Proc. of the European Safety and Reliability Conference (ESREL 2026)
Edited by José C. Matos, Paulo B. Lourenço, Michael Beer, Daniel Oliveira,
Edoardo Patelli, Hélder S. Sousa, Rui A. Silva, Monica Santamaria and Chiara Turco
© 2026 ESREL 2026 Organizers. *Published by* Research Publishing, Singapore.
ISBN: 978-981-94-6718-1 | doi: 10.3850/ESREL2026061419_esrel26-p26458-cd

On the value of information from inspections in condition-based maintenance of degrading units

Nicola Esposito

Scuola Superiore Meridionale, Italy. E-mail: n.esposito@ssmeridionale.it

Christophe Bérenguer

Université Grenoble Alpes, Grenoble INP, CNRS, GIPSA-Lab, Grenoble, France. E-mail: christophe.berenguer@grenoble-inp.fr

Massimiliano Giorgio

Università degli studi di Napoli Federico II, Italy. E-mail: Massimiliano.giorgio@unina.it

This paper addresses the problem of investigating the value of information gained via inspections in condition-based maintenance of degrading units. Namely, the maintenance model is constructed by assuming that the degradation process of the units of interest can be described by a gamma process with random effect. It is also assumed that inspections allow to know the exact level of degradation of the units. The study is performed by using a novel aperiodic sequential maintenance policy, where at each inspection a decision is made about whether to immediately replace the unit, to define its future replacement time, or to plan another inspection at a future time. All decisions are made adaptively based on the distribution of the remaining useful life, conditional to all the available measurements. The optimal policy is determined by adopting as performance index the long-run average cost rate. The study is replicated by considering several scenarios characterized by different values of inspection cost, logistic costs, and degradation model parameters. Under each scenario, once determined the optimal policy, the value of inspections is evaluated by: a) computing the probability that the decision to replace the unit will be made by performing only one or two inspections at most and b) evaluating the difference between the values of the long-run average cost rate obtained under the optimal policy, which allows for multiple inspections, and those obtained by using similar policies that allow either for only one or two inspections at most.

Keywords: Condition-based maintenance, Value of Information, random effect, adaptive decision-making, long-run average cost rate.

1. Introduction

Condition-based maintenance (CBM) strategies have become central in modern asset management, offering the possibility to reduce operational costs and improve system reliability by tailoring maintenance actions to the actual condition of individual units. A crucial aspect for a successful implementation of CBM is the acquisition and exploitation of information regarding the degradation state of the units, typically obtained through inspections or other condition monitoring techniques.

Richer information usually enables more accurate predictions of remaining useful life (RUL) and consequently more efficient maintenance decisions, but it also incurs additional costs. Being able to determine the best trade-off between cost and benefit of information obtained via inspections (i.e., its value) in terms of effectiveness of the decision process is an aspect of pivotal importance in maintenance optimization.

The implementation of CBM also relies on a mathematical model that describes the degradation process. Obviously, the suitability of the adopted model is a key for a successful implementation of CBM.

In this paper we focus on situations where the degradation process of the units of interest exhibits both temporal variability (i.e., the uncertainty associated with the progression over time of the degradation of a specific unit) and unit-to-unit variability (i.e., the uncertainty associated with the

difference existing between the evolution of degradation of different units).

We assume that the degradation process of the units is not decreasing and that it can be suitably described by the gamma process with random effect firstly proposed by (Lawless and Crowder 2004), which has proven to be suitable to model the degradation process of many real-world technological units.

Here, we adopt this model because it also has good mathematical tractability and very high flexibility, allowing for easy control and tuning of the amount of temporal variability and random effect (which describes the unit-to-unit variability).

The main objective of the paper is investigating to what extent the information gathered via inspections improve the effectiveness of a CBM policy and how the value of this information depends on the features of the degradation process and on inspection costs.

To this aim, we use a novel aperiodic sequential policy, where at each inspection time it is possible to decide whether to immediately replace the unit or to define an adaptive (i.e., based on the available measurements) intervention time where either a replacement or another inspection can be scheduled, depending on convenience. The optimal policy is determined by minimizing the long-run average cost rate.

Once determined the optimal policy, the value of inspections is evaluated by computing: a) the probability that the decision of replacing the unit is made by performing either one or two inspections only and b) the difference between the values of the long-run average cost rate obtained under the optimal policy which allows for a potentially unlimited number of inspections and the ones obtained by using similar policies that rely on only one or two inspections at most.

To investigate how the result depends on the experimental scenario, the analysis is replicated by considering different values of inspection cost, logistic cost, and degradation model parameters.

The rest of the paper is organized as follows. Section 2 details the adopted gamma process with random effect and defines the RUL. Section 3 describes the novel sequential maintenance policy, while Section 4 illustrates the formulation of the cost function. Section 5 reports the results of the numerical analysis. Section 6 concludes the paper.

2. The degradation model

The degradation model, say $\{W(t); t \geq 0\}$, we consider in this paper is the gamma process with random effect firstly proposed by (Lawless and Crowder 2004). This model describes the behavior of units pertaining to a heterogeneous population via a gamma process where the shape function is fixed and the scale parameter fluctuates randomly from unit to unit according to a gamma distribution. Hereinafter, we denote this random scale parameter by Λ and its realization by λ.

Under this setting, the conditional probability density function (pdf) of the increment $\Delta W(t, t+\Delta t) = W(t+\Delta t) - W(t)$ of the process $\{W(t); t \geq 0\}$, given $\Lambda = \lambda$, can be expressed as:

$$f_{\Delta W(t,t+\Delta t)|\Lambda}(\delta|\lambda) = \frac{\lambda^{\Delta\eta(t,t+\Delta t)} \cdot \delta^{\Delta\eta(t,t+\Delta t)-1}}{\Gamma\big(\Delta\eta(t,\ t+\Delta t)\big)} \cdot e^{-\lambda\cdot\delta},\ \ \delta > 0. \quad (1)$$

The pdf of Λ can be expressed as:

$$f_{\Lambda}(\lambda) = \frac{c^d \cdot \lambda^{d-1}}{\Gamma(d)} \cdot e^{-c\cdot\lambda},\ \ c, d, \lambda > 0, \quad (2)$$

where $\Gamma(y) = \int_0^\infty u^{y-1} \cdot e^{-u} \cdot \mathrm{d}u$ is the complete gamma function, $\Delta\eta(t,\ t+\Delta t) = \eta(t+\Delta t) - \eta(t)$, $\eta(t)$ is a monotonic increasing function (referred to as the age function). The initial condition at $t = 0$ is $W(0) = 0$.

The resulting marginal process is Markovian (Esposito et al. 2023). Hence, the functions (1) and (2), together with the initial condition, provide a complete probabilistic description of the process.

From Eqs. (1) and (2), the marginal pdf of $W(t)$ is:

$$f_{W(t)}(w) = \frac{c^d \cdot w^{\eta(t)-1}}{B(\eta(t), d) \cdot (w+c)^{\eta(t)+d}},\ \ w \geq 0, \quad (3)$$

and the corresponding marginal cumulative distribution function (cdf) is:

$$F_{W(t)}(w) = \mathbb{B}\left(\frac{w}{w+c};\ \eta(t), d\right),\ \ w \geq 0, \quad (4)$$

Where:

$$\mathbb{B}(z; \alpha, \beta) = \frac{\int_0^z u^{\alpha-1} \cdot (1-u)^{\beta-1} \cdot \mathrm{d}u}{B(\alpha, \beta)} \quad (5)$$

is the regularized beta function and

$$B(\alpha, \beta) = \frac{\Gamma(\alpha) \cdot \Gamma(\beta)}{\Gamma(\alpha+\beta)} \quad (6)$$

is the beta function. The mean and variance of $W(t)$ are equal to:

$$E\{W(t)\} = \frac{c \cdot \eta(t)}{d-1}, \quad (7)$$

$$V\{W(t)\} = \frac{c^2 \cdot \eta(t) \cdot [\eta(t) + d - 1]}{(d-1)^2 \cdot (d-2)}, \quad (8)$$

respectively. The mean exists only if $d > 1$ and the variance only if $d > 2$.

The conditional pdf of the degradation increment $\Delta W(t,\ t+\Delta t)$, given $W(t) = w$, can be formulated as:

$$f_{\Delta W(t,t+\Delta t)|W(t)}(\delta|w) = \frac{1}{B(\Delta\eta(t,t+\Delta t), \eta(t)+d)} \times \times \frac{(w+c)^{\eta(t)+d} \cdot \delta^{\Delta\eta(t,t+\Delta t)-1}}{(w+\delta+c)^{\eta(t+\Delta t)+d}}, \quad (9)$$

and the conditional cdf:

$$F_{\Delta W(t,t+\Delta t)|W(t)}(\delta|w) = \mathbb{B}\left(\frac{\delta}{\delta+w+c};\ \Delta\eta(t,t+\Delta t), \eta(t)+d\right). \quad (10)$$

Finally, from Eq. (9), it follows that the conditional pdf of $W(t+\Delta t)$, given $W(t) = w$, can be expressed as:

$$f_{W(t+\Delta t)|W(t)}(w_{t+\Delta t}|w) = \frac{1}{B(\Delta\eta(t,t+\Delta t), \eta(t)+d)} \times \times \frac{(w+c)^{\eta(t)+d} \cdot (w_{t+\Delta t}-w)^{\Delta\eta(t,t+\Delta t)-1}}{(w_{t+\Delta t}+c)^{\eta(t+\Delta t)+d}}, \quad (11)$$

and from Eq. (10), the corresponding conditional cdf is:

$$F_{W(t+\Delta t)|W(t)}(w_{t+\Delta t}|w) = \mathbb{B}\left(\frac{w_{t+\Delta t}-w}{w_{t+\Delta t}+c};\ \Delta\eta(t,t+\Delta t), \eta(t)+d\right). \quad (12)$$

Details about the derivations of these equations are given in (Esposito et al. 2023).

2.1. *The definition of failure*

In this paper, a unit is assumed to fail when its degradation level firstly passes a fixed failure threshold, say w_M, The useful life (i.e., the time to failure) X of the unit is defined as:

$$X = \inf\{x: W(x) > w_M\} \quad (13)$$

In the rest of this section, some results involving the lifetime X, needed to compute maintenance costs in the rest of the paper, are derived.

Let t_a and t_b denote two generic reference times, with $t_b > t_a > 0$. Then, exploiting the monotonic increasing property of $\{W(t); t \geq 0\}$, the conditional distribution of X given $W(t_b) = w_b$, in the case $w_b \leq w_M$ and $x > t_b$ is:

$$F_{X|W(t_b)}(x|w_b) = P(X \leq x|W(t_b) = w_b) = P(W(x) > w_M|W(t_b) = w_b) = 1 - F_{W(x)|W(t_b)}(w_M|w_b) \quad (14)$$

Likewise, by using Eqs. (3), (5), and (9), and some additional manipulations (Piscopo et al. 2024), the distribution of X given $W(t_a) = w_a$ and $W(t_b) = w_b$, in the case $w_a \leq w_M, w_b > w_M$, and $t_a < x \leq t_b$, can be expressed as:

$$F_{X|W(t_a),W(t_b)}(x|w_a, w_b) = P(X \leq x|W(t_a) = w_a, W(t_b) = w_b) = P(W(x) > w_M|W(t_a) = w_a, W(t_b) = w_b) = 1 - \mathbb{B}\left(\frac{w_M - w_a}{w_b - w_a}; \Delta\eta(t_a, x), \Delta\eta(x, t_b)\right) \quad (15)$$

3. The maintenance policy

The proposed maintenance policy is developed by assuming that the unit of interest is randomly selected from a heterogeneous population and that its degradation process can be described by the gamma process with random effect $\{W(t); t \geq 0\}$ presented in Section 2.

It is assumed that failures are not self-announcing (Bismut et al. 2022, Bautista et al. 2022), that failed units can continue to operate, and that operating failed units incurs additional cost. Corrective and preventive replacements are supposed to restore the unit to an "as good as new" state, and therefore time between successive replacements defines the cycle of a renewal process.

The sequential policy consists in scheduling maintenance inspection/intervention times, denoted as τ_j (where $j = 1, 2, ...$).

At inspection times an (instantaneous, non-destructive) inspection is performed, which provides the exact value of the degradation level of the unit.

If the inspection reveals that the unit is failed, then it is immediately (correctively) replaced. Otherwise, a condition-based rule is used to decide whether to preventively replace the unit at the current time, to postpone its replacement to a future time, or to inspect again the unit at the same future time. The procedure is applied sequentially, up to the replacement of the unit.

The condition-based rule is described in Table 1, where $U_j(X)$ is a reward function that

quantifies the economic utility of operating the unit in the interval (τ_j, τ_{j+1}), c_l is a logistic cost (i.e., a cost incurred every time a maintenance action is performed due to, for example, the cost of setting up a maintenance operation) and c_i is the cost of an inspection.

Table 1. Condition-based rule used at τ_j in case $w_j \leq w_M$.

Expected reward	Decision
$E\{U_j(X)\|W(\tau_j) = w_j\} \geq c_l + c_i$	Inspect at τ_{j+1}
$c_l \leq E\{U_j(X)\|W(\tau_j) = w_j\} < c_l + c_i$	Replace at τ_{j+1}
$E\{U_j(X)\|W(\tau_j) = w_j\} < c_l$	Replace at τ_j

The time τ_{j+1} is set to the value that satisfies the following equality:

$$F_{X|W(\tau_j)}(\tau_{j+1}|w_j) = q, \tag{16}$$

so that τ_{j+1} is the quantile of order q of the conditional Cdf of the useful life X, given $W(\tau_j) = w_j$.

The expression of $U_j(X)$ is given in Table 2, where c_c is the cost of a corrective replacement, c_p ($c_c > c_p$) is the cost of a preventive replacement, and c_r is the reward rate, which serves to evaluate the utility of operating an unfailed unit.

Table 2. Expression of the reward function $U_j(X)$.

X	$U_j(X)$
$X > \tau_{j+1}$	$c_r \cdot (\tau_{j+1} - \tau_j)$
$X \leq \tau_{j+1}$	$-(c_c - c_p) + c_r \cdot (X - \tau_j)$

The conditional expectation of $U_j(X)$ given $W(\tau_j) = w_j$ (where $w_j \leq w_M$) is:

$$\begin{aligned} &E\{U_j(X)|W(\tau_j) = w_j\} \\ &= \int_{\tau_j}^{+\infty} U_j(x) \cdot f_{X|W(\tau_j)}(x|w_j) \cdot \mathrm{d}x \\ &= \int_{\tau_j}^{\tau_{j+1}} \left(-(c_c - c_p) + c_r \cdot (x - \tau_j)\right) \times \\ &\quad \times f_{X|W(\tau_j)}(x|w_j) \cdot \mathrm{d}x \\ &+ \int_{\tau_{j+1}}^{\infty} c_r \cdot (\tau_{j+1} - \tau_j) \cdot f_{X|W(\tau_j)}(x|w_j) \cdot \mathrm{d}x, \end{aligned}$$

that, integrating by parts, and using the (14), can be written as:

$$\begin{aligned} &E\{U_j(X)|W(\tau_j) = w_j\} = \\ &= c_r \cdot \int_{\tau_j}^{\tau_{j+1}} F_{W(x)|W(\tau_j)}(w_M|w_j) \cdot \mathrm{d}x \\ &-(c_c - c_p) \cdot \left(1 - F_{W(\tau_{j+1})|W(\tau_j)}(w_M|w_j)\right) \end{aligned}$$

Basing the decision-making rule on the utility function $U_j(X)$ establishes a threshold structure for maintenance planning. Consequently, planning an inspection at time τ_{j+1} is undertaken only if the expected utility is greater than the cost of the inspection (i.e., $c_l + c_i$). Alternatively, if the expected utility is smaller than $c_l + c_i$ but greater than c_l, then performing an additional inspection is not justified cost-wise and the policy opts for a replacement instead of the inspection at the same future time τ_{j+1}.

Otherwise, if the expected utility is smaller than c_l, an immediate preventive replacement is performed at τ_j.

In addition, the expectation $E\{U_j(X)|W(\tau_j) = w_j\}$ depends both on the current degradation measurement w_j and on the intervention times τ_j and τ_{j+1}. Consequently, the maintenance decision is, implicitly, age- and state-dependent.

However, it is worth remarking that this condition-based rule is still based on a heuristic and therefore does not guarantee optimality. The sequential procedure is initialized by setting

$$\tau_1\colon\ F_X(\tau_1) = q.$$

The order of the quantile q is the only design parameter of the policy.

4. Formulation of the cost function

Maintenance costs are computed including the cost of a preventive replacement c_p, the cost of a corrective replacement c_c, the inspection cost c_i, the logistic cost c_l, the reward gained by operating an unfailed unit, and the downtime cost incurred by operating a failed unit. Reward and downtime cost are computed by multiplying fixed rates (i.e., c_r and c_d) by uptime and downtime respectively, where the uptime is the amount of time an unfailed unit is operated and the downtime is the time a failed unit is operated before being replaced. In this paper, for the sake of simplicity, yet without loss of generality, we assume that $c_d = c_r$ (assumption that has already been implicitly used to define the functional form of $U_j(X)$ in Table 2).

Table 3. Possible scenarios and associated cost at each decision made at time τ_j ($j = 0,1,2,\ldots,N$).

Expected utility	Lifetime X	ΔC_j
$E\{U(x)\|W(\tau_j) = w_j\} \leq c_l$	$-$	c_p
$c_l \leq E\{U(x)\|W(\tau_j) = w_j\} < c_l + c_i$	$\tau_j \leq X < \tau_{j+1}$	$c_l + c_c + c_d \cdot (\tau_{j+1} - X)$
$c_l \leq E\{U(x)\|W(\tau_j) = w_j\} < c_l + c_i$	$X \geq \tau_{j+1}$	$c_l + c_p$
$E\{U(x)\|W(\tau_j) = w_j\} > c_l + c_i$	$\tau_j \leq X < \tau_{j+1}$	$c_l + c_i + c_c + c_d \cdot (\tau_{j+1} - X)$
$E\{U(x)\|W(\tau_j) = w_j\} > c_l + c_i$	$X \geq \tau_{j+1}$	$c_l + c_i$

The performance measure adopted to define the optimal policy is the long-run average cost rate, computed via the renewal/reward theorem (see Ross 1983) as:

$$C_\infty(q) = \frac{E\{C(q)\}}{E\{T(q)\}} \tag{17}$$

Where $C(q)$ is the total cost incurred over a maintenance cycle and $T(q)$ is the length of the maintenance cycle. $C(q)$ is computed as:

$$C(q) = \sum_{j=0}^{N} \Delta C_j \tag{18}$$

where ΔC_j represents the cost incurred in the interval of time (τ_j, τ_{j+1}), N is the number of inspections carried out on the unit, and $\tau_0 = 0$.

$T(q)$ is the time at which the unit is replaced. So, for example, if the unit is replaced at the N-th inspection, then $T(q) = \tau_N$, whereas if at the N-th inspection the policy opts for postponing the replacement to τ_{N+1}, then $T(q) = \tau_{N+1}$. The value of ΔC_j corresponding to each possible scenario is reported in Table 3, where $\tau_0 = w_0 = 0$. The table applies only when $w_j \leq w_M$. If $w_j > w_M$ (a situation that can occur only when $j \geq 1$) the unit is subjected to a corrective replacement at τ_j. In Table 3, costs incurred due to corrective replacement at τ_j are included in those determined by the decision taken at τ_{j-1}.

The optimal policy is the one that minimizes the long-run average cost rate in Eq. (17). Hence, the optimal value of q, denoted as q^*, is:

$$q^* = \underset{q}{\operatorname{argmin}}\, C_\infty(q)$$

and the corresponding optimal cost is $C_\infty^* = C_\infty(q^*)$.

4.1. *Description of the simulator*

Due to the complexity arising from the formulation of the policy, we compute maintenance costs in Eq. (17) via Monte Carlo simulation.

The expectations are computed as $E\{C(q)\} = \sum_{k=1}^{K} C[k]/K$ and $E\{T(q)\} = \sum_{k=1}^{K} T[k]/K$, where K is the total number of simulated runs and $C[k]$ and $T[k]$ are the cost and cycle length in the k-th run. The pseudocode used to compute $C[k]$ and $T[k]$ is detailed in Table 4.

Obviously, the value of K should be determined balancing numerical accuracy and computational time, which both increase with K. In this paper, we settled on $K = 5 \cdot 10^5$, which allows us to obtain an estimate of the long-run average cost rate stable up to (at least) the fourth significant figure.

Table 4. Pseudocode used to compute $C_\infty(q)$

	Simulator
1	**For** $k = 1$ **to** K
2	$w_0 \leftarrow 0$, $\tau_0 \leftarrow 0$, $j = 1$, $T[k] \leftarrow 0$, $C[k] \leftarrow c_l + c_i$
3	τ_1: $F_X(\tau_1) = q$
4	**While** $true$
5	$\Delta w \leftarrow$*sample from* $\Delta W(\tau_{j-1}, \tau_j)$
6	$w_j \leftarrow w_{j-1} + \Delta w$
7	**If** $w_j > w_M$ **then**
8	$x \leftarrow$*sample from* $X\|W(\tau_{j-1}), W(\tau_i)$
9	$C[k] \leftarrow C[k] + c_c + c_d \cdot (\tau_j - x)$
10	$T[k] = \tau_j$
11	**Go to** line 32
12	**Else**
13	τ_{j+1} : $F_{X\|W(\tau_j)}(\tau_{j+1}\|w_j) = q$
14	$\beta \leftarrow E\{U_j(X)\|W(\tau_j) = w_j\}$
15	**If** $\beta \leq c_l$ **then**
16	$C[k] \leftarrow C[k] + c_p$
17	$T[k] \leftarrow \tau_j$

Go to line 32
ElseIf $c_l < \beta \leq c_l + c_i$
  $\Delta w \leftarrow$*sample from* $\Delta W(\tau_j, \tau_{j+1})$
  $w_{j+1} \leftarrow w_j + \Delta w$
  $T[k] \leftarrow \tau_{j+1}$
  If $w_{j+1} > w_M$ **then**
    $x \leftarrow$*sample from* $X|W(\tau_j), W(\tau_{j+1})$
    $C[k] \leftarrow C[k] + c_l + c_c + c_a \cdot (\tau_{j+1} - x)$
  Else
    $C[k] \leftarrow C[k] + c_l + c_p$
  Go to line 32
Else
  $C[k] \leftarrow C[k] + c_l + c_i$
$j \leftarrow j + 1$
$k \leftarrow k + 1$

5. Numerical example

In this section, we report the results of the numerical study we have conducted to investigate how the proposed policy behaves under four scenarios defined by setting in as many different ways the parameters of the cost and degradation models. The aim of the study is evaluating how sensitive the number of inspections is with respect to the inspection cost c_i, the logistic cost c_l, and to the relative amount of temporal variability and random effect. In this example, we assumed the shape of the age function to be $\eta(t) = a \cdot t$.

The values of the parameters used to define the four setups, denoted by the letters A, B, C, and D, are provided in Table 5. The values of the remaining parameters of the cost model, which are common to all scenarios, are $c_p = 2$, $c_c = 6$, $c_r = c_d = 0.05$.

Table 5. Parameters used to define the four Setups.

Setup	a	c	d	c_l	c_i
A	0.654	6.934	11.80	0.5	0.2
B				0.02	0.05
C	0.101	437.6	106.3	0.5	0.2
D				0.02	0.05

Setups A and B are characterized by high unit-to-unit variability and small temporal variability, while the opposite occurs in the Setups C and D. Within these two pairs, Setups differ in inspection and logistic costs, that are high in Setups A and C and low in Setups B and D.

It is worth remarking that the parameters used to define the Setups reported in Table 5 allow the resulting degradation models to have the same mean function and comparable variance functions.

For illustrative purposes, Figure 1 presents ten simulated degradation paths under the Setups A and C, shown in the upper and lower panels, respectively. These Setups differ only in the relative amounts of unit-to-unit and temporal variability, and the figure is intended to compare these two variability structures (Setups B and D are not shown, as these share the same degradation process parameters with Setups A and C, respectively).

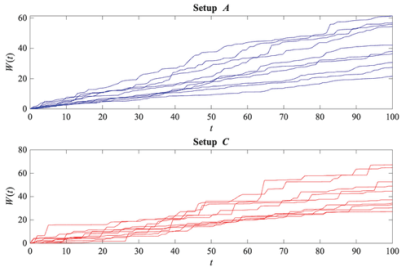


Fig. 1. Ten simulated degradation paths under the Setup A (upper panel) and the Setup C (lower panel).

Figure 1 illustrates that, under the Setup A (characterized by high unit-to-unit variability and low temporal variability), the degradation paths intersect much more infrequently then under Setup C (where temporal variability is high and unit-to-unit variability is low).

To assess the value of the inspections, we compared the performance of the proposed aperiodic sequential policy (hereinafter denoted by P_S) with similar policies that use one or two inspections at most (hereinafter denoted by P_1 and P_2, respectively). These policies are obtained from P_S by constraining the maximum number of inspections to one and two, respectively.

The constraint is imposed by modifying the decision rule used at τ_H (where H is equal to 1 for P_1 and 2 for P_2) as in Table 6, where w_H is the state of the unit at inspection time τ_H. Every other element of the structure of the policy remains the same as under policy P_S.

Table 6. Condition-based rule used at τ_H, if $w_H \leq w_M$.

Expected reward	Decision
$E\{U_H(X)\|W(\tau_H) = w_H\} \geq c_l$	Replace at τ_{H+1}
$E\{U_H(X)\|W(\tau_H) = w_H\} < c_l$	Replace at τ_H

Table 7 reports the minimum long-run average cost rate obtained under policies P_1, P_2, and P_S for each considered Setup. Similarly, Table 8 reports the optimal value q^* of the sole design parameter q of the considered policies.

Table 7. Values of the minimum long-run average cost rate for each policy and each Setup.

Setup	P_1	P_2	P_S
A	4.550	4.651	4.679
B	3.233	2.988	2.988
C	5.458	5.514	5.518
D	4.120	3.947	3.837

Table 8. Value of q^* under each policy and Setup.

Setup	P_1	P_2	P_S
A	0.053	0.053	0.053
B	0.038	0.022	0.022
C	0.131	0.147	0.147
D	0.069	0.038	0.022

Table 7 shows that the policy P_S may or may not outperform the policies P_1 and P_2 depending on the Setup. In fact, under the Setup A, where c_i and c_l are high, unit-to-unit variability is high and the temporal is low, the policy P_1 provides the best performance (i.e., lower optimal cost).

Moving to Setup B, where the inspection and logistic costs decrease, performing additional inspections becomes more appealing. However, even in this case, it does not seem to be economically beneficial to perform more than two inspections, as the policy P_2 and P_S yield (at least up to the fourth significant figure) the same optimal long-run average cost rate.

The one-inspection policy P_1 provides the best performances also under Setup C, where c_l and c_i are high, temporal variability is high and unit-to-unit variability is low. Conversely, under the Setup D, where inspection and logistic costs are low, temporal variability is low, and unit-to-unit variability is high, the policy P_S yields the lowest optimal cost, followed in order by P_2 and P_1.

These results can be intuitively explained by the properties of the adopted degradation model, the gamma process with random effect. In fact, under this model, if unit-to-unit variability is high and temporal variability is low, the unit-specific degradation paths become highly predictable, once the value of the unit-specific scale parameter is known. Of course, this value cannot be directly observed, but especially when temporal variability is low (see Esposito et al. 2023), even a single degradation measurement can give very good information about its value. Therefore, the one-inspection policy P_1 can define adaptive replacement times that are well aligned with the predicted degradation behavior. Additional inspections, especially when c_l and c_i are high, do not provide meaningful economic benefits.

On the other hand, when temporal variability is high and unit-to-unit variability is low, under Setups C and D, the degradation paths become less predictable and performing many inspections provides more valuable information. Hence, when c_l and c_i are low, the multiple-inspections policy P_S becomes more convenient.

To confirm these insights, let us now focus on the number N of inspections that the policy P_S carries out in a renewal cycle. Given the structure of the policy P_S, N is a discrete random variable whose distribution depends, for a given setup, also on q. Figure 2 reports, in the upper and lower panel, respectively, the probability mass function (pmf) of N in the case of Setups A and B when q is set to its optimum value q^* (see Table 8). Similarly, Figure 3 reports in the upper and lower panel the pmf of N in the case of Setups C and and D, respectively when $q = q^*$. Figures 2-3 confirm the aforementioned insights. In fact, when c_l and c_i are high (i.e., Setups A and C), the policy P_S will very frequently carry out only one or two inspections, going up to three or four in very limited occasions. Conversely, when c_l and c_i are low, P_S will resort to significantly more inspections.

Moreover, comparing the pmf of N under Setups B and D, we can see that P_S performs many more inspections under Setup D than under Setup B, even though the values of c_l and c_i do not change between these two Setups.

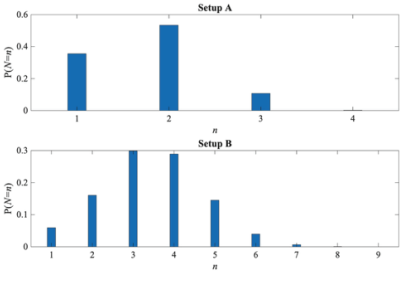


Fig. 2. Pmf of N under Setup A and B.

As already explained, this is because when temporal variability is high and the unit-to-unit variability is low (as under Setup D), the degradation path of the unit is less predictable, and thus it is more convenient to inspect the unit more than once.

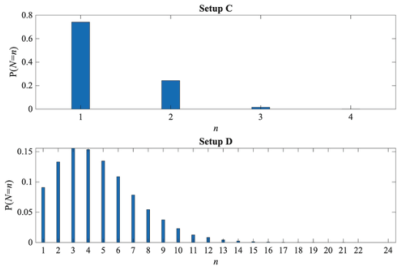


Fig. 3. Pmf of N under Setup C and D.

This numerical analysis also highlights a limitation of the proposed policy P_S. In fact, we can see that, when the most convenient option is to perform a single inspection (e.g., Setup A), P_S still performs some (albeit few) unnecessary extra inspections. This is due to the stringency determined by the circumstance that under P_S, the postponement time is equal to a quantile of the RUL whose order is constrained to be the same at all decision epochs. Defining a more flexible sequential decision rule, which could allow to reduce the number of unnecessary inspections, will be object of future research.

6. Conclusions

This paper investigated the value of information provided by inspections in condition-based maintenance of degrading units in the simultaneous presence of temporal variability and unit-to-unit variability. The degradation process was modeled by a gamma process with random effect.

To this aim, a novel condition-based maintenance policy was proposed that allows for multiple inspections, with their timing being adaptively defined based on the conditional distribution of the useful life given the latest degradation measurement.

The adopted performance measure, used to define the optimal policy, is the long-run average cost rate, which accounts for inspection, logistic, preventive replacement, corrective replacement, and downtime costs. The performances of the proposed policy are compared with those of two similar policies defined by constraining the maximum number of inspections to one or two, respectively. A numerical analysis highlighted that the economic benefit of the inspections strongly depends on logistic and inspection costs (i.e., costs of acquiring information), as well as on the amount of temporal variability and unit-to-unit variability.

In particular, obtained results show that when unit-to-unit variability is high and temporal variability is low, performing numerous inspections appears to not be advantageous. When this scenario is paired with high inspection and logistic costs, then even a one-inspection policy becomes convenient.

Conversely, when temporal variability is high and unit-to-unit variability is low, especially when inspection and logistic costs are low, then carrying out multiple inspections can provide significant economic benefits.

Future research should explore refined structures to minimize unnecessary inspections and relax the assumption of perfect inspection, for example by incorporating measurement error, to better reflect real-world conditions.

References

Lawless, J., and M. Crowder (2004). Covariates and random effects in a gamma process model with application to degradation and failure. *Lifetime Data Analysis 10(3)*, 213-227.

Bismut, E., M. D. Pandey, and D. Straub (2022). Reliability based inspection and maintenance planning of a nuclear feeder piping system. *Reliability Engineering & System Safety, 224, 108521*, 1-16.

Bautista, L., I. T. Castro, and L. Landesa (2022). Condition-based maintenance for a system subject to multiple degradation processes with stochastic arrival intensity. *European Journal of Operational Research 302(2),* 560-574.

Ross S. M. (1983). *Stochastic Processes*. John Wiley, New York.

Esposito, N., A. Mele, B. Castanier, and M. Giorgio (2023). A new gamma degradation process with random effect and state-dependent measurement error. Proceedings of the institution of Mechanical Engineers, Part O: *Journal of Risk and Reliability 237(5)*, 868-885.

Piscopo, A., N. Esposito, B. Castanier, and M. Giorgio (2024). Maintenance of degrading units: new policy and some critical considerations about the value of inspections. In proceedings of the 34th European Safety and Reliability conference (ESREL2024), 145-154. Polish Safety and Reliability Association.