Journal of Interdisciplinary Engineering, Computing and Automation
Vol. 01, Issue 01, July–December 2026, pp. 25-36

No Diagnostic Substitutes for Structure That Is Not There: When Age-Based and Condition-Based Maintenance Actually Pay

Vynotdni Rathinasamy

School of Civil Engineering, Universiti Sains Malaysia, Jalan Bukit Panchor, 14300 Nibong Tebal, Pulau Pinang, Malaysia.

Email: vynotdni@usm.my

ABSTRACT

Condition monitoring with machine-learning diagnostics is deployed on the premise that predicting failure reduces cost. Whether it can depends on whether the failure process contains predictable structure. We identify the preconditions for age-based and condition-based maintenance and quantify what each is worth when they are met and when they are not. Failure times were Weibull with scale 1000 hours and shape varied from 1 to 5. Long-run cost per hour was computed analytically for age-based replacement and by simulating 40,000 renewal cycles for condition-based replacement under a delay-time model, in which a defect arrives and the asset fails after a further interval during which a diagnostic may detect it. Cost ratio, detection probability, defect window and false alarm rate were swept. Age-based replacement saved 0 per cent over run-to-failure at shape 1, where no finite optimal interval exists, rising to 40.7 per cent at shape 2 and 72.0 per cent at shape 5. Condition-based replacement was almost independent of shape, saving 71.0 to 71.2 per cent across the range, because it exploits the defect window rather than wear-out. That window behaved as a gate: with no detectable defective stage the policy returned -4.1 per cent, worse than run-to-failure once monitoring cost is included, while a window of 2 per cent of mean life returned 74.0 per cent and further widening added nothing. Detection probability moved the saving from -4.9 to 80.5 per cent, while raising the false alarm rate to one per 1,000 hours cost only 5.3 points. At a 2:1 cost ratio no policy saved more than 22 per cent. The two policies exploit different properties of the failure process, and each has a precondition that engineering effort cannot supply. Establishing which is present is cheap and should precede any diagnostic investment.

Keywords: predictive maintenance; condition monitoring; reliability engineering; delay-time model; cost-rate optimisation; industrial automation.

1. INTRODUCTION

Industrial condition monitoring programmes are typically justified by the accuracy of their diagnostics. A model is shown to classify bearing faults, detect gearbox wear or forecast remaining useful life, its accuracy is reported, and a business case follows from the assumption that better prediction means lower cost [1-3,17,18].

The assumption is not always true, and the reason is a property of the asset rather than of the model. Maintenance policies work by exploiting structure in the failure process. Age-based replacement exploits wear-out: if the hazard rate rises with age, replacing before old age avoids

failures. Condition monitoring exploits a different structure: if a detectable defect precedes failure by some interval, the defect can be found and the replacement planned [15,16].

Where neither structure is present, neither policy helps. An asset whose failures arrive as a Poisson process is as good as new at every moment until it fails; there is no age at which replacement is advantageous, and there is no degradation signal to observe. This is textbook reliability theory [4-6,22], and it is routinely overlooked when monitoring programmes are specified, because the diagnostic accuracy question is more tractable and more interesting than the question of whether there is anything to diagnose [25,26].

What is less often quantified is the magnitude of each effect and, in particular, how sharply the value of condition monitoring depends on the size of the detectable window. We compute cost rates for both policies over a range of failure processes and diagnostic qualities, in order to identify the conditions under which an investment in monitoring can be expected to return anything [23,24].

1.1 Contributions

We quantify the value of age-based replacement as a function of the Weibull shape parameter, including the degenerate case where no optimal interval exists; we quantify the value of condition-based replacement under a delay-time model and show that it is governed by the defect window rather than by wear-out; we establish that the window behaves as a threshold rather than a continuous parameter; and we compare the sensitivity of the result to detection probability, false alarm rate and the cost ratio.

2. METHOD

2.1 Failure Model and Costs

Time to failure is Weibull with scale 1000 hours and shape β varied from 1 to 5. A shape of 1 gives an exponential distribution with constant hazard — a memoryless process with no wear-out — and larger values give increasing hazard [22]. The scale is held fixed, so mean life varies mildly across the range and all comparisons are made against the run-to-failure cost rate for the same shape.

A planned replacement costs 1 unit and an unplanned failure 8 units, a ratio of 8:1 intended to represent the downtime, secondary damage and expediting costs that accompany an unplanned stoppage. Monitoring costs 0.0004 units per operating hour. The cost ratio is swept separately in Section 3.4.

2.2 Age-Based Replacement

Under age-based replacement a unit is replaced on failure or on reaching age T, whichever comes first. The long-run cost per unit time follows the classical renewal-reward result [4,5]:

$$C(T)=\frac{\left[c_p R(T)+c_f\left(1-R(T)\right)\right]}{\int_0^T R(t)\,dt} \tag{1}$$

where R is the reliability function, c_p the planned replacement cost and c_f the failure cost. We minimise this numerically over T. For β = 1 the expression is monotone and the minimum lies at infinity, meaning the policy degenerates to run-to-failure; we report this case as having no optimal interval rather than reporting a large arbitrary one.

2.3 Condition-Based Replacement

Condition monitoring is modelled with a delay-time formulation, which is the standard way to represent an inspectable degradation process [7,8]. The asset is sound until a defect arrives, then fails after a further interval. A diagnostic can detect the defect only while it exists, that is, during that interval.

We parameterise the mean interval as a fraction of mean life, so that a fraction of zero represents an asset that fails without any detectable precursor and larger fractions represent progressively more observable degradation. On detection the asset is replaced at a planned cost at a uniformly distributed point within the remaining window.

We emphasise this choice because it is what makes the comparison meaningful. A diagnostic modelled as issuing a warning at a fixed fraction of the true time to failure would be an oracle, not a detector, and would appear to add value even for a process containing nothing detectable. The delay-time model ties the diagnostic to an underlying physical state [9].

False alarms occur at a constant hourly rate and trigger a planned replacement of a sound asset. Parameters are given in Table 1 and the complete script accompanies this paper. Figure 1 summarises the structures the two policies exploit.

Table 1: Model parameters.

Parameter	Value	Note
Failure distribution	Weibull, scale 1000 h	Shape varied from 1 to 5
Shape $\beta = 1$	Constant hazard	Memoryless; no wear-out to exploit
Planned replacement cost	1	Normalised
Unplanned failure cost	8	Includes downtime and secondary damage
Monitoring cost	0.0004 per hour	Charged whenever monitoring is in use
Defect window	Mean delay time as a fraction of mean life	Zero means no detectable precursor
Detection probability	0 to 0.99	Probability the defect is found while it exists
Renewal cycles simulated	40,000	Per configuration

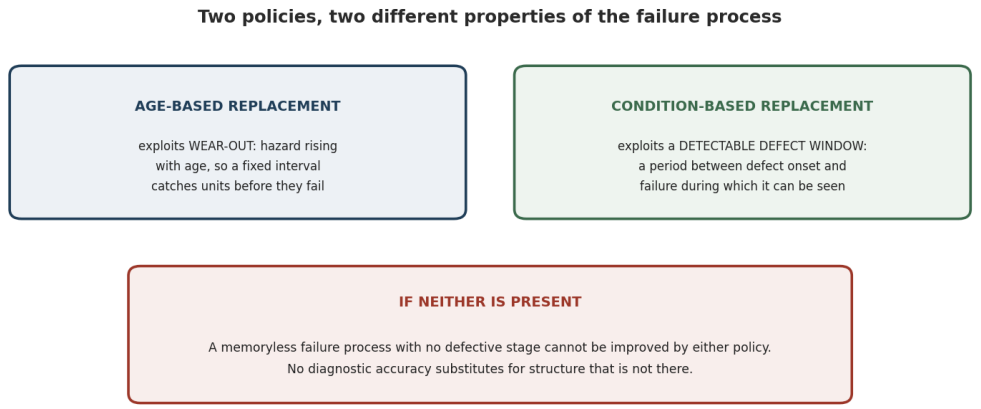


Fig.1: The distinct structures the two policies exploit, and the case in which neither is present.

3. RESULTS

3.1 Age-Based Replacement Needs Wear-Out

Figure 2 plots the cost-rate curve against replacement interval for three shapes. At β = 2 and β = 3 the curve has a clear interior minimum, at 382 and 416 hours respectively. At β = 1 it has none: the curve descends monotonically toward the run-to-failure rate and never falls below it.

Table 2 and Figure 3 give the resulting savings. Age-based replacement returned 0 per cent at β = 1, 20.4 per cent at β = 1.5, 40.7 per cent at β = 2, 59.4 per cent at β = 3 and 72.0 per cent at β = 5.

The β = 1 result is exact rather than numerical. For a memoryless process a surviving unit is statistically identical to a new one, so replacing it discards remaining life of exactly the same expected value as the replacement provides, at the cost of a planned intervention. No interval can be advantageous, and this is a property of the failure process that no scheduling sophistication alters.

The practical consequence is that a preventive replacement programme on an asset class whose failures are close to random is a pure cost. Fitting a Weibull distribution to existing failure records is a routine exercise and settles the question before any programme is designed.

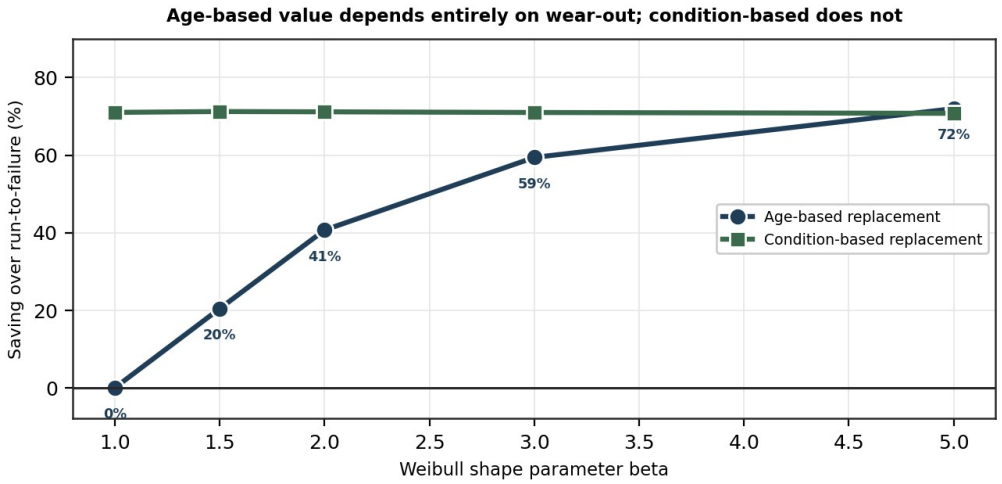


Fig.3: Saving over run-to-failure for both policies against the shape parameter.

Table 2: Policy value against the Weibull shape parameter.

β	Mean life (h)	Run-to-failure cost rate	Optimal T (h)	Age-based saving	Condition-based saving
1.0	1000	0.00800	none	0.0%	71.0%
1.5	903	0.00886	451	20.4%	71.2%
2.0	886	0.00903	382	40.7%	71.2%
3.0	893	0.00896	416	59.4%	71.0%
5.0	918	0.00871	514	72.0%	70.8%

Condition-based figures use detection probability 0.90, defect window 20% of mean life and a false alarm rate of 1e-5 per hour.

3.2 Condition Monitoring Needs a Detectable Window

Table 2 also shows the contrast that motivates the paper. Condition-based replacement saved between 70.8 and 71.2 per cent across the entire shape range, including at β = 1 where age-based replacement is worthless. The two policies are not competing versions of the same idea; they exploit different structure.

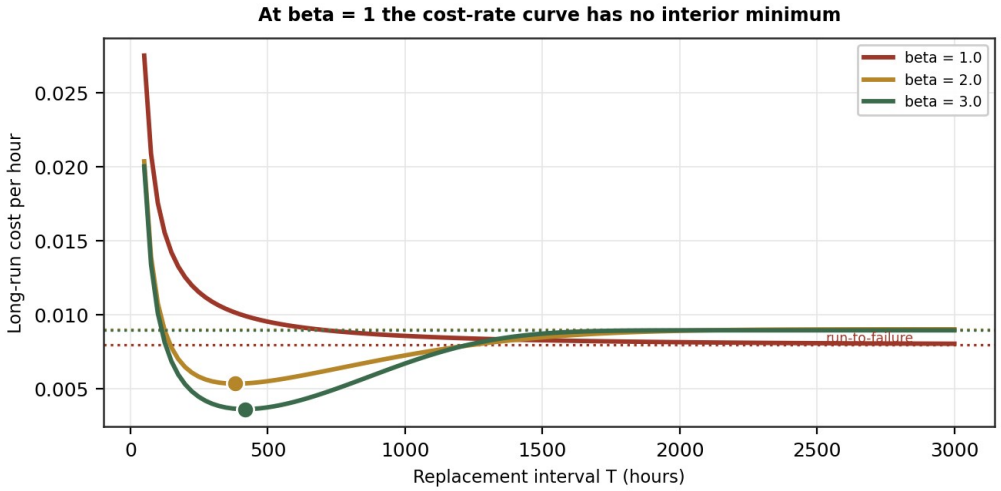


Fig.2: Cost rate against replacement interval. At β = 1 there is no interior minimum.

Figure 4 and Table 3 vary the defect window. With no detectable defective stage the policy returned -4.1 per cent — a loss, because the monitoring cost is incurred and nothing is ever detected. A window of only 2 per cent of mean life returned 74.0 per cent, and widening it to 40 per cent changed the figure only to 70.3 per cent.

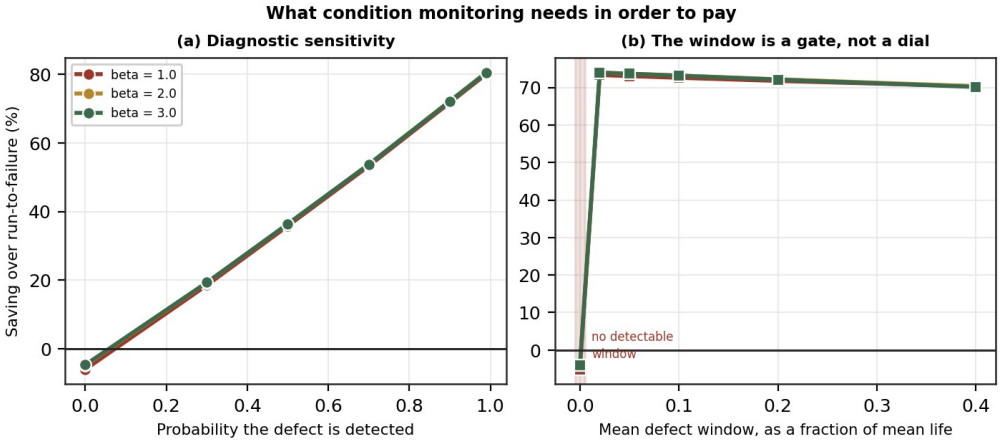


Fig.4: Condition-based saving against detection probability and against the defect window.

The window therefore behaves as a gate rather than a dial. What matters is whether a detectable precursor exists at all; given that it does, its duration is a second-order consideration, and a longer window even slightly reduces the saving because detection occurs earlier and more remaining life is discarded.

This reframes the usual specification question. A monitoring programme is often specified by asking how much warning time the diagnostic must provide. On these results the prior question — whether the failure mode has a physically detectable precursor at all — is the

one that determines whether the programme can succeed, and it is answered by failure mode analysis rather than by data science [12-14].

Table 3: Effect of the detectable defect window.

Defect window (fraction of mean life)	Saving, β = 1	Saving, β = 2	Saving, β = 3
0.00	-5.1%	-4.1%	-4.1%
0.02	73.3%	74.0%	74.0%
0.05	73.0%	73.7%	73.7%
0.10	72.6%	73.2%	73.2%
0.20	71.8%	72.2%	72.1%
0.40	70.3%	70.3%	70.1%

3.3 Sensitivity Matters More Than False Alarms

Figure 4 and Table 4 vary detection probability. The relationship is close to linear, from -4.9 per cent at zero detection — the cost of monitoring with no benefit — to 80.5 per cent at 0.99.

Table 5 and Figure 5 vary the false alarm rate. Raising it from zero to one per 1,000 operating hours reduced the saving at β = 2 only from 72.2 to 66.9 per cent.

Table 5: Effect of false alarms.

False alarms per hour	Mean time between false alarms	Saving, β = 2	Saving, β = 3
none	—	72.2%	72.1%
1e-5	100,000 h	72.1%	72.1%
5e-5	20,000 h	72.0%	72.1%
1e-4	10,000 h	71.9%	71.9%
3e-4	3,333 h	70.9%	71.4%
1e-3	1,000 h	66.9%	68.3%

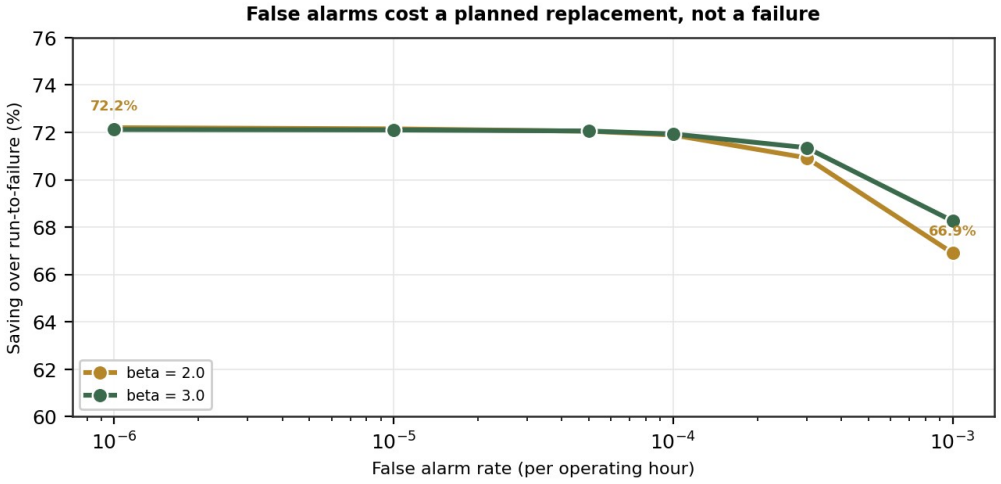


Fig.5: Saving against false alarm rate.

The asymmetry has a simple cause and a useful implication. A missed detection costs a failure; a false alarm costs a planned replacement, which in this setting is an eighth as expensive and also resets the asset. Diagnostics for maintenance should therefore be tuned toward sensitivity far more aggressively than a symmetric classification metric would suggest, and reporting a balanced accuracy or F-score for such a model obscures the asymmetry that governs its economic value [19].

We note the limit of this argument. Where a false alarm triggers a costly intervention, an unnecessary shutdown of a process line, or erodes operator trust to the point that genuine alarms are ignored, the asymmetry narrows or reverses [27]. Our model captures only the direct replacement cost.

3.4 The Cost Ratio Gates Everything

Table 6 and Figure 6 vary the ratio of failure cost to planned replacement cost. At 2:1, age-based replacement saved 3.3 per cent at $\beta = 2$ and condition-based 21.9 per cent. At 32:1 the same policies saved 69.1 and 84.9 per cent.

Table 4: Effect of detection probability.

Detection probability	Saving, $\beta = 1$	Saving, $\beta = 2$	Saving, $\beta = 3$
0.00	-6.0%	-4.9%	-4.7%
0.30	18.5%	19.3%	19.3%
0.50	35.8%	36.4%	36.4%
0.70	53.3%	53.8%	53.8%
0.90	71.7%	72.2%	72.2%
0.99	80.1%	80.5%	80.5%

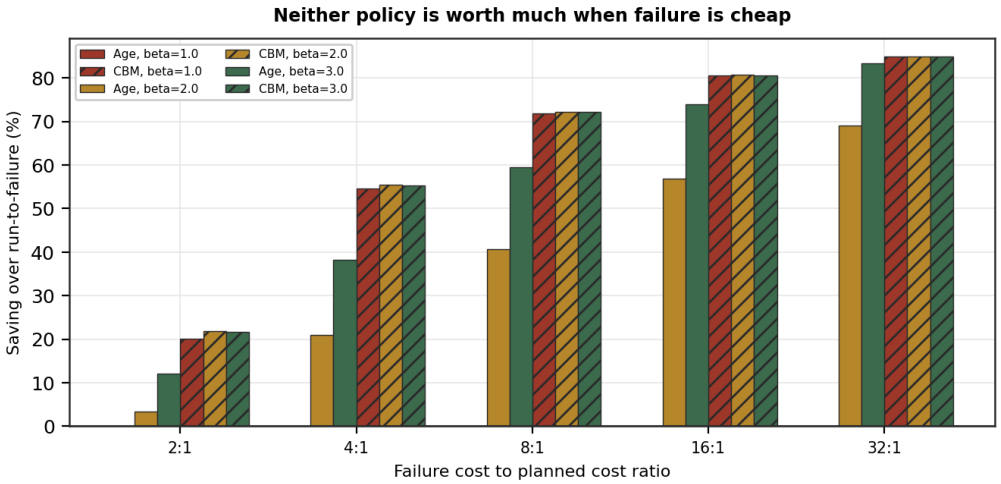


Fig.6: Saving against the failure-to-planned cost ratio for both policies.

Where an unplanned failure costs little more than a planned replacement, there is little for any policy to recover, and run-to-failure is the correct engineering answer rather than a failure of maintenance planning [10,11]. The cost ratio is usually the easiest of all these parameters to estimate and the one most often left implicit in a business case.

Note also that the ratio interacts with the failure process. At $\beta = 1$ the age-based column is zero at every ratio, because the policy cannot help regardless of how expensive failure is. A

high cost ratio makes a viable policy more valuable; it does not make an inapplicable policy applicable.

Table 6: Effect of the cost ratio.

Cost ratio	Age, β=1	CBM, β=1	Age, β=2	CBM, β=2	Age, β=3	CBM, β=3
2:1	0.0%	20.2%	3.3%	21.9%	12.0%	21.7%
4:1	0.0%	54.7%	21.1%	55.5%	38.3%	55.4%
8:1	0.0%	71.9%	40.7%	72.3%	59.4%	72.2%
16:1	0.0%	80.6%	56.9%	80.7%	73.9%	80.6%
32:1	0.0%	84.9%	69.1%	84.9%	83.4%	84.8%

4. DISCUSSION

4.1 Principal Findings

Four findings follow. Age-based replacement returned 0 per cent at β = 1, where no optimal interval exists, rising monotonically to 72.0 per cent at β = 5. Condition-based replacement was nearly independent of β, confirming that the two policies exploit different structure. The detectable defect window acted as a gate: -4.1 per cent with no window against 74.0 per cent with a window of 2 per cent of mean life, and little further gain thereafter. And detection probability dominated false alarm rate by roughly an order of magnitude in its effect on cost.

The unifying claim is that both policies have preconditions which are properties of the asset, not of the maintenance organisation. Diagnostic accuracy is worth improving only once the precondition is established, and a programme that omits that check can invest heavily in a model that could not have paid for itself at any accuracy.

4.2 Relation to Existing Work

The age-replacement cost-rate result and the non-existence of an optimal interval for a constant hazard are classical [4-6]. Delay-time modelling of inspection and condition monitoring is an established body of work with substantial methodological development [7-9], and the economics of maintenance policy selection has been treated extensively [10,11,23,24]. Reliability-centred maintenance frameworks make failure-mode analysis the entry point to policy selection for precisely the reasons set out here [12-14].

None of the individual results is new to reliability engineering. What we offer is a quantification in a form aimed at the machine-learning side of a condition monitoring project, where the diagnostic is usually the object of attention and the asset's failure characteristics are treated as background [17,18,20,21]. The finding that the defect window behaves as a threshold, and that this threshold rather than model accuracy is the first thing to establish, is the part we would most want a project team to take away.

4.3 Recommendations

Figure 7 and Table 7 set out the guidance. Two points deserve comment.

Fit the failure distribution before commissioning a model. Estimating a Weibull shape parameter from maintenance records is a few hours of work and determines whether age-based replacement can help at all. Where the estimate is near unity, a preventive replacement programme should not be built, and the case for condition monitoring rests entirely on whether a physical precursor exists.

Table 7: Preconditions and the decision they imply.

Asset property	Age-based	Condition-based	Recommended
Wear-out present, detectable precursor present	Effective	Effective	Either; compare cost of monitoring
Wear-out present, no precursor	Effective	Loss-making	Age-based replacement
No wear-out, detectable precursor	Worthless	Effective	Condition monitoring
No wear-out, no precursor	Worthless	Loss-making	Run to failure; invest in redundancy or design

Establish the precursor physically, not statistically. Whether a failure mode has a detectable degradation stage is a question for failure-mode analysis, materials evidence and instrumented trials, not for a model fitted to whatever sensor data happen to exist. A model can always find some correlation in historical data; whether it corresponds to a precursor with usable lead time is a separate and prior question.

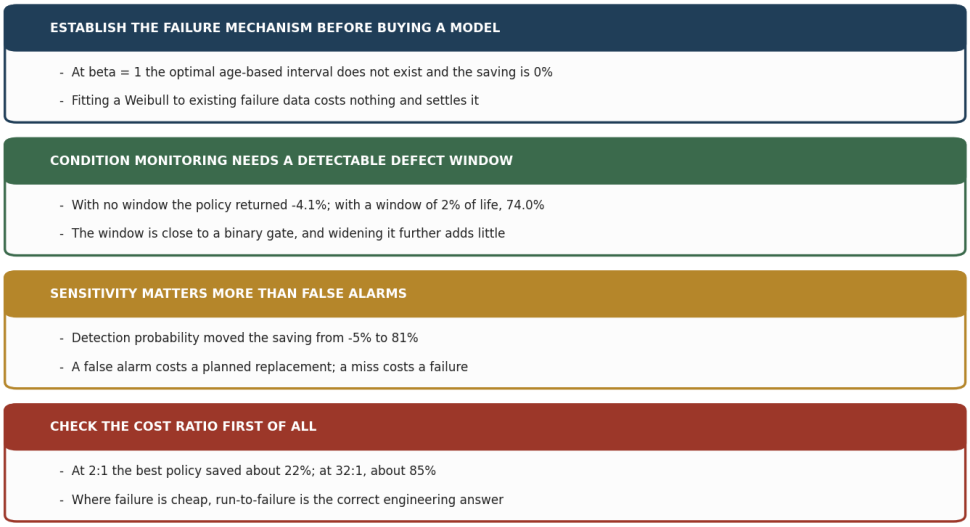


Fig.7: Proposed sequence for specifying a maintenance programme.

5. LIMITATIONS

The delay-time model assumes an exponentially distributed window independent of the failure time, and the diagnostic is characterised by a single detection probability applied uniformly within that window. Real detectability rises as a defect grows, so a constant probability understates the advantage of frequent inspection and overstates the value of a single early observation.

Costs are constant and known, replacement is instantaneous and restores the asset to new, and there is no opportunistic grouping of maintenance across a fleet, no spare-parts constraint and no production scheduling interaction. Each of these is material in practice and each would change the numbers, though none affects the existence of the preconditions we identify.

We compare policies at a single set of nominal parameters within each sweep rather than optimising each policy jointly over all of its parameters. A fully optimised condition-based

policy with an inspection interval chosen against cost would perform somewhat better than reported here.

The monitoring cost is a single per-hour figure and excludes installation, integration, model development and the ongoing effort of investigating alarms, which in practice often dominate. Our condition-based savings should therefore be read as an upper bound on what the policy could return, before those costs are subtracted.

Finally, the Weibull shape parameter is treated as known. In practice it is estimated from censored maintenance records of modest size, and an estimate near unity carries wide uncertainty. A decision to abandon preventive replacement on the strength of a point estimate would be premature; the appropriate response to β near 1 is to widen the analysis, not to conclude immediately.

6. CONCLUSION

Age-based replacement saved 0 per cent over run-to-failure when the failure process was memoryless and 72.0 per cent when strongly wearing out. Condition-based replacement saved around 71 per cent irrespective of wear-out, because it exploits a detectable defect window rather than age.

That window is a gate. With no detectable precursor the policy returned -4.1 per cent, a loss; with a precursor occupying 2 per cent of mean life it returned 74.0 per cent, and widening it further added nothing. Detection probability then set the value within that regime, while false alarms at realistic rates cost a few percentage points.

Both policies therefore rest on a property of the asset that no amount of modelling supplies. Determining which property is present costs a Weibull fit and a failure-mode review, and it decides whether a condition monitoring investment can return anything at all. That determination belongs at the start of such a project rather than in its post-mortem.

DECLARATIONS

Ethics approval and consent to participate. Not applicable. This is a simulation and analytical study using synthetic data generated by the author. It did not involve human participants, animals or personal data, and institutional review board approval and informed consent were therefore not required.

Consent for publication. Not applicable.

Data availability statement. No external data were used. All results were generated by the simulation and analysis script described in Section 2, developed by the author in Python using NumPy, SciPy and Matplotlib [28-30], with all random seeds fixed. Every model parameter is stated in full in Table 1, and every figure and every value in Tables 2 to 6 can be regenerated exactly by executing that script. The script is available from the corresponding author upon reasonable request and will be deposited in a public repository with a permanent identifier on acceptance.

Conflict of interest. The author declares no competing interests, financial or non-financial, in relation to the subject matter of this manuscript. The author has no financial or advisory relationship with any vendor of condition monitoring hardware, maintenance management software or predictive analytics services.

Funding. This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

Author's contribution. Vynotdni Rathinasamy is the sole author of this work and was responsible for the conception and design of the study, the derivation of the cost-rate analysis, the implementation and verification of the simulation, the preparation of all figures and tables,

the formulation of the recommendations, and the drafting, revision and final approval of the manuscript. The author is accountable for all aspects of the work.

Acknowledgements. None declared.

REFERENCES

1. Lei Y, Yang B, Jiang X, Jia F, Li N, Nandi AK. Applications of machine learning to machine fault diagnosis: a review and roadmap. Mech Syst Signal Process. 2020;138:106587.
2. Zhang W, Yang D, Wang H. Data-driven methods for predictive maintenance of industrial equipment: a survey. IEEE Syst J. 2019;13(3):2213-27.
3. Zonta T, da Costa CA, da Rosa Righi R, de Lima MJ, da Trindade ES, Li GP. Predictive maintenance in the Industry 4.0: a systematic literature review. Comput Ind Eng. 2020;150:106889.
4. Barlow RE, Proschan F. Mathematical theory of reliability. Philadelphia: SIAM; 1996.
5. Ross SM. Introduction to probability models. 12th ed. London: Academic Press; 2019.
6. Rausand M, Hoyland A. System reliability theory: models, statistical methods, and applications. 2nd ed. Hoboken: Wiley; 2004.
7. Christer AH. Developments in delay time analysis for modelling plant maintenance. J Oper Res Soc. 1999;50(11):1120-37.
8. Wang W. An overview of the recent advances in delay-time-based maintenance modelling. Reliab Eng Syst Saf. 2012;106:165-78.
9. Jardine AKS, Lin D, Banjevic D. A review on machinery diagnostics and prognostics implementing condition-based maintenance. Mech Syst Signal Process. 2006;20(7):1483-510.
10. Wang H. A survey of maintenance policies of deteriorating systems. Eur J Oper Res. 2002;139(3):469-89.
11. de Jonge B, Scarf PA. A review on maintenance optimization. Eur J Oper Res. 2020;285(3):805-24.
12. Moubray J. Reliability-centred maintenance. 2nd ed. Oxford: Butterworth-Heinemann; 1997.
13. SAE International. Evaluation criteria for reliability-centered maintenance (RCM) processes. SAE JA1011. Warrendale: SAE; 2009.
14. Nowlan FS, Heap HF. Reliability-centered maintenance. Washington: US Department of Defense; 1978.
15. Ahmad R, Kamaruddin S. An overview of time-based and condition-based maintenance in industrial application. Comput Ind Eng. 2012;63(1):135-49.
16. Alaswad S, Xiang Y. A review on condition-based maintenance optimization models for stochastically deteriorating system. Reliab Eng Syst Saf. 2017;157:54-63.
17. Lee J, Wu F, Zhao W, Ghaffari M, Liao L, Siegel D. Prognostics and health management design for rotary machinery systems: reviews, methodology and applications. Mech Syst Signal Process. 2014;42(1-2):314-34.
18. Si XS, Wang W, Hu CH, Zhou DH. Remaining useful life estimation: a review on the statistical data driven approaches. Eur J Oper Res. 2011;213(1):1-14.
19. Saxena A, Celaya J, Saha B, Saha S, Goebel K. Metrics for offline evaluation of prognostic performance. Int J Progn Health Manag. 2010;1(1):4-23.
20. Vachtsevanos G, Lewis F, Roemer M, Hess A, Wu B. Intelligent fault diagnosis and prognosis for engineering systems. Hoboken: Wiley; 2006.
21. Kim NH, An D, Choi JH. Prognostics and health management of engineering systems. Cham: Springer; 2017.

22. Elsayed EA. Reliability engineering. 3rd ed. Hoboken: Wiley; 2021.
23. Dekker R. Applications of maintenance optimization models: a review and analysis. Reliab Eng Syst Saf. 1996;51(3):229-40.
24. Scarf PA. On the application of mathematical models in maintenance. Eur J Oper Res. 1997;99(3):493-506.
25. Bousdekis A, Lepenioti K, Apostolou D, Mentzas G. Decision making in predictive maintenance: literature review and research agenda for Industry 4.0. IFAC-PapersOnLine. 2019;52(13):607-12.
26. Tiddens W, Braaksma J, Tinga T. Exploring predictive maintenance applications in industry. J Qual Maint Eng. 2022;28(1):68-85.
27. Bokrantz J, Skoogh A, Berlin C, Wuest T, Stahre J. Smart maintenance: a research agenda for industrial maintenance management. Int J Prod Econ. 2020;224:107547.
28. Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, et al. Array programming with NumPy. Nature. 2020;585(7825):357-62.
29. Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. Nat Methods. 2020;17(3):261-72.
30. Hunter JD. Matplotlib: a 2D graphics environment. Comput Sci Eng. 2007;9(3):90-95.