	*	*	*	*	*	*	*	*	*	*	*	*
1	50	10	C	2	2	2	2	2	2	1	1	1	1	1	1
1	10	10	D	7	7	7	5	5	5	2	2	2	1	2	1
1	1	10	E	25	25	27	19	19	21	11	12	16	10	11	*
1	10	50	F	7	7	7	5	5	5	2	2	2	1	2	1
1	10	1	G	7	7	7	5	5	5	2	2	2	2	2	2
50	1	1	H	248	317	*	*	*	*	*	*	*	*	*	*
10	1	1	I	87	95	117	65	74	124	*	*	*	*	*	*
1	1	1	J	24	24	25	14	18	19	9	9	11	6	7	14

*For these combinations inspection was never economically justified.

- A notable exception to the above point is scenario B. In this case, because of the imbalance in the penalty costs if we do not inspect, we will prefer the units to be non-conforming. Thus, as p_c increases for fixed p_n (scenarios II/IV, III/V, VIII/X, and IX/XI) the expected average cost per unit decreases.
- Another exception occurs when the inspection costs are so high that no inspections are called for (see Table 5) and $C_P = C_S$. In these situations, for a fixed p_c, the expected average cost per unit at first goes up with p_n and then down. The intuition for this is that when no inspections are called for the units must be disposed of without inspection. When $p_c = p_n$ (the situation with the highest cost) the probability that the last unit is conforming/non-conforming asymptotically approaches one-half (see Equation (1)) and thus the expected cost of disposing of the unit without inspection is high. However, when p_n is decreased (respectively, increased) the probability that the last unit is conforming asymptotically approaches one-third (respectively, two-thirds) and the last unit can be rejected (respectively, accepted) with a reasonable expectation that the decision is the correct one.
- When the inspection costs are relatively high (scenarios H and I) and the process is less reliable (high values of p_c), the expected average cost per unit is insensitive to the value of the inspection cost. This can be explained by the fact that in these cases no inspections are called for (as discussed below); the units are classified based on the other information available. However, when the inspection costs are not high enough to cause no inspections to occur then the expected average cost per unit is more sensitive to changes in C_I.

4.2.3. *Need for inspection*

An important operational question when implementing the optimal policy is whether any inspection is economically justifiable. To address this issue, we calculated the minimum batch size that justifies inspection for different parameter values. Recall that for small batches it may be worthwhile to dispose of the units without inspection. As the batch size grows, we have less information about the status of the units in the batch and thus, in some informal sense, inspection gives us more information for a larger number of units. The threshold batch size was obtained by finding the minimum batch size for which the optimal policy called for inspecting at least one unit. Table 5 presents the summary of these results for all 120 combinations of our cost and probability scenarios. Examining the table the following observations can be made:

- When the reliability of the process decreases (larger values of p_c and lower values of p_n) the need for inspection increases, that is, inspection is justified for smaller batches.
- When the inspection costs increase relative to the other cost parameters, the need for inspection decreases, that is, larger batches are required to justify inspection (see scenarios J, I and H).
- The need for inspection is more sensitive to changes in the penalty for incorrect acceptance C_P than changes in the penalty for incorrect rejection C_S (see scenarios E, D, C and G, D, F).

4.2.4. *Expected number of inspections*

In addition to economic optimization managers are interested in operational aspects of the production process, which includes workload analysis. The expected workload on the inspection facility is important for production and inspection planning. To address this issue, we calculated the expected number of inspections for all 120 different combinations of the cost and probability scenarios. Our calculations were performed using the recursive formulas below. First, new notation is defined:

$\nu^{S_b S_e}(K)$ = The unit number to be inspected under the optimal inspection/disposition policy in a batch of size K, given that before (after) the batch started (completed) the process was in the S_b (S_e) state. By convention, if the batch is disposed of without inspection we let $\nu^{S_b S_e}(K) = 0$.

$I^{S_b S_e}(K)$ = The expected number of inspections under the optimal inspection/disposition policy in a batch of size K, given that before (after) the batch started (completed) the process was in the S_b (S_e) state. Note that $I^{S_b S_e}(K)$ is *not* the minimal expected number of inspections possible, but rather the expected number of inspections using the policy that minimizes the expected average cost per unit.

Applying the definition of $\nu^{S_b S_e}(K)$ we have:

$$\nu^{S_b S_e}(K) = \begin{cases} 0 & \text{if the minimum in Equation (3) is } W^{S_b S_e}(K), \\ \arg\min\limits_{1 \le j \le K} \left\{ G_j^{S_b S_e}(K) \right\} & \text{otherwise.} \end{cases}$$

Thus, the recursive formula for calculating the expected number of inspections is:

$$I^{S_b S_e}(K) = \begin{cases} 0 & \text{if } \nu^{S_b S_e}(K) = 0, \\ \alpha^{S_b S_e}_{\nu^{S_b S_e}(K)}(K) \big(I^{S_b c}\big(\nu^{S_b S_e}(K)\big) \\ \quad + I^{c S_e}\big(K - \nu^{S_b S_e}(K)\big)\big) \\ \quad + \big(1 - \alpha^{S_b S_e}_{\nu^{S_b S_e}(K)}(K)\big) \\ \quad \times \big(I^{S_b n}\big(\nu^{S_b S_e}(K)\big) \\ \quad + I^{n S_e}\big(K - \nu^{S_b S_e}(K)\big)\big) & \text{otherwise.} \end{cases}$$

Table 6 summarizes the results for the expected number of inspections, $I^{S_b S_e}(K)$, for a batch size of 500 and all 120 combinations of the cost and probability scenarios. The table reveals trends similar to the results we observed regarding the expected costs and the threshold batch size.

- When the stability of the process decreases, as identified by higher values of p_c and p_n, the expected number of inspections generally increases. Intuitively this is because a single inspection gives less information (see explanation in section 4.2.2). Since each individual inspection gives less information, more inspections are generally needed.
- When the penalties for acceptance/rejection errors are large, more units are to be inspected (compare scenarios E, D, C and G, D, F).
- On the other hand, the expected number of inspections increases as the inspection cost decreases (H, I and J).

Another issue of interest is the expected number of inspections for different batch sizes. Table 7 presents the expected number of inspections for the optimal policy, $I^{S_b S_e}(K)$, as a function of the batch size for scenario V for each of the different cost scenarios.

Raz *et al.* (2000) stated with respect to their model (i.e., without recovery, $p_c = 0$) and for their graph corresponding to Table 7 that "The last part of the curves appears to converge to asymptotic values. This is intuitively plausible: as the batch becomes larger, the probability that the last units are non-conforming becomes practically one, and there is no need for additional inspection to make the correct disposition decision". The present model shows the same phenomenon, albeit the asymptote is to a linear function and not a constant. This is due to the fact that in the present model, as the batch becomes larger, the probability that "middle" units will be non-conforming reaches an asymptote and thus the attention it needs (number of inspections) also reaches an asymptote. In other words, if we examine the expected average number of inspections per unit, it will approach a constant. This phenomenon is consistent with the observation above that the expected average cost per unit approaches a constant as the batch size grows.

Table 6. Expected number of inspections for a large batch (i.e., $N = 500$)

				p_c											
				0.005			*0.01*			*0.05*			*0.1*		
			p_n	*0.0025*	*0.005*	*0.01*	*0.005*	*0.01*	*0.02*	*0.025*	*0.05*	*0.1*	*0.05*	*0.1*	*0.2*
C_I	C_P	C_S	*scenario*	*I*	*II*	*III*	*IV*	*V*	*VI*	*VII*	*VIII*	*IX*	*X*	*XI*	*XII*
1	∞	∞	A	500	500	500	500	500	500	500	500	500	500	500	500
1	∞	1	B	0	0	0	0	0	0	0	0	0	0	0	0
1	50	10	C	32	41	52	48	65	81	132	176	249	205	263	433
1	10	10	D	25	31	39	41	50	63	114	149	178	179	235	264
1	1	10	E	18	21	23	30	34	35	81	82	70	121	104	0
1	10	50	F	31	37	44	51	62	69	143	175	203	248	263	301
1	10	1	G	18	24	31	26	36	47	58	84	122	72	105	181
50	1	1	H	1	1	0	0	0	0	0	0	0	0	0	0
10	1	1	I	4	4	5	5	6	5	0	0	0	0	0	0
1	1	1	J	13	15	18	19	24	26	42	50	55	56	72	55

Table 7. Expected number of inspections for probability scenario V

				Batch size										
C_I	C_P	C_S	*Scenario*	*1*	*50*	*100*	*150*	*200*	*250*	*300*	*350*	*400*	*450*	*500*
1	∞	∞	A	1	50	100	150	200	250	300	350	400	450	500
1	∞	1	B	0	0	0	0	0	0	0	0	0	0	0
1	50	10	C	0	7.08	13.7	20.2	26.7	33.0	39.3	45.6	51.9	58.2	64.6
1	10	10	D	0	4.70	10.4	15.1	19.8	25.5	30.2	35.0	40.6	45.4	50.1
1	1	10	E	0	2.37	5.36	9.39	12.6	16.4	19.8	23.3	26.9	30.4	33.9
1	10	50	F	0	5.11	11.5	17.5	23.9	30.1	36.4	42.8	49.0	55.3	61.7
1	10	1	G	0	3.95	7.91	11.6	15.1	18.7	22.2	25.8	29.3	32.8	36.4
50	1	1	H	0	0	0	0	0	0	0	0	0	0	0
10	1	1	I	0	0	1	1.45	2.43	2.88	2.92	4.32	4.36	5.75	5.80
1	1	1	J	0	1.84	4.62	6.79	8.96	11.7	13.8	16.5	18.7	20.9	23.6

4.2.5. *Effect of imperfect estimation of the model parameters*

Until now we have assumed that all the model parameters are given. In practice, however, this is not always the situation. In fact, if the true cost of inspection, C_{I}, was 1.0, due to the difficulty in estimating the cost parameters, it would not be too surprising if one estimates the inspection cost to be $\hat{C}_{\mathrm{I}} = 1.1$ (we will use the symbol to denote perceived costs). In this subsection we investigate the effects of such estimation errors.

Consider the situation where the costs are estimated as $\hat{C}_{\mathrm{I}}$, $\hat{C}_{\mathrm{P}}$ and $\hat{C}_{\mathrm{S}}$ when the true costs are C_{I}, C_{P} and C_{S}. In this situation, the inspection/disposition policy would be calculated using $\hat{C}_{\mathrm{I}}$, $\hat{C}_{\mathrm{P}}$ and $\hat{C}_{\mathrm{S}}$. This policy would then be implemented, but the actual costs incurred would be C_{I}, C_{P} and C_{S}. We call the ratio of the cost of this policy over the optimal policy (i.e., the policy calculated and implemented using the true costs C_{I}, C_{P} and C_{S}) the error factor due to imperfect cost estimation.

In Table 8 we report the error factor due to imperfect cost estimation for selected combinations of the cost and probability scenarios for a batch size of 500. We examine the cases where the estimated costs differed from the actual costs by 10, 20 and 30%. As can be seen from the table, the system is more sensitive to errors in C_{I} and more sensitive to underestimation, but this is not to say that the system is sensitive. In fact we conclude from Table 8 that the proposed method is robust to imperfect cost estimation. The largest error factor is found to be only 1.07 while the mean error

Table 8. Error factor due to imperfect cost estimation

					p_c											
					0.005			*0.01*			*0.05*			*0.1*		
				p_n	*0.0025*	*0.005*	*0.01*	*0.005*	*0.01*	*0.02*	*0.025*	*0.05*	*0.1*	*0.05*	*0.1*	*0.2*
C_I	C_P	C_S	*Scenario*	*Deviation*	*I*	*II*	*III*	*IV*	*V*	*VI*	*VII*	*VIII*	*IX*	*X*	*XI*	*XII*
1	10	10	D	+10% in C_{P}	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.01
				−10% in C_{P}	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
				+20% in C_{P}	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.02
				−20% in C_{P}	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
				+30% in C_{P}	1.00	1.01	1.01	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.03
				−30% in C_{P}	1.00	1.01	1.01	1.00	1.00	1.01	1.01	1.01	1.01	1.00	1.01	1.00
				+10% in C_{S}	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
				−10% in C_{S}	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
				+20% in C_{S}	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
				−20% in C_{S}	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
				+30% in C_{S}	1.00	1.00	1.00	1.00	1.00	1.00	1.01	1.00	1.00	1.00	1.00	1.00
				−30% in C_{S}	1.00	1.00	1.00	1.01	1.00	1.00	1.00	1.01	1.00	1.01	1.01	1.00
10	1	1	I	+10% in C_{I}	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
				−10% in C_{I}	1.00	1.00	1.00	1.00	1.01	1.00	1.00	1.01	1.00	1.00	1.00	1.00
				+20% in C_{I}	1.01	1.01	1.00	1.01	1.00	1.01	1.00	1.00	1.00	1.00	1.00	1.00
				−20% in C_{I}	1.01	1.01	1.01	1.01	1.02	1.01	1.00	1.03	1.00	1.00	1.00	1.00
				+30% in C_{I}	1.02	1.01	1.01	1.01	1.00	1.02	1.00	1.00	1.00	1.00	1.00	1.00
				−30% in C_{I}	1.02	1.03	1.02	1.02	1.03	1.02	1.00	1.07	1.00	1.00	1.00	1.00

factor is a mere 1.00 (1.004 if we use more significant figures) and the median error factor is also 1.00 (1.001 if we use more significant figures).

5. Heuristic solution

The optimal inspection/dispostion policy was found with a dynamic programming algorithm and each decision depends on the results of previous inspections. This means:

- The operational complexity is high. Before each inspection the inspector needs to consult the policy and use the results of previous inspections as input. This operational complexity may cause the optimal solution to be unimplementable for many real-life production processes.
- Only one inspection can be carried out at a time. As pointed out by Herer and Raz (2000) there are potential advantages of carrying out inspections in parallel.

On the other hand, the two immediately available heuristic policies, namely the Inspect-All and No-Inspection Heuristics may be unacceptably costly. For these reasons, we develop a solution method that is cost efficient, simple to implement, and allows the inspection process on all units to begin simultaneously. Below we develop such a heuristic which we will call the "end-point heuristic".

5.1. *The end-point heuristic*

One of the better known heuristics for on-line inspection is to inspect the process at constant intervals. That is, to set a positive integer ℓ and inspect every ℓ units that exit the machine. See for example Lee and Rosenblatt (1987), Badía *et al.* (2001) and Kim *et al.* (2001). Our End-Point Heuristic can be seen as taking the constant interval concept and applying it to off-line inspection. Basically, our End-Point Heuristic takes a batch of N units and divides it into sub-batches of ℓ units each, for some $\ell = 1 \ldots N$. If N is not divisible by ℓ, then we let each sub-batch have ℓ units except for the last one. For a given ℓ, we inspect the last unit of each sub-batch. The units in each sub-batch are then disposed of in the optimal manner using the information on the status of the process at the end-points of the sub-batches. The expected cost, which we denote cost_ℓ, is equal to the sum of inspection costs and the expected penalty costs of each of the sub-batches. We evaluate the expected cost for each $\ell = 1 \ldots N$ and choose the ℓ with the lowest cost. The following pseudo-code presents the End-Point Heuristic.

input N, S_b, S_e
for $\ell = 1$ to N
  $m = 1$, $\text{cost}_\ell = 0$
  while $m + \ell < N$ do begin

$$\begin{aligned}\text{cost}_\ell = \text{cost}_\ell &+ C_\text{I} + \alpha_m^{S_\text{b}S_\text{e}}(N)\alpha_{m+\ell-1}^{S_\text{b}S_\text{e}}(N)W^{cc}(\ell)\\ &+ \alpha_m^{S_\text{b}S_\text{e}}(N)\big(1-\alpha_{m+\ell-1}^{S_\text{b}S_\text{e}}(N)\big)W^{cn}(\ell)\\ &+ \big(1-\alpha_m^{S_\text{b}S_\text{e}}(N)\big)\alpha_{m+\ell-1}^{S_\text{b}S_\text{e}}(N)W^{nc}(\ell)\\ &+ \big(1-\alpha_m^{S_\text{b}S_\text{e}}(N)\big)\big(1-\alpha_{m+\ell-1}^{S_\text{b}S_\text{e}}(N)\big)W^{nn}(\ell)\end{aligned}$$

  $m = m + \ell$
  end while

$$\begin{aligned}\text{cost}_\ell = \text{cost}_\ell &+ C_\text{I} + \alpha_m^{S_\text{b}S_\text{e}}(N)\alpha_N^{S_\text{b}S_\text{e}}(N)W^{cc}(N-m+1)\\ &+ \alpha_m^{S_\text{b}S_\text{e}}(N)\big(1-\alpha_N^{S_\text{b}S_\text{e}}(N)\big)W^{cn}(N-m+1)\\ &+ \big(1-\alpha_m^{S_\text{b}S_\text{e}}(N)\big)\alpha_N^{S_\text{b}S_\text{e}}(N)W^{nc}(N-m+1)\\ &+ \big(1-\alpha_m^{S_\text{b}S_\text{e}}(N)\big)\big(1-\alpha_N^{S_\text{b}S_\text{e}}(N)\big)W^{nn}(N-m+1)\end{aligned}$$

end for

The cost of the End-Point Heuristic is $\min_\ell \text{cost}_\ell$.

5.2. *Computational study of the end-point heuristic*

In this section we compare, through a computational study, the performance (i.e., the expected cost relative to the optimal policy) of the three heuristics presented in this paper, namely the End-Point, Inspect-All, and No-Inspection Heuristics.

Inclusion of the Inspect-All and No-Inspection Heuristics into the scope of the computational study is important since it allows operations managers to evaluate the trade-off between cost efficiency and the computational effort needed for the optimal and End-Point Heuristic policies. One of the purposes of the comparisons presented in this section is to analyze the characteristics of the production process that affect the relative gain of the optimal over the heuristic solution. Table 9 presents the ratio of the heuristic policy cost to the optimal cost (see Table 4), for the three policies for a batch size of 500 units and for all 120 combinations of the cost and probability scenarios. Examining this table we can make following observations:

- As the inspection cost grows, the Inspect-All Heuristic becomes more expensive relative to the optimal policy (compare scenarios J, I and H). Although this is a trivial result, it has important managerial application: in the presence of high inspection costs, managers should allow the parts to be accepted and/or rejected without knowing their status with certainty.
- In general, the performance of the heuristics relative to the optimal solution improves with higher values of the failure and recovery probabilities, p_c and p_n. Our intuitive explanation for this finding is as follows. The advantage of the optimal policy results from the fact that it is adaptive, in the sense that at any point in time the decision regarding whether to accept, reject or inspect, and if so, which unit, is made on the basis of all the information available at that moment, as reflected by S_b and S_e. In contrast, the No-Inspection and Accept-All Heuristics disregard that information and the End-Point heuristic only utilizes it in a limited manner, since according to this heuristic the inspection interval is fixed and cannot be adjusted in response to the actual values observed.

Table 9. Ratio of expected costs for different policies to the cost of the optimal policy

					p_c											
					0.005			*0.01*			*0.05*			*0.1*		
				p_n	*0.0025*	*0.005*	*0.01*	*0.005*	*0.01*	*0.02*	*0.025*	*0.05*	*0.1*	*0.05*	*0.1*	*0.2*
C_I	C_P	C_S	*Scenario*	*policy*	*I*	*II*	*III*	*IV*	*V*	*VI*	*VII*	*VIII*	*IX*	*X*	*XI*	*XII*
1	∞	∞	A	End-point	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
				No-inspection	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
				Inspect-all	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
1	∞	1	B	End-point	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
				No-inspection	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
				Inspect-all	1.98	1.67	1.41	2.38	1.82	1.45	2.86	1.96	1.49	2.93	1.98	1.50
1	50	10	C	End-point	11.81	9.05	6.96	7.77	5.82	4.42	2.84	2.10	1.61	1.89	1.41	1.09
				No-inspection	55.56	50.86	46.69	31.36	30.95	29.59	9.85	10.63	10.77	6.43	7.07	7.30
				Inspect-all	11.81	9.05	6.96	7.77	5.82	4.42	2.84	2.10	1.61	1.69	1.41	1.09
1	10	10	D	End-point	14.86	11.87	9.66	9.40	7.44	6.03	3.28	2.59	2.15	2.14	1.69	1.41
				No-inspection	54.05	47.66	27.96	32.95	33.55	18.80	11.06	12.69	7.07	7.17	8.38	4.68
				Inspect-all	14.86	11.87	9.66	9.40	7.44	6.03	3.28	2.59	2.15	2.14	1.69	1.41
1	1	10	E	End-point	4.57	4.00	3.24	3.09	2.84	2.38	1.56	1.49	1.21	1.27	1.24	1.00
				No-inspection	9.56	6.65	4.28	6.81	4.57	2.90	2.64	1.78	1.21	1.75	1.24	1.00
				Inspect-all	19.35	16.55	14.78	11.76	10.13	9.29	4.07	3.63	3.68	2.66	2.50	3.01
1	10	50	F	End-point	11.87	9.89	8.46	7.23	6.09	5.26	2.47	2.12	1.91	1.62	1.41	1.30
				No-inspection	58.66	39.74	24.49	41.87	27.46	16.41	16.05	10.40	6.28	10.68	7.00	4.32
				Inspect-all	11.87	9.89	8.46	7.23	6.09	6.26	2.47	2.12	1.91	1.62	1.41	1.30
1	10	1	G	End-point	4.48	3.68	2.81	3.24	2.73	2.12	1.70	1.49	1.26	1.23	1.25	1.12
				No-inspection	9.40	8.54	7.79	5.36	5.15	4.89	1.79	1.82	1.77	1.23	1.25	1.19
				Inspect-all	19.27	14.72	11.27	13.00	9.53	7.20	5.14	3.57	2.65	3.61	2.48	1.78
50	1	1	H	End-point	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
				No-inspection	1.11	1.12	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
				Inspect-all	152.96	139.74	172.66	142.59	110.86	160.37	148.39	101.83	151.72	149.15	110.81	150.70
10	1	1	I	End-point	1.25	1.17	1.11	1.11	1.06	1.05	1.00	1.00	1.00	1.00	1.00	1.00
				No-inspection	2.31	2.17	1.42	1.58	1.69	1.10	1.00	1.00	1.00	1.00	1.00	1.00
				Inspect-all	63.48	54.13	49.20	45.13	37.54	35.42	29.68	20.37	30.34	29.83	20.16	30.14
1	1	1	J	End-point	3.39	2.91	2.34	2.44	2.14	1.77	1.30	1.21	1.12	1.12	1.06	1.05
				No-inspection	9.72	8.73	5.27	6.06	6.19	3.60	2.19	2.51	1.48	1.49	1.74	1.07
				Inspect-all	26.71	21.74	18.20	17.28	13.72	11.54	6.51	5.11	4.49	4.46	3.51	3.23

Now, that advantage is most valuable when the failure and recovery rates are relatively small, i.e., the process is relatively stable, which means that the inferences regarding the status of the other units in the batch that are made based on S_b and S_e are valid. Conversely, the advantage of the optimal policy with respect to the heuristics declines when the failure and recovery rates are higher, i.e., the process is relatively unstable, and the information provided by S_b and S_e is less valuable.

- Not surprisingly, the End-Point Heuristic outperforms the other two heuristics. This is due to the fact that it utilizes information about S_b and S_c, which the other heuristics ignore.

6. Concluding remarks

Off-line inspection following an unreliable production process is an appropriate quality assurance tactic for certain systems. In this paper we developed a model that supports decision-making regarding which units should be inspected and how the units that were not inspected should be disposed of, in order to minimize the sum of the expected inspection and disposition error costs. The model is based on a dynamic programming algorithm that has a low computational complexity, $O(N^2)$, and that can be implemented on electronic spreadsheets without major effort.

The study also included a sensitivity analysis under a variety of cost and probability scenarios, supplemented by an analysis of the smallest batch that requires inspection, the expected number of inspections, and the performance of an easy to implement heuristic.

The main contribution of this work is that it provides a practical tool to assist operations managers and quality managers in managing the performance of a type of processes that thus far has not been fully addressed by the research literature. The value of the tool is enhanced by its ease of implementation and by the various insights and implications obtained from the analyses.

One of the assumptions of our model is that only conforming (non-conforming) units are produced in the IN (OUT) state. Relaxing this assumption would mean that the problem would no longer "separate" when an item is inspected. For example, even if units 7–22 were all non-conforming, this could be due to "bad luck" in the IN state. For this reason the size of the state space would become 3^n.

References

Badía, F.G., Berrade, M.D. and Campos, C.A. (2001) Optimization of inspection intervals based on cost. *Journal of Applied Probability*, **38**(4), 872–881.

Ben-Gal, I., Herer, Y.T. and Raz, T. (2002). Self-correcting inspection procedure under inspection errors. *IIE Transactions*, **34**, 529–540.

Fine, C.H. (1983) Quality control and learning in productive systems. Ph.D. thesis, Graduate School of Business, Stanford University, Stanford, CA, USA.

Hassin, R. (1984) A dichotomous search for a geometric random variable. *Operations Research*, **32**, 423–439.

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

Herer, Y.T. and Raz, T. (2000) Optimal amount of parallel inspection for finding the first nonconforming unit in a batch—an information theoretic approach. *Management Science*, **46**, 845–857.

Kim, C.H., Hong, Y.S. and Chang, S.Y. (2001) Optimal production run length and inspection schedules in a deteriorating production process. *IIE Transactions*, **33**(5), 421–426.

Lee, H.L. and Rosenblatt, M.J. (1987) Simultaneous determination of production cycle and inspection schedules in a production system. *Management Science*, **33**(9), 1125–1136.

Raz, T., Herer, Y.T. and Grosfeld-Nir, A. (2000) Economic optimization of off-line inspection. *IIE Transactions*, **32**, 205–217.

Biographies

Alexander Finkelshtein holds a B.A. in Economics from the Hebrew University in Israel (1994) and an MBA from Tel-Aviv University (2000). Currently he is a PhD student at Business Management School of Tel-Aviv University. Alexander is working in the Economic Department of Cellcom Telecom in field of Managerial Accounting and Financial Analysis.

Yale T. Herer, B.S. (1986), M.S. (1990), Ph.D. (1990), Cornell University, Department of Operations Research and Industrial Engineering. Yale joined the Faculty of Industrial Engineering and Management at the Technion—Israel Institute of Technology in 1990 immediately after the completion of his graduate studies. In 1997 Yale joined the Department of Industrial Engineering at Tel-Aviv University and in 2001 he returned to the Technion. Recently, Yale is on sabbatical at Northwestern University in the Department of Industrial Engineering and Management Sciences. Yale also briefly visited Cornell University in the Department of Operations Research and Industrial Engineering. Recently, Yale served as the Head of the Industrial Engineering Area at the Technion. He has worked for several industrial concerns, both as a consultant and as an advisor to project groups. In 1996 Yale received the *IIE Transactions* best paper award. Yale is a member of the Institute for Operations Research and Management Sciences (INFORMS), Institute of Industrial Engineers (IIE), and the Operations Research Society of Israel (ORSIS). He serves on the Editorial Board of *IIE Transactions*. His research interests include supply chain management, especially when integrated with transshipments. He is also interested in production control and production system design.

Tzvi Raz holds B.Sc, M.A.Sc and Ph.D degrees in Industrial and Management Engineering. He is a Professor of Technology and Information Systems Management at the Leon Recanati Graduate School of Business Administration at Tel Aviv University and Director of the Henry Crown Institute of Business Research in Israel. Previously, he managed a technology insertion program at an IBM software development laboratory, and was on the Industrial Engineering faculties of the University of Iowa, US and Ben Gurion University, Israel. Dr. Raz has published over 60 research papers in refereed international journals and is on the editorial boards of Computers and Operations Research and the Project Management Journal.

Irad Ben-Gal is Senior Lecturer at the Department of Industrial Engineering in Tel-Aviv University. He holds a B.Sc. (1992) degree from Tel-Aviv University, M.Sc. (1996) and Ph.D. (1998) degrees from Boston University. Irad is a member of the Institute for Operations Research and Management Sciences (INFORMS) and the Institute of Industrial Engineers (IIE). He is the head of the Computer Integrated Manufacturing (CIM) lab at Tel-Aviv University. Irad has worked for several years in various industrial organizations as a consultant and a project manager. His research interests include quality control, design-of-experiments, testing procedures, and application of information theory to industrial and bioinformatics problems.

Contributed by the Online Quality Engineering Department