3

Quality Control for Products Supplied with Warranty

The objective of this chapter is to develop an inspection procedure for end products that are supplied to customers with some type of warranty (or service contract), which obliges the manufacturer to provide repair, replacement, or, in some cases, refund to the consumer for a product that has failed within a certain period of time as specified by the contract.

Consider the context of batch manufacturing. Suppose we are dealing with a batch of N units, and the quality of the batch is characterized by Θ, the proportion of defective units. Assume Θ is only known through its distribution. Both defective and non-defective units have random lifetimes with given distributions.

Suppose there is an inspection-repair procedure that can identify and repair all the defective units. Hence, if we follow a 100% inspection, we can guarantee that all N units are non-defective before they are shipped. However, inspection and repair do not come free. In the very least, they will consume production capacity. The essence of the problem here is to strike a balance between the warranty cost and the inspection-repair cost.

The policy that we shall identify and prove to be optimal has a simple, sequential structure: It is characterized by a sequence of threshold values, $d_{n_0} \leq \ldots \leq d_n \leq \ldots \leq d_{n_1}$, such that if D_n denotes the number of defective units among n inspected units, then the optimal policy is to stop inspection at the first n that satisfies $D_n < d_n$.

The key that underlies the optimality of this policy is a simple and intuitive monotone property: the higher the defective rate – not in terms of Θ, but in terms of its posterior estimate, given the outcome of the inspection – the more inspection an optimal policy will call for. It turns out that this monotone property is a direct consequence of the warranty cost, as a function of the number of inspected units and the conditional defective rate, satisfying a so-called K-submodularity property, which is a strengthening of the usual notion of submodularity. Based on this property, we are able to identify several structural results of the optimal policy, and eventually characterize the policy itself in terms of certain simple thresholds.

In §3.1 we spell out the precise details of the problem. We then elaborate on the K-submodularity property of the expected warranty cost in §3.2, and related properties of the conditional distribution of the defective rate in §3.3. The optimal control problem is formulated in §3.4, where several key structural properties of the optimal policy are established in Theorems 3.11, 3.12, and 3.14, which lead to a statement of the optimal policy in Theorem 3.16. A special case, the individual warranty model, is studied in §3.5. Two numerical examples and possible extensions are presented in §3.6.

3.1 Warranty Cost Functions

A batch of N units of a certain product has been completed from the production line. The units will supply customer demand, under some kind of warranty that will be specified below. We want to devise an inspection-repair procedure so as to ensure quality, and to balance inspection-repair cost on the one hand and warranty cost on the other hand.

Assume each unit in the batch of N is either defective or non-defective. A non-defective unit has a lifetime of X, and a defective unit has a lifetime of Y. Both X and Y are random variables. Suppose X and Y are ordered under stochastic ordering, $X \geq_{st} Y$, i.e.,

$$\mathsf{P}[X \geq a] \geq \mathsf{P}[Y \geq a] \text{ for all } a \geq 0.$$

(Refer to Definition 2.1.)

Assume an inspection procedure can identify whether a unit is defective or not, at a cost of c_i per unit. Each defective unit identified by the inspection is repaired, at a cost of c_r per unit, and becomes a non-defective unit.

The quality of the batch, before any inspection and repair, is represented by Θ, the proportion of defective units in the batch. Here Θ is assumed to be a random variable, with a known distribution function. Without loss of generality, assume $\Theta \in [\theta_0, \theta_1]$, where θ_0 and θ_1 are two given constants, $0 \leq \theta_0 \leq \theta_1 \leq 1$. (This essentially follows the quality model of Mamer [60].) Note that letting $\theta_0 = \theta_1 = \theta$ models the special case of a deterministic $\Theta \equiv \theta$. However, this special case restricts the number of defectives in n items to a binomial distribution, with a squared coefficient of variation equal to $(1-\theta)/(n\theta)$, much too small – when n is large – for many applications.

Since there is no *a priori* discernible information about the quality of any units in the batch, we assume that each inspection will identify a defective unit with probability Θ and a non-defective unit with probability $1 - \Theta$. Note here Θ is itself a random variable. That is, we do not have the exact information about the quality of the batch, although each inspection will improve our estimation of Θ, in the sense of obtaining an updated *conditional* distribution.

Let $Z(\theta)$ denote a random variable that is equal in distribution to Y (resp. X) with probability θ (resp. $1 - \theta$). That is, given $\Theta = \theta$, $Z(\theta)$ denotes the lifetime of a unit in the batch that has not been inspected.

Let $C(t)$ denote the warranty cost, a function of the lifetime of the units. (The functional form of $C(t)$ depends, of course, on the type of the warranty; see the examples below.) Specifically, suppose $\Theta = \theta$, and exactly n units in the batch are inspected (and the defectives repaired). Then, the expected warranty cost is

$$\phi(n, \theta) := \mathsf{E}[C(X_1 + \cdots + X_n + Z_{n+1}(\theta) + \cdots + Z_N(\theta))], \tag{3.1}$$

where X_i's and Z_j's are i.i.d. random variables that follow the distributions of X and $Z(\theta)$, respectively. We assume that $C(t)$ is decreasing and convex in t (for the obvious motivation presented in the last section).

An example is in order. Consider, for instance, the so-called "cumulative warranty" ([8]): it covers the batch as a whole, with a warranty period (for the entire batch) of NW time units, where W is a given positive constant. This type of warranty applies mostly to reliability systems, where spare parts (in cold stand-by) are used extensively. Let T denote the argument of $C(\cdot)$ in (3.1). Suppose the warranty cost takes the following form:

$$C(T) = (cN)[NW - T]^+/(NW) = c[N - T/W]^+, \tag{3.2}$$

where $[x]^+$ denotes $\max\{x, 0\}$, and $c > 0$ is the selling price of each unit. Under this model, the manufacturer pays back part of the selling price on a pro rata basis. Here C is obviously a decreasing and convex function.

Another case of interest is when $C(\cdot)$ is an *additive* function. That is, (3.1) takes the following form:

$$\begin{aligned} \phi(n, \theta) \quad = \quad & \mathsf{E}[C(X_1)] + \cdots + \mathsf{E}[C(X_n)] + \mathsf{E}[C(Z_{n+1}(\theta))] \\ & + \cdots + \mathsf{E}[C(Z_N(\theta))]. \end{aligned} \tag{3.3}$$

This is the usual individual warranty model, i.e., the warranty applies to each individual unit instead of the batch as a whole. (For instance, let $N = 1$ in (3.2) and apply the function to each individual unit.) In this case, we only need to assume the decreasing property of $C(\cdot)$, the convexity being replaced by the additivity.

In either case, the expected total (inspection, repair and warranty) cost associated with a batch in which exactly $n \leq N$ units are inspected can be expressed as follows:

$$\Pi(n, \theta) = c_i n + c_r n\theta + \phi(n, \theta). \tag{3.4}$$

In the rest of this section and the following three sections (§3.2 through §3.4) we will focus on the cost model in (3.1) and derive the optimal policy. In §3.5 we will show that the same optimal policy, in a simplified form, applies to the individual warranty model in (3.3) as well.

To conclude this section, we illustrate two points: (a) the ϕ function in (3.1) preserves the decreasing convexity of C, and (b) a relationship between the repair cost and the warranty cost is implicit in our model.

We will make frequent use of the following notation:

$$X_{1,n} := X_1 + \cdots + X_n, \quad Z_{n,N}(\theta) := Z_n(\theta) + \cdots + Z_N(\theta).$$

Lemma 3.1 Given θ, $\phi(n, \theta)$ is decreasing and convex in n.

Proof. Clearly, $X \geq_{st} Y$ implies $X \geq_{st} Z(\theta)$ (for any $\theta \in [\theta_0, \theta_1]$). For simplicity, below we omit the argument θ. Since the random variables are independent, we have

$$X_{1,n} + Z_{n+1,N} \leq_{\text{st}} X_{1,n+1} + Z_{n+2,N},$$

and hence

$$\mathsf{E}[C(X_{1,n} + Z_{n+1,N})] \geq \mathsf{E}[C(X_{1,n+1} + Z_{n+2,N})],$$

since $C(t)$ is a decreasing function. That is, ϕ is decreasing in n.

To establish convexity, we use coupling, similar to the approach in Examples 2.16 and 2.17. Since $X \geq_{st} Z$, we can have, for $j = 1, 2$, $X^j \geq Z^j$ almost surely (a.s.), with X^j and Z^j equal in distribution to X and Z, respectively. Let $\tau := X_{1,n-1} + Z_{n+2,N}$. Maintain independence wherever necessary. Since $C(t)$ is convex, we have

$$C(\tau + X^1 + X^2) + C(\tau + Z^1 + Z^2) \geq C(\tau + X^1 + Z^2) + C(\tau + X^2 + Z^1) \quad \text{a.s.}$$

Taking expectations on both sides yields

$$\phi(n+1) + \phi(n-1) \geq \phi(n) + \phi(n) = 2\phi(n),$$

which is the required convexity. □

Clearly, in the above proof, we have actually established that the warranty cost $C(X_{1,n} + Z_{n+1,N})$, as a function of n, is SICX-sp, as in Definition 2.9, which is stronger than the convexity of the expected cost.

In the earlier model description, we have assumed that each defective unit identified by the inspection must be repaired. While this assumption itself appears reasonable and innocuous, it does impose certain restrictions on the warranty cost. Imagine that if the warranty cost were sufficiently low, or the repair cost were relatively high, then the manufacturer might choose to repair only some of the defective units, while taking a chance on the others. Hence, in order to be consistent with our assumption that all defective units are repaired, we will insist that the following condition be satisfied:

$$c_r \leq \mathsf{E}[C(X_{1,n-1} + Y + Z_{n+1,N}(\theta)] - \mathsf{E}[C(X_{1,n-1} + X + Z_{n+1,N}(\theta)]$$

for all $n \leq N - 1$ and all θ. The condition says that suppose we have identified a certain defective unit in the batch, it pays to have it repaired, since the repair cost plus the subsequent warranty cost will not exceed the warranty cost with the defective unit shipped unrepaired. Below, in Lemma 3.5 we will show that the right side above is decreasing in n. Hence, the condition reduces to a single inequality (the case of $n = N - 1$):

$$c_r \leq \mathsf{E}[C(X_{1,N-1} + Y)] - \mathsf{E}[C(X_{1,N})]. \tag{3.5}$$

Throughout our study of the cumulative warranty model (from here through §3.4), we assume that the condition in (3.5) is always in force.

3.2 K-Submodularity

When the defective rate Θ is a known constant, say θ, the quality control problem becomes a *static* optimization problem: we want to find the optimal $n^*(\theta)$ that minimizes the expected total cost $\Pi(n, \theta)$ in (3.4), for a given θ. (The problem is static, because there is nothing to be gained after each inspection, in terms of estimating the defective rate.)

Specifically, we want to establish the following monotone property of the optimal solution: $n^*(\theta') \geq n^*(\theta)$ for all $\theta' \geq \theta$. That is, the lower the quality of the batch (in terms of a larger θ), the more we need to inspect. It turns out that the key to this is the notion of K-submodularity defined below.

Recall the isotone property involved in minimizing a submodular function (§2.3): suppose $x^*(y)$ is the optimal solution to the minimization problem, $\min_x g(x, y)$, for a given y; then $x^*(y)$ is increasing in y. However, here we are interested in a slightly different problem: $\min_x[Kxy + g(x, y)]$, where $K > 0$ is a constant [cf. $\Pi(n, \theta)$ in (3.4)]. Since Kxy is a *super*modular function, the submodularity of g will not guarantee the increasingness of the optimal solution $x^*(y)$ in y. In order to maintain the isotone property of the optimal solution, we need to strengthen the submodularity of g.

Definition 3.2 A bivariate function, $g(x, y)$, is called K-submodular, if for some $K \geq 0$, we have,

$$[g(x_1, y_2) + g(x_2, y_1)] - [g(x_1, y_1) + g(x_2, y_2)] \geq K(x_1 - x_2)(y_1 - y_2),$$

for all $x_1 \geq x_2$ and $y_1 \geq y_2$.

Remark 3.3 (1) Following the above definition, a K-submodular function has the following geometric property: Consider its values on the four corner points, (x_1, y_1), (x_1, y_2), (x_2, y_1) and (x_2, y_2), of a rectangle; the off-diagonal sum is greater than the diagonal sum by at least K times the area of the rectangle. Obviously, K-submodularity specializes to submodularity with $K = 0$.
(2) The term, K-submodularity, is inspired by the notion of K-convexity, which plays a key role in proving the optimality of (s, S) inventory policies; refer to Scarf [81].

Clearly, from Definition 3.2, $g(x, y)$ is K-submodular if and only if $Kxy + g(x, y)$ is submodular. Hence, the next lemma follows from the known isotone property in minimizing a submodular function mentioned earlier. However, we still give a proof below, because we need the details to support the later extension of the result to the stochastic setting (see Proposition 3.8).

Lemma 3.4 Let $x^*(y)$ be an optimal solution to $\min_x[Kxy + g(x, y)]$, for a given y. Then, $x^*(y)$ is increasing in y, if $g(x, y)$ is K-submodular. (In the event of multiple optimal solutions, then $x^*(y)$ is taken to be the largest one.)

Proof. Denote $x_1 = x^*(y)$. Pick $y' > y$. We show that if $x_2 < x_1$, then x_2 cannot provide a better solution than x_1 at y', hence x_2 cannot be $x^*(y')$. Use contradiction. Suppose x_2 yields a better solution at y', i.e.,

$$Kx_2y' + g(x_2, y') < Kx_1y' + g(x_1, y').$$

From the optimality of x_1 at y, we have,

$$Kx_1y + g(x_1, y) \leq Kx_2y + g(x_2, y).$$

Summing up the two inequalities, we have

$$[g(x_1, y) + g(x_2, y')] - [g(x_1, y') + g(x_2, y)] < K(x_1 - x_2)(y' - y),$$

contradicting the K-submodularity of g. $\square$

Lemma 3.5 The expected warranty cost $\phi(n, \theta)$ is K-submodular in (n, θ) with $K = c_r$, i.e.,

$$[\phi(n, \theta) + \phi(n-1, \theta')] - [\phi(n, \theta') + \phi(n-1, \theta)] \geq c_r(\theta' - \theta),$$

for all n and all $\theta' > \theta$.

Proof. Consider $\phi(n-1, \theta) - \phi(n, \theta)$. Conditioning on $Z_n(\theta) = X$ or Y, we have

$$\begin{aligned} & \phi(n-1, \theta) - \phi(n, \theta) \\ = \; & \theta\{\mathsf{E}[C(X_{1,n-1} + Y + Z_{n+1,N}(\theta))] - \mathsf{E}[C(X_{1,n-1} + X + Z_{n+1,N}(\theta))]\} \\ := \; & \theta\Delta(n, \theta). \end{aligned} \tag{3.6}$$

We want to show $\theta'\Delta(n,\theta') - \theta\Delta(n,\theta) \geq c_r(\theta' - \theta)$, for all $n \leq N$ and all $\theta' > \theta$.

We first note that $\Delta(n,\theta)$ is increasing in θ. This follows easily from the definition of Δ in (3.6) and a coupling argument (similar to the one in the proof of Lemma 3.1), taking into account that (a) $X \geq_{st} Y$, (b) $Z(\theta)$ is (stochastically) decreasing in θ, and (c) $C(t)$ is decreasing and convex. Next, also note that the decreasing convexity of ϕ in n (Lemma 3.1) implies that $\Delta(n,\theta)$ is decreasing in n [cf. (3.6)].

Making use of the above two properties, we have

$$\begin{aligned}
& \theta'\Delta(n,\theta') - \theta\Delta(n,\theta) \\
\geq\ & (\theta' - \theta)\Delta(n,\theta) \\
\geq\ & (\theta' - \theta)\Delta(N,\theta) \\
=\ & (\theta' - \theta)\{\mathsf{E}[C(X_{1,N-1} + Y)] - \mathsf{E}[C(X_{1,N})]\} \\
\geq\ & (\theta' - \theta)c_r,
\end{aligned}$$

where the last inequality follows from (3.5). □

Clearly, $c_i n + \phi(n,\theta)$ is also K-submodular. From Lemma 3.4, we have

Proposition 3.6 The optimal solution $n^*(\theta)$ that solves $\min_{0\leq n\leq N} \Pi(n,\theta)$ for any given θ is increasing in θ.

Next, suppose instead of $\Pi(n,\theta)$, we want to minimize

$$\bar{\Pi}(n,\Theta) := \mathsf{E}[\Pi(n,\Theta)] = c_i n + c_r n\mathsf{E}[\Theta] + \mathsf{E}[\phi(n,\Theta)]. \tag{3.7}$$

(Note that this is still a static optimization problem: the defective rate, although a random variable, does not change with respect to n.)

Lemma 3.5 can be readily adapted to the stochastic setting; refer to §2.3, Definition 2.18 in particular. From Lemma 3.5, we know that $\phi(n-1,\theta) - \phi(n,\theta) - c_r\theta$ is increasing in θ, and hence, the following:

Proposition 3.7 The warranty cost $\phi(n,\Theta)$ is stochastically K-submodular, with $K = c_r$, in the following sense:

$$\{\mathsf{E}[\phi(n-1,\Theta')] + \mathsf{E}[\phi(n,\Theta)]\} - \{\mathsf{E}[\phi(n-1,\Theta)] + \mathsf{E}[\phi(n,\Theta')]\} \geq c_r\mathsf{E}[\Theta' - \Theta],$$

for all n and $\Theta' \geq_{st} \Theta$. □

Making use of the above inequality and mimicking the proof of Lemma 3.4, we have

Proposition 3.8 Let $n^*(\Theta)$ be the solution to $\min_{0\leq n\leq N} \bar{\Pi}(n,\Theta)$ for a give Θ. Then,
(1) $\Theta' \geq_{st} \Theta$ implies $n^*(\Theta') \geq n^*(\Theta)$; and in particular,
(ii) $n_0^* \leq n^*(\Theta) \leq n_1^*$, where for $j = 0,1$, n_j^* is the optimal solution to $\min_{0\leq n\leq N} \Pi(n,\theta_j)$.

3.3 Conditional Distribution for Defectives

Let D_n denote the number of defectives uncovered through inspecting n units. We are interested in the conditional distribution of Θ given $D_n = d$. Let $\Theta_n(d) := [\Theta | D_n = d]$. This is the quantity that embodies the sequential nature of the original quality control problem. Therefore, to make preparations for deriving the optimal policy in the next section, here we establish, in the next two lemmas, (a) the monotone properties of $\Theta_n(d)$ with respect to n and d, and (b) K-submodularity properties similar to Proposition 3.7, with Θ replaced by $\Theta_n(d)$.

We want to show next that the likelihood ratio ordering applies to $\Theta_n(d)$, as n and d varies. (Following Lemma 2.3), this then implies the stochastic ordering, which, although weaker, is in this case more cumbersome to prove directly.) For convenience, assume Θ has a density function $f_\Theta(x)$, and denote the density function of $\Theta_n(d)$ as $f_{\Theta_n(d)}(x)$. Then, we have

$$\begin{aligned} f_{\Theta_n(d)}(x) &= \mathsf{P}[\Theta \in dx | D_n = d] \\ &= \frac{\mathsf{P}[D_n = d | \Theta = x] f_\Theta(x)}{\int_\theta \mathsf{P}[D_n = d | \Theta = \theta] f_\Theta(\theta)\, d\theta} \\ &= \frac{x^d(1-x)^{n-d} f_\Theta(x)}{\mathsf{E}[\Theta^d(1-\Theta)^{n-d}]}, \quad x \in [\theta_0, \theta_1]. \end{aligned} \tag{3.8}$$

Lemma 3.9 For all n and $d \le n$, we have

$$\Theta_n(d+1) \ge_{\mathrm{lr}} \Theta_{n+1}(d+1) \ge_{\mathrm{lr}} \Theta_n(d) \ge_{\mathrm{lr}} \Theta_{n+1}(d).$$

In particular, $\Theta_n(d)$ is increasing in d and decreasing in n, both in the sense of the likelihood ratio ordering.

Proof. We prove the third inequality; the other two are similarly proved. From (3.8), we have, for all $x \ge y$,

$$\begin{aligned} \frac{f_{\Theta_n(d)}(x)}{f_{\Theta_{n+1}(d)}(x)} &= \frac{1}{1-x} \cdot \frac{\mathsf{E}[\Theta^d(1-\Theta)^{n+1-d}]}{\mathsf{E}[\Theta^d(1-\Theta)^{n-d}]} \\ &\ge \frac{1}{1-y} \cdot \frac{\mathsf{E}[\Theta^d(1-\Theta)^{n+1-d}]}{\mathsf{E}[\Theta^d(1-\Theta)^{n-d}]} = \frac{f_{\Theta_n(d)}(y)}{f_{\Theta_{n+1}(d)}(y)}, \end{aligned}$$

hence the desired likelihood ratio ordering. □

Lemma 3.10 For all $n < N$ and $d < n$, we have

$$\begin{aligned} & c_r \mathsf{E}[\Theta_n(d)] + \mathsf{E}[\phi(n+1, \Theta_n(d))] - \mathsf{E}[\phi(n, \Theta_n(d))] \\ & \ge c_r \mathsf{E}[\Theta_n(d+1)] + \mathsf{E}[\phi(n+1, \Theta_n(d+1))] \\ & \qquad - \mathsf{E}[\phi(n, \Theta_n(d+1))], \end{aligned} \tag{3.9}$$

and

$$\begin{aligned}
&c_r\mathsf{E}[\Theta_n(d)] + \mathsf{E}[\phi(n+1,\Theta_n(d))] - \mathsf{E}[\phi(n,\Theta_n(d))] \\
&\geq c_r\mathsf{E}[\Theta_{n-1}(d)] + \mathsf{E}[\phi(n,\Theta_{n-1}(d))] \\
&\qquad -\mathsf{E}[\phi(n-1,\Theta_{n-1}(d))]. \qquad (3.10)
\end{aligned}$$

Proof. Since $\Theta_n(d+1) \geq_{st} \Theta_n(d)$, which follows from the likelihood ratio ordering established in Lemma 3.9, the inequality in (3.9) follows from Proposition 3.7.

Similarly, to prove (3.10), from Lemma 3.9, we have $\Theta_{n-1}(d) \geq_{st} \Theta_n(d)$, and hence

$$\begin{aligned}
& c_r\mathsf{E}[\Theta_n(d)] + \mathsf{E}[\phi(n+1,\Theta_n(d))] - \mathsf{E}[\phi(n,\Theta_n(d))] \\
\geq\ & c_r\mathsf{E}[\Theta_{n-1}(d)] + \mathsf{E}[\phi(n+1,\Theta_{n-1}(d))] - \mathsf{E}[\phi(n,\Theta_{n-1}(d))] \\
\geq\ & c_r\mathsf{E}[\Theta_{n-1}(d)] + \mathsf{E}[\phi(n,\Theta_{n-1}(d))] - \mathsf{E}[\phi(n-1,\Theta_{n-1}(d))],
\end{aligned}$$

where the first inequality makes use of Proposition 3.7 and the second one makes use of the decreasing convexity of ϕ in n (Lemma 3.1). □

3.4 Optimal Policy

Let $V_n(d)$ be the expected total remaining cost, following an optimal policy, after n units are inspected and d units are found defective. Then the optimal cost for the original problem is $V_0(0)$.

Let

$$\Phi_n(d) := \mathsf{E}[\phi(n,\Theta_n(d))],$$

and

$$\begin{aligned}
\Psi_n(d) \quad := \quad & c_i + [c_r + V_{n+1}(d+1)]\mathsf{P}[D_{n+1} = d+1 | D_n = d] \\
& + V_{n+1}(d)\mathsf{P}[D_{n+1} = d | D_n = d].
\end{aligned}$$

Clearly, $\Phi_n(d)$ and $\Psi_n(d)$ represent the expected cost associated with the two actions we can take in stage n and state d: either stop inspection (i.e., ship the batch without inspecting the remaining $N-n$ units), or continue inspecting one more unit.

Hence, we have the following recursion:

$$V_n(d) = \min\{\Phi_n(d), \Psi_n(d)\},$$

for $0 \leq n \leq N-1$ and $V_N(d) = \Phi_N(d)$. Furthermore, from standard results in dynamic programming, (e.g., Ross [79], Chapter 1), we know that the optimal policy that minimizes $V_0(0)$ has the following general structure: at each stage n, stop in state d if $\Phi_n(d) < \Psi_n(d)$; continue inspecting more units if $\Phi_n(d) > \Psi_n(d)$; and choose either action when $\Phi_n(d) = \Psi_n(d)$. Below we will reveal more structural properties of the optimal policy and eventually establish its threshold nature.

Observing that

$$\mathsf{P}[D_{n+1} = d+1 | D_n = d] = \mathsf{E}[\Theta_n(d)] = \frac{\mathsf{E}[\Theta^{d+1}(1-\Theta)^{n-d}]}{\mathsf{E}[\Theta^d(1-\Theta)^{n-d}]},$$

we can also express $\Psi_n(d)$ as

$$\Psi_n(d) = c_i + [c_r + V_{n+1}(d+1)]\mathsf{E}[\Theta_n(d)] + V_{n+1}(d)(1 - \mathsf{E}[\Theta_n(d)]).$$

Yet another expression for $\Psi_n(d)$, which will be used below, is:

$$\Psi_n(d) = c_i + c_r\mathsf{E}[\Theta_n(d)] + \mathsf{E}[V_{n+1}(D_{n+1})|D_n = d]. \tag{3.11}$$

For each $0 \le n \le N-1$, define

$$\mathbf{S}_n := \{d : \ 0 \le d \le n, \Phi_n(d) \le \Psi_n(d)\},$$

$$\overline{\mathbf{S}}_n := \{d : \ 0 \le d \le n, \Phi_n(d) \ge \Psi_n(d)\}.$$

That is, $\mathbf{S}_n$ is the set of states in which it is optimal to stop (after inspecting n units), while $\overline{\mathbf{S}}_n$ is the set of states in which it is optimal to continue the inspection. (Note that here $\overline{\mathbf{S}}_n$ is not just the complement of $\mathbf{S}_n$.)

In what follows, we present the structure of the optimal policy in Theorems 3.11 through 3.14, which lead to a statement of the optimal policy in Theorem 3.16.

The first theorem specifies the range for the number of inspected units: the optimal policy must inspect a minimum of n_0^* units and a maximum of n_1^* units, where n_0^* and n_1^* are characterized in Proposition 3.8.

Theorem 3.11 For each n and all $0 \le d \le n$, (i) $d \in \overline{\mathbf{S}}_n$ if $n < n_0^*$, and (ii) $d \in \mathbf{S}_n$ if $n \ge n_1^*$.

Proof. Based on Lemma 3.1 and Propositions 3.6, we know that $\Pi(n,\theta)$ is convex in n and reaches its minimum at $n^*(\theta)$ that falls in between n_0^* and n_1^*. Hence, it is decreasing in n for $n < n_0^*$, i.e., $\Pi(n+1,\theta) \le \Pi(n,\theta)$, for $n < n_0^*$. Hence, replacing θ by $\Theta_n(d)$ and taking expectations, we have $\bar{\Pi}(n+1, \Theta_n(d)) \le \bar{\Pi}(n, \Theta_n(d))$. That is, from (3.7),

$$\begin{aligned} & c_i(n+1) + c_r(n+1)\mathsf{E}[\Theta_n(d)] + \mathsf{E}[\phi(n+1, \Theta_n(d))] \\ \le \quad & c_i n + c_r n\mathsf{E}[\Theta_n(d)] + \mathsf{E}[\phi(n, \Theta_n(d))], \end{aligned}$$

which simplifies to

$$c_i + c_r\mathsf{E}[\Theta_n(d)] + \mathsf{E}[\phi(n+1, \Theta_n(d))] \le \mathsf{E}[\phi(n, \Theta_n(d))] = \Phi_n(d).$$

Observing that

$$\begin{aligned}\mathsf{E}[\phi(n+1,\Theta_n(d))] &= \mathsf{E}[\phi(n+1,\Theta)|D_n=d] \\ &\geq \mathsf{E}[V_{n+1}(D_{n+1})|D_n=d],\end{aligned}$$

where the inequality is implied by the optimality of V_{n+1}, we have

$$\begin{aligned}\Psi_n(d) &= c_i + c_r\mathsf{E}[\Theta_n(d)] + \mathsf{E}[V_{n+1}(D_{n+1})|D_n=d] \\ &\leq c_i + c_r\mathsf{E}[\Theta_n(d)] + \mathsf{E}[\phi(n+1,\Theta_n(d))] \leq \Phi_n(d).\end{aligned}$$

Hence, $d \in \overline{\mathbf{S}}_n$, when $n < n_0^*$.

To show that $d \in \mathbf{S}_n$, for $n \geq n_1^*$, we use induction. Clearly, $V_N(d) = \Phi_N(d)$, i.e., $d \in \mathbf{S}_N$. Suppose $d \in \mathbf{S}_{n+1}$, i.e., $V_{n+1}(d) = \Phi_{n+1}(d)$ for any $d \leq n+1$. We then have,

$$\begin{aligned}\Psi_n(d) &= c_i + c_r\mathsf{E}[\Theta_n(d)] + \mathsf{E}[V_{n+1}(D_{n+1})|D_n=d] \\ &= c_i + c_r\mathsf{E}[\Theta_n(d)] + \mathsf{E}[\phi(n+1,\Theta_{n+1}(D_{n+1}))|D_n=d] \\ &= c_i + c_r\mathsf{E}[\Theta_n(d)] + \mathsf{E}[\phi(n+1,\Theta_n(d))] \geq \Phi_n(d),\end{aligned}$$

where the inequality follows from $\bar{\Pi}(n+1,\Theta_n(d)) \geq \bar{\Pi}(n,\Theta_n(d))$, for $n \geq n_1^*$, which in turn follows from the fact that $\Pi(n,\theta)$ is increasing in n for $n \geq n_1^*$ (since $\Pi(n,\theta)$ is convex in n and reaches its minimum at $n^*(\theta) \leq n_1^*$). This implies $d \in \mathbf{S}_n$. $\square$

The next theorem establishes the following monotone property of the optimal policy: at each stage n, if it is optimal to continue inspection in state d, then it is also optimal to continue in state $d+1$. (This follows intuitively from Lemma 3.9, since state $d+1$ implies a poorer quality, in terms of an increased estimate of the defective rate.) This monotone property then leads to a threshold structure of the optimal policy: once the number of defectives identified, D_n, exceeds a threshold, it is optimal to continue with more inspections.

Theorem 3.12 For each n: $n_0^* \leq n < n_1^*$, and all $d < n$, if $d \in \overline{\mathbf{S}}_n$, then $(d+1) \in \overline{\mathbf{S}}_n$.

Proof. We want to show that for each n: $n_0^* \leq n < n_1^*$, and each $d < n$,

$$\Psi_n(d) \leq \Phi_n(d) \quad \Rightarrow \quad \Psi_n(d+1) \leq \Phi_n(d+1).$$

Instead, we prove via induction a stronger result: that $\Psi_n(d) - \Phi_n(d)$ is decreasing in d, for each given n.

When $n = n_1^* - 1$, since $(n_1^*, d) \in \mathbf{S}_n$, we have

$$\mathsf{E}[V_{n+1}(D_{n+1})|D_n] = \mathsf{E}[\phi(n+1,\Theta_n(d))]$$

in (3.11). Hence,

$$\Psi_n(d) - \Phi_n(d) = c_i + c_r\mathsf{E}[\Theta_n(d)] + \mathsf{E}[\phi(n+1,\Theta_n(d))] - \mathsf{E}[\phi(n,\Theta_n(d))],$$

which is decreasing in d, following (3.9). Consequently, when $n = n_1^* - 1$, $V_n(d) = \Phi_n(d)$ for $d < d^*$, and $V_n(d) = \Psi_n(d) \leq \Phi_n(d)$ for $d \geq d^*$, with d^* defined as

$$d^* := \min\{d \leq n : \psi_n(d) \leq \phi_n(d)\}.$$

Hence, $V_n(d) - \Phi_n(d)$ is also decreasing in d.

Next, consider $n < n_1^* - 1$. As induction hypothesis, assume that $V_{n+1}(d) - \Phi_{n+1}(d)$ is decreasing in d. Then,

$$\begin{aligned}
& \Psi_n(d) - \Phi_n(d) \\
= \;& c_i + c_r \mathsf{E}[\Theta_n(d)] + \mathsf{E}[V_{n+1}(D_{n+1})|D_n = d] - \mathsf{E}[\phi(n, \Theta_n(d))] \\
= \;& \{c_i + c_r \mathsf{E}[\Theta_n(d)] + \mathsf{E}[\phi(n+1, \Theta_n(d))] - \mathsf{E}[\phi(n, \Theta_n(d))]\} \\
+ \;& \{\mathsf{E}[V_{n+1}(D_{n+1})|D_n = d] - \mathsf{E}[\phi(n+1, \Theta_n(d))]\}.
\end{aligned} \tag{3.12}$$

It remains to show that the two parts (in braces) in (3.12) are both decreasing in d. This is obvious for the first part, following (3.9). For the second part, note that the two terms can be written as

$$\mathsf{E}[V_{n+1}(d + I_n(d)) - \Phi_{n+1}(d + I_n(d))|D_n = d], \tag{3.13}$$

where $I_n(d)$ is a 0-1 binary random variable that equals 1 with probability $\mathsf{E}[\Theta_n(d)]$, and hence is stochastically increasing in d (Lemma 3.9). Hence, the required decreasing (in d) property follows from the induction hypothesis, which implies that $V_{n+1}(d) - \Phi_{n+1}(d)$ is decreasing in d. □

Corollary 3.13 For each n: $n_0^* \leq n < n_1^*$, let

$$d_n := \min\{d \leq n : \Psi_n(d) \leq \Phi_n(d)\}.$$

Then, d_n is well defined, with $d \in \mathbf{S}_n$ for $d < d_n$, and $d \in \overline{\mathbf{S}}_n$ for $d_n \leq d \leq n$.

Complementing Theorem 3.12 above, we next show the other half of the monotone property of the optimal policy: if it is optimal to stop at stage n in state $D_n = d$, then it is also optimal to stop at stage $n+1$ in the same state d. Consequently, the threshold values d_n in Corollary 3.13 must be increasing in n.

Theorem 3.14 For each n: $n_0^* \leq n < n_1^*$, and all $d \leq n$, if $d \in \mathbf{S}_n$ then $d \in \mathbf{S}_{n+1}$.

Proof. Similar to the proof of Theorem 3.12, here it suffices to show that $\Psi_n(d) - \Phi_n(d)$ is increasing in n, i.e.,

$$\psi_n(d) - \phi_n(d) \leq \psi_{n+1}(d) - \phi_{n+1}(d) \tag{3.14}$$

for any given $d \leq n$, $n_0^* \leq n < n_1^*$. The induction steps are exactly the same, except here we use (3.10) instead of (3.9). In particular, the first part in (3.12) is increasing in n, following (3.10). To show that (3.13) is also increasing in n, note that it dominates, via the induction hypothesis, $\mathsf{E}[V_n(d+I_n(d)) - \Phi_n(d+I_n(d))|D_n = d]$, which, in turn, dominates

$$\mathsf{E}[V_n(d+I_{n-1}(d)) - \Phi_n(d+I_{n-1}(d))|D_n = d],$$

since $I_n(d)$ is decreasing in n and $V_n(d) - \Phi_n(d)$ is decreasing in d. This yields the desired increasing (in n) property. □

Since $\Psi_n(d) - \Phi_n(d)$ is increasing in n, $d < d_n$ will always implies $d < d_{n+1}$ for any d. Therefore, we have

Corollary 3.15 The d_n values in Corollary 3.13 are increasing in n, i.e., $d_n \leq d_{n+1}$.

Summarizing the results in Theorems 3.11, 3.12 and 3.14, we have

Theorem 3.16 The optimal policy that minimizes $V_0(0)$ is to start from inspecting n_0^* units, continue to inspect one more unit at a time, and stop as soon as the total number of inspected units n satisfies: $D_n < d_n$, or when $n = n_1^*$.

3.5 The Individual Warranty Model

We now return to the individual warranty model introduced in §3.1. Recall in this case the cost function $C(\cdot)$ is additive. From (3.3), we have

$$\phi(n,\theta) = N\mathsf{E}[C(X)] + N\theta(\mathsf{E}[C(Y)] - \mathsf{E}[C(X)]) - n\theta(\mathsf{E}[C(Y)] - \mathsf{E}[C(X)]).$$

(Note here we do not need to assume the convexity of $C(\cdot)$. The decreasingness of $C(\cdot)$, however, is still needed. This ensures that

$$\mathsf{E}[C(Y)] - \mathsf{E}[C(X)] \geq 0,$$

for $X \geq_{st} Y$.) Consequently,

$$\begin{aligned}\Pi(n,\theta) \quad = \quad & N\mathsf{E}[C(X)] + N\theta(\mathsf{E}[C(Y)] - \mathsf{E}[C(X)]) \\ & +c_i n - n\theta(\mathsf{E}[C(Y)] - \mathsf{E}[C(X)] - c_r). \quad (3.15)\end{aligned}$$

When $\mathsf{E}[C(Y)] - \mathsf{E}[C(X)] - c_r > 0$, define

$$\hat{\theta} := \frac{c_i}{\mathsf{E}[C(Y)] - \mathsf{E}[C(X)] - c_r},$$

which guarantees that $\Pi(n,\theta)$ is increasing in n for $\theta \leq \hat{\theta}$ and decreasing in n for $\theta \geq \hat{\theta}$.

Lemma 3.17 Let n^*_θ be the optimal solution for $\min_{0\le n\le N}\Pi(n,\theta)$ of (3.15). Then,
(i) $n^*_\theta = 0$ for any θ when $c_r \ge \mathsf{E}[C(Y)] - \mathsf{E}[C(X)]$;
(ii) $n^*_\theta = 0$ when $c_r < \mathsf{E}[C(Y)] - \mathsf{E}[C(X)]$, and $\theta \le \hat{\theta}$;
(iii) $n^*_\theta = N$ when $c_r < \mathsf{E}[C(Y)] - \mathsf{E}[C(X)]$, and $\theta > \hat{\theta}$.

Proof. $\Pi(n,\theta)$ is increasing in n under the conditions in (i) and (ii), and decreasing in n under the conditions in (iii). □

Note that while the condition, $c_r < \mathsf{E}[C(Y)] - \mathsf{E}[C(X)]$, in (ii) and (iii) of the above lemma, is what (3.5) specializes to here, the condition in (i) goes in the opposite direction. Regardless, however, since $\Pi(n,\theta)$ here is linear in n, in all three cases, there is a meaningful solution. Specifically, provided that $\Theta \equiv \theta$ is a known constant, the optimal (static) policy is either inspect all units in the batch, or not inspect at all, depending on the relationship among the cost and quality data in question. This is consistent with the recommendation of Tapiero and Lee in [99], as well as what is usually followed in practice.

Intuitively, Lemma 3.17 recommends that if the repair cost is high ($c_r \ge \mathsf{E}[C(Y)] - \mathsf{E}[C(X)]$) or if the defective rate is low ($\theta \le \hat{\theta}$), then do no inspection at all; otherwise (i.e., if the repair cost is low *and* the defective rate is high), do 100% inspection. Note the quantity $\hat{\theta}$ plays the role of a threshold for the defective rate.

Next consider the original problem with Θ being a random variable, and the optimal policy that minimizes $V_0(0)$. From (3.15), obviously $\Pi(n,\theta)$ is linear in n and linear in θ. Besides, it is easy to see that, when $c_r \le \mathsf{E}[C(Y)] - \mathsf{E}[C(X)]$, $\Pi(n,\theta)$ is submodular in (n,θ). Consequently, $\bar{\Pi}(n,\Theta)$ is linear in n, linear in $\mathsf{E}[\Theta]$, and submodular in $(n,\mathsf{E}[\Theta])$ if $c_r \le \mathsf{E}[C(Y)] - \mathsf{E}[C(X)]$. Hence, Proposition 3.8 and Theorem 3.11 also hold here. These, along with Lemma 3.17, lead to

Proposition 3.18 (i) When $c_r \ge \mathsf{E}[C(Y)] - \mathsf{E}[C(X)]$, then $n^*_0 = n^*_1 = 0$, and the optimal policy is not to inspect any unit.
(ii) When $c_r < \mathsf{E}[C(Y)] - \mathsf{E}[C(X)]$, there are two cases:

(a) if $\theta_0 \ge \hat{\theta}$, then $n^*_0 = n^*_1 = N$, and the optimal policy is to inspect all N units;

(b) If $\theta_1 \le \hat{\theta}$, then $n^*_0 = n^*_1 = 0$, and the optimal policy is not to inspect any unit.

Note that $\theta_0 \ge \hat{\theta}$ implies $\Theta \ge \hat{\theta}$ and $\theta_1 \le \hat{\theta}$ implies $\Theta \le \hat{\theta}$. Hence, Proposition 3.18 confirms that the optimal rules for a constant Θ in Lemma 3.17 are also optimal in the general setting of a random Θ.

Now, what remains is the most interesting case of $\theta_0 < \hat{\theta} < \theta_1$ under $c_r < \mathsf{E}[C(Y)] - \mathsf{E}[C(X)]$. First note that $n^*_0 = 0$ and $n^*_1 = N$ in this case, following Lemma 3.17. Hence, in principle all N units could be subject to inspection. We can do better, however; in particular, in improving upon n_0^*.

Lemma 3.19 (i) Consider the case of $c_r < \mathsf{E}[C(Y)] - \mathsf{E}[C(X)]$ and $\hat{\theta} \in (\theta_0, \theta_1)$. After inspecting n units, if $\mathsf{E}[\Theta_n(d)] \geq \hat{\theta}$, then it is optimal to continue inspection.
(ii) Denote $\hat{n} := \max\{n : 0 \leq n \leq N, \mathsf{E}[\Theta_n(0)] \geq \hat{\theta}\}$, with the understanding that $\hat{n} = -1$ if the set is empty. Then, it is optimal to inspect at least $n^* := \hat{n} + 1$ units.

Proof. Note that

$$\Phi_n(d) = N\mathsf{E}[C(X)] + (N-n)\mathsf{E}[\Theta_n(d)](\mathsf{E}[C(Y)] - \mathsf{E}[C(X)]),$$

and conditioning upon the quality of the $(n+1)^{st}$ unit, we have

$$\mathsf{E}[\Phi_{n+1}(D_{n+1})|D_n = d] = \Phi_n(d) - \mathsf{E}[\Theta_n(d)](\mathsf{E}[C(Y)] - \mathsf{E}[C(X)]),$$

or

$$\Phi_n(d) = \mathsf{E}[\Theta_n(d)](\mathsf{E}[C(Y)] - \mathsf{E}[C(X)]) + \mathsf{E}[\Phi_{n+1}(D_{n+1})|D_n = d].$$

On the other hand,

$$\begin{aligned} \Psi_n(d) &= c_i + c_r\mathsf{E}[\Theta_n(d)] + \mathsf{E}[V_{n+1}(D_{n+1})|D_n = d] \\ \Psi_{n+1}(D_{n+1})\}| &\leq c_i + c_r\mathsf{E}[\Theta_n(d)] + \mathsf{E}[\Phi_{n+1}(D_{n+1})|D_n = d], \end{aligned}$$

where the inequality is from the definition of $V_{n+1}(\cdot)$. If $\mathsf{E}[\Theta_n(d)] \geq \hat{\theta}$, we have

$$c_i + c_r\mathsf{E}[\Theta_n(d)] \leq \mathsf{E}[\Theta_n(d)](\mathsf{E}[C(Y)] - \mathsf{E}[C(X)]).$$

Therefore, $\Psi_n(d) \leq \Phi_n(d)$. i.e., it is optimal to continue inspection if $\mathsf{E}[\Theta_n(d)] \geq \hat{\theta}$.

Note that $d = 0$ when the inspection starts, and $\mathsf{E}[\Theta_n(0)]$ is decreasing in n, from Lemma 3.9. Since, for any n, $n \leq \hat{n}$ implies $\mathsf{E}[\Theta_n(0)] \geq \hat{\theta}$ and it is optimal to continue inspection, at least $\hat{n}+1$ units should be inspected. This is the conclusion in (ii). □

For n with $n^* \leq n < N$, the optimal decision follows the sequence of thresholds in Corollary 3.13. To summarize, we have

Proposition 3.20 In the case of $c_r < \mathsf{E}[C(Y)] - \mathsf{E}[C(X)]$ and $\hat{\theta} \in (\theta_0, \theta_1)$, the optimal policy is to start from inspecting a sample of n^* units (where n^* follows the definition in Lemma 3.19), and then continue with one unit at a time, and stop as soon as the total number of inspected units n satisfies: $D_n < d_n$, or $n = N$. □

Since here the only requirement for the cost function $C(\cdot)$ is the decreasing property, specializing the function in different ways, we can model a wide range of individual warranty types, e.g., those with free replacement or rebate replacement (e.g., [8], or Lie and Chun [58]).

3.6 Examples and Extensions

Consider a batch of $N = 30$ units. Suppose the defective rate Θ is uniformly distributed between $\theta_0 = 5\%$ and $\theta_1 = 30\%$. The lifetime of a good unit, X, is uniformly distributed on [70, 110]; and the lifetime of a defective unit, Y, is uniformly distributed on [30, 70]. The inspection and repair costs are $c_i = 0.5$ and $c_r = 1$. Suppose there is a cumulative, pro-rata rebate warranty of the type in (3.2) associated with the batch, with the unit price $c = 100$ and a warranty period $W = 82$ per unit.

Following Theorem 3.16, the following thresholds can be computed:

$$d_n = 0,\ n \leq 8; \quad d_9 = d_{10} = d_{11} = 1; \quad d_{12} = d_{13} = d_{14} = 2;$$

$$d_{15} = d_{16} = 3; \quad d_{17} = 4; \quad d_{18} = 5; \quad d_{19} = 6; \quad d_{20} = 8; \quad d_{21} = 12.$$

These thresholds guide the inspection at each step. For instance, as long as there is at least one defect among 11 inspected units, inspection should continue; whereas inspection can be terminated if there are fewer than 4 defectives among 17 inspected units. Under no circumstances should inspection continue beyond $n_1 = 22$ inspected units ($n_0 = 0$).

Following this optimal policy, the expected number of inspected units is 14.79, uncovering an average of 2.76 defective units. The total expected cost is 12.964. In contrast, the expected costs under zero-inspection and full-inspection policies are 36.41 and 20.25, respectively.

As a second example, consider the individual warranty model. Use the above data, and apply the pro-rata rebate warranty to each individual unit. The only change is to set the unit price at $c = 11$. The computed thresholds are:

$$d_n = 0,\ n \leq 7; \quad d_n = 1,\ 8 \leq n \leq 12; \quad d_n = 2,\ 13 \leq n \leq 18;$$

$$d_n = 3,\ 19 \leq n \leq 23; \quad d_n = 4,\ 24 \leq n \leq 28; \quad d_{29} = 5.$$

(Here, $n_0 = 0$ and $n_1 = 30$.) This optimal policy yields an expected number of 20.73 inspected units, with an average of 4.01 identified defective units. The total expected cost is 26.65. This is much closer (than the case of cumulative warranty) to the expected costs under zero-inspection and full-inspection policies: 28.51 and 27.49, respectively. (When the unit price is $c = 100$, this falls into the special case of Proposition 3.18 (i), and the optimal policy is to do full inspection.)

Imperfect inspection and/or repair can be easily incorporated into our model. The only change is to modify the distributions of X and Y, so that each inspected (and repaired) unit will have a lifetime that is a mixture of X and Y. This will not affect the structure of the model or the form of the optimal policy established earlier. (Note here we assume an imperfect inspection will only identify a defective unit with a certain probability, but will never mistake non-defective units as defective. Similarly, we assume an imperfect repair will only transform a defective unit into a non-defective unit with a certain probability, but will never make a non-defective unit defective. If these assumptions do not hold, obviously the optimal policy will be structurally different. For instance, it might become justified *not* to repair a unit that is identified as defective by the inspection.)

The imperfection of inspection and repair will be reflected in the threshold values. For instance, if the inspection and repair are very ineffective (in identifying and correcting defective units), then the threshold value n_0^* will be low, while d_n's will be high, so that the optimal policy will stop inspection early on, or conduct no inspection at all.

The random variables, X and Y, can be replaced by random vectors without affecting any structure of the optimal policy. The random vectors can model, for instance, the so-called two-dimensional warranties in [91]. That is, in addition to lifetime, there is a second dimension that explicitly accounts for usage (e.g., six years or six hundred thousand miles).

3.7 Notes

In the literature of statistical quality control (or, statistical process control), there are several approaches to quality assurance through inspection and repair. The simplest one is to do some back-of-envelope calculation based on the given cost data and the *average* defective rate ($\mathsf{E}[\Theta]$), and then choose from two actions: either do 100% inspection on the batch or do no inspection at all. A better, but more involved approach is to inspect a small sample (of size n, say), and if the number of defectives identified exceeds the expected value, $n\mathsf{E}[\Theta]$ by a certain number of "sigma" (standard deviation), then inspect all the remaining units; otherwise, stop inspection and accept the batch. The most sophisticated existing approach is perhaps the CUSUM technique (refer to, e.g., Thompson and Koronacki [102], Chapter 4): Inspect the units one at a time; at each step update the cumulative sum of the log-likelihood estimate of the defective rate; continue if the sum falls within a prespecified interval, say $[\alpha, \beta]$; stop inspection and ship the whole batch, if the sum falls below α; and inspect all the remaining units if the sum exceeds β.

Common to all these approaches is that the *policy* is pre-specified; the issue then becomes essentially a parametric design problem of finding one or two threshold values: the breakeven point in the back-of-envelope analysis, the upper limit on the acceptable number of defective units along with the sample size in the acceptance-rejection approach, and the α and β values in the CUSUM technique.

In contrast, our focus here (also refer to [20]) is on identifying a *policy* that is provenly optimal, instead of merely finding optimal parameters of pre-specified policies. Our approach is based on dynamic programming, or more specifically, sequential analysis. The basic idea of sequential analysis (Wald [105]), in our context, calls for identifying a procedure (i.e., policy) by which items are inspected one by one, and each time after inspecting a unit, we decide either to stop inspection or to continue, depending on the outcome of the inspection up to that point.

There is an extensive body of literature that studies the statistical, social-economical and behavioral aspects of warranty; see, e.g., the survey articles of Blischke [8], and Singpurwalla and Wilson [91]. There are also studies on various optimization issues that arise in warranties and related services, e.g., Mamer [60], Djamaludi, Murthy and Wilson [33], Murthy and Nguyen [64], Murthy, Djamaludi and Wilson [65], Nguyen and Murthy [66], and Thomas [101]. The model studied in Tapiero and Lee [99] is quite similar to the individual warranty model in §3.5. For the static optimization problem, i.e., with $\Theta \equiv \theta$, a given constant, Tapiero and Lee demonstrated that the optimal policy is either 0% or 100% inspection. They also pointed out that when Θ is a random variable, the optimal policy will be different from these extreme-point rules. This optimal policy is completely characterized in §3.5 here, as a special case of our general model.