Health Index and Markov Prediction Modelling of Maintenance Repair Rates: A Square-Root Law of Intervention Intensity in Transformer Asset Management

Muhammad Sharil Yahaya[1]*, Zul Hasrizal Bohari[1], Raja Saharuddin Raja Daud[2]

[1] Faculty of Electrical Technology and Engineering, Universiti Teknikal Malaysia Melaka, Hang Tuah Jaya, 76100 Durian Tunggal, Melaka, Malaysia

[2] TNB Research Sdn. Bhd., Kawasan Institut Penyelidikan, 43000 Kajang, Selangor, Malaysia

***Corresponding Author**

DOI: https://doi.org/10.47772/IJRISS.2026.101000534

Received: 25 October 2026; Accepted: 30 October 2026; Published: 07 November 2026

ABSTRACT

The Health Index (HI) quantifies the present condition of a transformer population, and the Markov Prediction Model (MPM) projects how that population will migrate between condition states once a pre-determined maintenance repair rate is applied to the transition matrix. Together they forecast deterioration and confirm that intervention improves the forecast. What they seldom characterize is the economic structure of that improvement: how benefit scales with repair rate, where on the deterioration path intervention is most productive, and what rate is efficient. These are allocation questions, and the maintenance-optimization models that address them typically require cost parameters that utilities cannot specify credibly — long identified as the principal barrier to their adoption in practice. This paper shows that the economic structure can instead be recovered from the physical response of a validated deterioration model, without cost data. We re-implement the published HI-MPM of a Malaysian distribution transformer population (3,195 oil samples, 373 units, 33/11 kV, 30 MVA), verify it against its published transition matrices and reported outcomes, and use it to generate a fine-grained repair-rate response surface spanning 2% to 50%. Three findings follow. First, the average improvement in the HI performance condition curve scales as a power law in repair rate with exponent 0.44 (R^2 = 0.994 on published values; 0.51 on the simulated surface), establishing an approximate square-root law: the second decile of intensity returns 36% of the effectiveness of the first and the third only 18%. Monte Carlo analysis over 4,000 randomized fleet compositions and zone weightings confirms diminishing returns in 100% of cases, with tight dispersion (interquartile range 0.34–0.37 and 0.17–0.19 respectively). Second, the placement of intervention dominates its intensity: directing effort at the poor state rather than the good state yields a median benefit ratio of 8.8 (5th percentile 3.7), and the single published comparison of 2.60 is conservative. Third, the power-law form admits a closed-form efficient intensity, $\rho^* = (\alpha A/c)^{(1/(1-\alpha))}$, expressed entirely through the cost-to-benefit ratio and therefore computable without monetary data. Expected dwell times recovered directly from the MPM transition probabilities additionally reveal a 7.2-fold divergence between network zones in the persistence of HI state 3 (fair). The study contributes a transferable method for extracting decision-relevant economic structure from HI-based Markov models already in service, together with an allocation rule for constrained maintenance budgets: spread a modest repair rate widely, target the degraded cohort, and differentiate by zone.

Keywords: health index; Markov prediction model; maintenance repair rate; transformer asset management; diminishing marginal returns; resource allocation

INTRODUCTION

The maintenance budget encodes two decisions. How much intervention to buy, and where to point it. Condition-based maintenance (CBM) has transformed the information available for both, and two instruments carry most of that load in utility practice. The Health Index (HI) compresses many condition parameters into a single normalized score, so a heterogeneous population becomes comparable and therefore rank able. The Markov Prediction Model (MPM) takes those states and projects how the population will migrate between them, and — by inserting a pre-determined maintenance repair rate into the transition matrix — how that migration changes when the utility intervenes. Diagnostic coverage has expanded alongside both (Jardine et al., 2006; Alsumaidaee et al., 2022). What has not kept pace is the layer that turns these forecasts into spending decisions.

The gap is a specific one. A typical CBM study reports that intervention improves the forecast condition of a population and quantifies that improvement at one or two assumed intervention levels. This establishes that maintenance works. It says nothing about what the next unit of maintenance buys relative to the last, or whether that unit does good on assets already degraded than on assets still sound, or at what point further intervention stops being worth funding. Those are questions about marginal products and allocative efficiency. They are also the questions a budget holder must answer.

Maintenance-optimization theory has addressed these questions since Barlow and Hunter (1960), and the literature is extensive (Wang, 2002; Van Horenbeek et al., 2010). Uptake in practice has nonetheless been thin. Dekker (1996) put this down mainly to the data such models demand, and above all to cost parameters — the cost of an unplanned failure, or of lost supply — that few organizations can state with confidence. Three decades on, the diagnosis still fits. Utilities tend to hold rich condition data and sparse, contested cost data.

We propose a route around that constraint. The economic structure of maintenance — how benefit scales with the repair rate, and how productive the same rate is when applied at different HI states — is a property of the physical system's response rather than of its costs. Where a validated HI-based Markov model already exists, that structure can be recovered from its response surface alone. Costs are then needed only to locate a point on a curve whose shape is already known, and even that can be handled parametrically, through a cost-to-benefit ratio rather than absolute money.

We demonstrate the approach on the HI-MPM developed for a Malaysian distribution transformer population by Yahaya et al. (2018), in which 3,195 oil samples drawn from 373 transformers rated 33/11 kV and 30 MVA were used to estimate state-transition probabilities, and pre-determined maintenance repair rates of 10%, 20% and 30% were applied to examine their effect on health-index state distribution and the performance condition curve. That study showed that intervention improves the forecast, and that its effect is largest when applied at the poor state. What it did not do was characterize how the effect scales, how robust that scaling is to fleet composition, how large the efficiency gap between intervention points really is, or what intensity would be efficient.

Four research questions follow. (RQ1) What functional form relates the maintenance repair rate to the improvement in the HI performance condition curve, and does it exhibit diminishing returns? (RQ2) How robust is that form to the baseline HI state distribution, which in practice is rarely known precisely? (RQ3) How much more productive is the same repair rate applied at HI state 4 (poor) than at state 2 (good)? (RQ4) Can an efficient repair rate be derived from the MPM response alone, without monetary cost data?

Relative to the baseline secondary dataset of Yahaya et al. (2018), which reports the HI-MPM and its outcomes at three discrete repair rates, the novelty of this study is theoretical rather than empirical: no new data are introduced, and every contribution is a property derived from that model. We make four contributions. First, we establish an approximate square-root law in the maintenance repair rate for this asset class, estimating the exponent twice over — once from the published outcomes, once from a re-implemented response surface — which recasts the three published outcomes as points on a concave benefit function with diminishing marginal returns. Second, we show by Monte Carlo analysis that the diminishing-returns

structure does not depend on the baseline HI state distribution, which is what makes it usable as a planning heuristic rather than a one-off result. Third, we quantify the efficiency gap between intervention points, find the previously reported comparison to be conservative, and show how the gap narrows, but persists, under imperfect repair. Fourth, we formulate the choice of repair rate as a net-benefit maximization problem and derive a closed-form efficient repair rate, confirmed as a unique maximum by its first- and second-order conditions, that needs only a cost-to-benefit ratio, which speaks directly to the data barrier Dekker (1996) identified. We also reinterpret the MPM diagonal transition probabilities as expected dwell times per HI state, and these turn out to diverge sharply between network zones. None of this is specific to transformers: any HI-based Markov model that can be run at several repair rates admits the same treatment.

The remainder of the paper is organized as follows. The next section reviews the relevant literature. The research methodology describes the source model, its re-implementation and verification, and the derived measures. Results and discussion follow, after which the managerial and economic implications are developed, limitations stated and conclusions drawn.

LITERATURE REVIEW

Maintenance as an Allocation Problem

Asset management is defined as the coordinated activity by which an organization realizes value from its assets, value being the balance of cost, risk and performance across the life cycle (ISO, 2014). The definition is economic rather than technical, framing maintenance as allocation under constraint. Brown and Humphrey (2005) apply this to electricity networks; Schneider et al. (2006) observe that the binding difficulty for utilities is not the absence of condition data but the absence of a consistent basis for converting it into prioritized expenditure.

The formal treatment of intervention intensity originates with the preventive-replacement policies of Barlow and Hunter (1960) and is surveyed for deteriorating systems by Wang (2002). Under any such policy, expected cost comprises intervention performed on assets that did not require it plus the expected cost of failures occurring between interventions, and the optimum lies where marginal benefit equals marginal cost. Van Horenbeek et al. (2010) emphasize that realized benefit depends on the criteria translating condition signals into action; Jardine et al. (2006) note that this decision layer is substantially less developed than the diagnostics feeding it. Alsumaidaee et al. (2022) reach the same conclusion for medium-voltage switchgear, where detection capability has outrun the frameworks for acting on it.

Two observations from this literature motivate what follows. Diminishing marginal returns are the shape theory predicts for the intensity–benefit relationship, yet utility asset studies almost always report a single intervention scenario, from which no marginal structure can be inferred at all. The second observation is the more useful one: because that structure belongs to the physical response rather than to the cost side, it can be estimated even where cost parameters are missing. That is precisely the condition Dekker (1996) identified as blocking adoption.

The Health Index and the Markov Prediction Model

The health index compresses multiple condition parameters into a single normalized score, rendering a heterogeneous population comparable and therefore rank able within a budget cycle. The approach was established for transformers by Naderian et al. (2008) and Jahromi et al. (2009), and developed by Haema and Phadungthin (2012), Scatiggio and Pompili (2013) and Azmi et al. (2017). Panmala et al. (2022) extend it to gas-insulated switchgear by combining the index with conditional risk factors. In all of this work the index describes present condition.

Forecasting future conditions requires a second instrument. The MPM treats deterioration as a discrete-state stochastic process in which an asset either holds its HI state or advances to the next inferior one in each interval, the probabilities being estimated from observed condition data rather than from failure records, which are scarce for long-lived plants. Markov formulations of this kind have a long record in infrastructure

deterioration modelling (Jiang et al., 1988; Micevski et al., 2002; Li et al., 2016). Cesare et al. (1993) and Li et al. (2016) show further that maintenance can be represented inside the model itself, by inserting a repair rate into the transition matrix so that a proportion of a state's population is returned to a superior state each interval. Yahaya et al. (2017) brought the approach to transformer populations, and Yahaya et al. (2018) applied repair-rate adjustment to examine its effect on HI state distribution and on the performance condition curve. That second study supplies the HI-MPM and the data re-analyzed here.

Research Gap

Condition-forecasting studies in this tradition report the effect of intervention but stop short of its economic structure. Three decision-relevant quantities are left implicit even where sufficient information has been published to recover them: the marginal return to successive increments of intervention, the relative productivity of intervening at different points on the deterioration path, and the planning horizon each condition state affords. A fourth question — whether these properties survive a change in the baseline HI state distribution, which is rarely known precisely — is not posed at all. This paper addresses all four and, in doing so, proposes a general method for extracting allocation guidance from HI-based Markov models already in service.

RESEARCH METHODOLOGY

Source Model and Data Provenance

This is a secondary analysis of a HI-MPM published. We collected no new condition data and altered no part of the underlying model. The data, the Markov Prediction Model and the base outcomes are those published by Yahaya et al. (2018) under a Creative Commons Attribution 4.0 licenses, reproduced here with attribution for the purpose of re-analysis. That study drew 3,195 oil samples from 373 distribution transformers rated 33/11 kV and 30 MVA, aged 1 to 25 years, estimated the MPM transition probabilities by nonlinear minimization separately for two network zones, and applied pre-determined maintenance repair rates of 10%, 20% and 30% to examine their effect on HI state distribution and on the performance condition curve. The present contribution lies entirely in the analyses performed on that model, which are set out in the subsections below.

Health Index States and the Markov Prediction Model

The Health Index is expressed on the five-state scale of Table 1, each state carrying a defined maintenance action (Naderian et al., 2008; Jahromi et al., 2009). The scale is simultaneously a technical classification and an allocation rule, since each band represents a distinct claim on maintenance resource. In the MPM these five bands are the state space over which the chain evolves.

Table 1. Health index bands, condition states and corresponding maintenance action.

Health Index	Condition	State	Required Maintenance Action
85 – 100%	Very Good	1	Normal maintenance
70 – 84%	Good	2	Normal maintenance
50 – 69%	Fair	3	Increase diagnostic testing; remedial work or replacement depending on criticality
30 – 49%	Poor	4	Begin planning to replace or rebuild, considering risk and consequences of failure
0 – 29%	Very Poor	5	Immediately assess risk; replace or rebuild based on assessment

The MPM treats deterioration as a monotonic Markov process in which an asset either remains in its HI state or advances to the next inferior state in each annual interval, with state 5 absorbing. The HI state distribution

D(t) and the aggregate performance condition curve HI(t) then evolve as:

$$D(t+1) = D(t)\, P_m \qquad (1)$$

$$HI(t+1) = D(t+1)\, R^T, \quad R = [100 \;\; 84 \;\; 69 \;\; 49 \;\; 29] \qquad (2)$$

Maintenance enters the MPM by inserting a pre-determined repair rate ρ into row i of the transition matrix, redirecting a proportion of that HI state's population to a superior state j while scaling the remaining transitions proportionally (Cesare et al., 1993; Li et al., 2016):

$$P_m[i, j] = \rho, \quad P_m[i, k] = (1 - \rho)\, P[i, k] \;\; for \; k \neq j \qquad (3)$$

Table 2. Published transition probabilities by network zone (Yahaya et al., 2018).

Zone	P_{11} (Very good)	P_{22} (Good)	P_{33} (Fair)	P_{44} (Poor)
Zone 1	0.2889	0.7717	0.9626	0.99
Zone 2	0.5051	0.8923	0.7315	0.99

Re-Implementation and Verification

We re-implemented the HI-MPM independently in Python from equations (1) to (3) and the published transition matrices, then verified it twice. First, the matrix-update rule of equation (3) was checked against the updated matrices tabulated in the source study: for zone 1 at a 10% rate, the computed row for the good-to-very-good adjustment is [0.1000, 0.6945, 0.2055, 0, 0] and for the poor-to-fair adjustment [0, 0, 0.1000, 0.8910, 0.0090], both reproducing the published values exactly. Second, the benefit metric was identified by reproducing a published outcome: the reported 3.94% average improvement for intervention at the good state under a 10% rate is recovered as 3.99% when improvement is computed as the mean relative difference between the maintained and unmaintained performance curves over the 30-year horizon,

$$I(\rho) = (1/T)\, \Sigma_t \, [HI_m(t) - HI_0(t)] \,/\, HI_0(t) \times 100\% \qquad (4)$$

The source study does not report the baseline HI state distribution D(0) behind its performance-curve calculations. We tried to recover it by fitting the published outcomes and obtained a degenerate solution — roughly half the fleet in state 1 and half in state 5 — which cannot describe a real population. We therefore make no claim to reproduce every published figure exactly. Rather than assume the problem away, we turned it into the design: since D(0) is unidentified, the Monte Carlo analysis below asks whether the structural findings depend on it. They do not. That is a stronger result than point reproduction would have been.

Derived Measures and Robustness Design

Four measures are computed from the HI-MPM response. The marginal return to the repair rate over successive increments is

$$MR_i = [I(\rho_i) - I(\rho_{i-1})] \,/\, (\rho_i - \rho_{i-1}) \qquad (5)$$

The functional form is estimated by fitting a power law I(ρ) = Aρ^α in logarithms, where α < 1 indicates diminishing returns. The placement ratio compares benefit from an identical rate applied at two different states, $\Pi = I_{41}(\rho) \,/\, I_{21}(\rho)$, where the subscripts give the state repaired and its target state: I_{41} returns a proportion ρ of the state 4 (poor) population to state 1 (very good) and I_{21} returns state 2 (good) to state 1, so both assume perfect, as-good-as-new repair (j = 1 in equation (3)). Expected dwell time in state i, the mean number of annual intervals before transition under the geometric distribution implied by the Markov structure, is

$$E[T_i] = 1 \,/\, (1 - P_{ii}) \qquad (6)$$

We assessed robustness to the unidentified baseline distribution by Monte Carlo simulation, taking four thousand draws. Each draw paired a baseline HI state distribution sampled from a Dirichlet(0.8) over the five states — a deliberately dispersed prior, admitting fleets from near-new to heavily degraded — with a zone weighting drawn uniformly on [0,1]. For every draw we recomputed the marginal-return ratios and the placement ratio. We also swept the response surface at 2% intervals from 2% to 50%, which characterizes the functional form over a far wider range than the three published points allow.

Parametric Efficient Repair Rate

Where benefit follows I(ρ) = Aρ^α and intervention cost is approximately linear in intensity, C(ρ) = cρ, which is the usual case for labor- and materials-bound maintenance, the efficient repair rate ρ* is the one that maximises the net benefit NB(ρ) of intervention:

$$max\ NB(\rho) = I(\rho) - C(\rho) = A\rho^{\alpha} - c\rho, \quad 0 < \rho \leq 1 \qquad (7)$$

with A > 0, c > 0 and 0 < α < 1, where NB, I and C are measured in the same unit (percentage points of HI improvement), so that c is the cost of one unit of repair rate expressed in units of condition benefit. The first-order condition (FOC) gives the stationary point

$$\frac{dNB}{d\rho} = \alpha A\rho^{(\alpha-1)} - c = 0 \quad \Rightarrow \quad \rho^* = (\alpha A / c)^{[1/(1-\alpha)]} = (\alpha / \lambda)^{[1/(1-\alpha)]}, \quad \lambda = c / A \qquad (8)$$

and the second-order condition (SOC) confirms that it is a maximum:

$$\frac{d^2NB}{d\rho^2} = \alpha(\alpha - 1)A\rho^{(\alpha-2)} < 0 \quad for\ all\ \rho > 0 \qquad (9)$$

Because 0 < α < 1, NB(ρ) is strictly concave, so ρ* is the unique global maximum of the net benefit; with the fitted α = 0.438 the SOC holds at every repair rate. Should ρ* exceed 1, the constraint binds and ρ = 1 applies.

Equation (8) expresses the optimum entirely through the dimensionless ratio λ of cost to benefit scale. No absolute monetary values are required, only the organization's view of what a unit of condition improvement is worth relative to what it costs — a judgement utility can supply even where they cannot cost an outage.

RESULT AND DISCUSSION

A Square-Root Law in the Maintenance Repair Rate

Fitting a power law to the three published HI performance-curve improvements (10.24%, 14.23% and 16.48% at repair rates of 10%, 20% and 30%) gives I(ρ) = 28.24ρ^0.438 with R^2 = 0.994. Fitting the same form to the independently simulated response surface across 2% to 50% gives an exponent of 0.510 with R^2 = 0.924. Two estimates, arrived at by different routes, put the exponent near one half. HI improvement scales roughly as the square root of the maintenance repair rate (Figure 1).

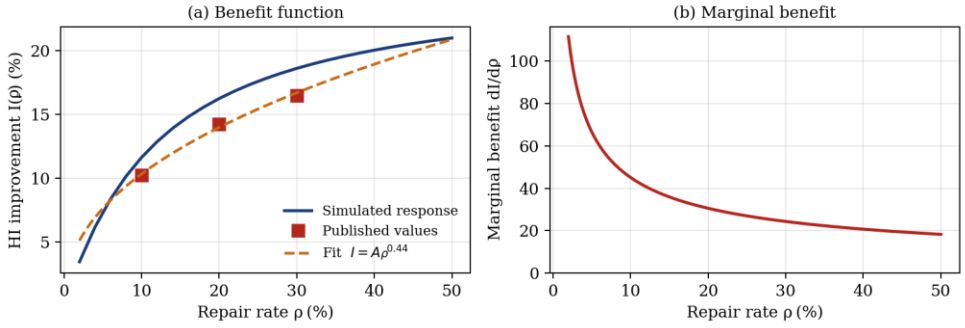


Figure 1. (a) Benefit function: simulated response surface, published outcomes and fitted power law. (b) Marginal benefit dI/dρ, showing rapid decay with intensity.

The marginal structure implied is steep (Table 3). The second decile of repair rate returns 36% of the effectiveness of the first and the third only 18%. Tripling the rate from 10% to 30% therefore raises HI improvement by only 61%, and by the fifth decile the marginal return has fallen to roughly 11% of its initial value.

Table 3. Marginal return to maintenance intensity derived from published outcomes.

Increment	Cumulative benefit I(ρ) (%)	Incremental gain (pp)	Marginal return (pp per 1% of ρ)	Relative to first decile
0 → 10%	10.24	10.24	1.024	1
10 → 20%	14.23	3.99	0.399	0.39
20 → 30%	16.48	2.25	0.225	0.22

Robustness to the Baseline HI State Distribution

Because the baseline HI state distribution is unidentified, we tested the structural findings across 4,000 randomized baseline distributions and zone weightings (Figure 2, Table 4). Every single draw showed diminishing returns. The dispersion is narrow: the second-decile ratio has an interquartile range of 0.34 to 0.37 and the third-decile ratio of 0.17 to 0.19, with 5th–95th percentile ranges of 0.31–0.38 and 0.15–0.20 respectively. The square-root law is therefore a property of the MPM transition structure rather than of any fleet, which is what permits its use as a planning heuristic.

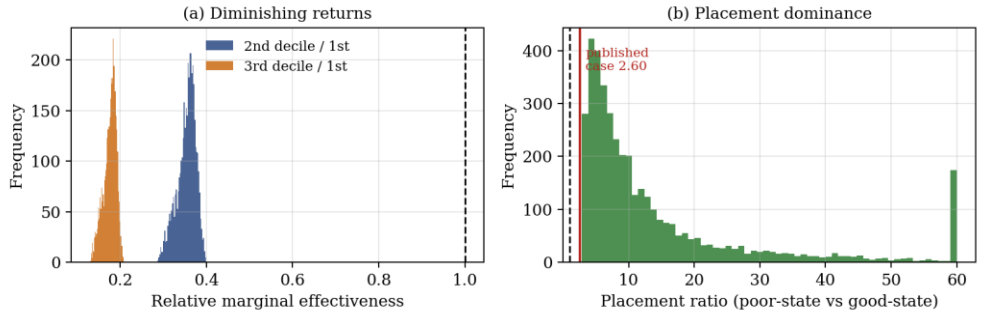


Figure 2. Monte Carlo robustness over 4,000 randomized initial distributions and zone weightings. (a) Relative marginal effectiveness of the second and third deciles. (b) Placement ratio; values above 60 are truncated for display.

Table 4. Robustness of structural findings over 4,000 randomized fleet compositions.

Quantity	Median	Interquartile range	Share of draws confirming
2nd decile / 1st decile	0.36	0.34 – 0.37	100% show decline
3rd decile / 1st decile	0.18	0.17 – 0.19	100% show decline
Placement ratio Π	8.8	5.49 – 16.16	100% exceed unity

The HI State at Which Repair Is Applied Dominates the Rate

At an identical 10% repair rate, the published outcomes give 10.24% HI improvement when repair is applied at state 4 (poor) against 3.94% at state 2 (good), a ratio of 2.60. Across randomized fleet compositions the median ratio is 8.80, with the 5th percentile at 3.67 — the published case is therefore conservative, and the advantage of targeting the degraded cohort is larger in most plausible fleets than the single reported comparison suggests.

Two mechanisms drive this. State 4 holds the largest share of the population under the baseline MPM forecast, so a given repair rate simply moves more assets. The chain is also path-dependent: an asset

returned from state 4 to state 1 re-enters the deterioration sequence at its start and keeps contributing to every later interval, whereas an asset returned from state 2 gains very little remaining path.

The placement result assumes perfect repair, in which a repaired state 4 asset re-enters the chain at state 1. In practice, refurbishment may restore a poor-state transformer only to state 3 (fair) or state 2 (good). This is represented in the same model by setting j = 3 or j = 2 in equation (3), and it reduces the placement ratio through the path-dependence mechanism above: the repaired asset re-enters the deterioration sequence part-way along it, gaining a smaller HI increment (20 or 35 points on the R scale of equation (2), rather than 51) and fewer subsequent intervals in superior states. Re-running the Monte Carlo design with these targets, the median placement ratio falls from 8.8 (repair to state 1) to about 7.6 (state 2; interquartile range 4.7–14.4) and 3.9 (state 3; 2.3–7.8), with 5th percentiles of 3.1 and 1.4; even with repair to state 3, 99% of draws still favour intervention at state 4. Placement therefore continues to dominate, but the size of its advantage depends materially on repair quality. The single published comparison of 2.60 should be read against the target state of the source scenario: the matrix-update check in the methodology uses a poor-to-fair (state 4 to state 3) adjustment, and against the corresponding distribution 2.60 lies in the lower quartile rather than below the 5th percentile.

Set against the intensity result, this gives us the paper's central allocation finding. Redirecting repair effort from state 2 to state 4 at a constant 10% rate buys 6.30 percentage points of extra HI improvement without increasing intervention volume at all. Tripling the repair rate at fixed placement buys 6.24 points — at three times the volume. The two gains are almost identical. One of them is free.

Efficient Repair Rate without Cost Data

Substituting the fitted parameters A = 28.24 and α = 0.438 into equation (8) yields the efficient repair rate as a function of the cost-to-benefit ratio alone (Figure 3). The curve falls steeply: at λ = 0.6 the efficient rate is 57%, at λ = 1.0 it is 23%, and at λ = 1.5 it is 11%. Intervention is only worth pursuing at high intensity where a unit of condition benefit is cheap relative to its value; across most of the plausible range the efficient rate sits at the low end of the interval examined in the source study.

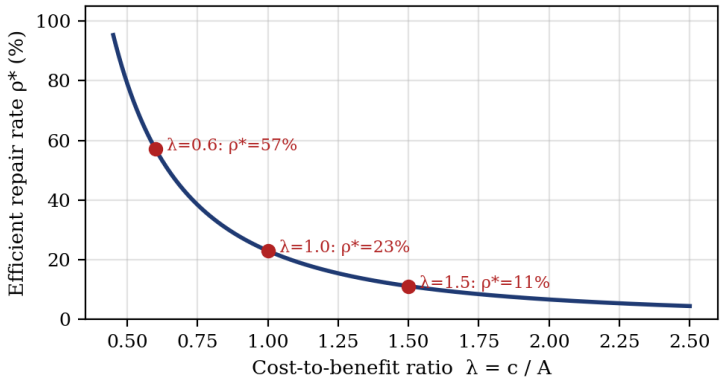


Figure 3. Efficient repair rate ρ* as a function of the cost-to-benefit ratio λ, from equation (8).

The practical value of equation (8) is that it inverts the usual data requirement. Rather than asking a utility to cost an outage, it asks only how one unit of cost compares with one unit of benefit — a relative judgement that can be elicited, debated and revised without pretending to a precision the underlying cost records do not support.

Expected Dwell Time per HI State and Zone Divergence

Applying equation (6) to the published MPM transition probabilities converts each diagonal element into an expected residence time per HI state (Table 5, Figure 4). The zones diverge sharply. In zone 1 an asset entering state 3 (fair) remains there for an expected 26.7 years; in zone 2 it passes through in 3.7 years, a

7.2-fold difference. The pattern reverses one state higher, zone 2 assets dwelling 9.3 years in state 2 (good) against 4.4 years in zone 1.

Table 5. Expected dwell time per condition state, derived from published transition probabilities.

Zone	State 1 (Very good)	State 2 (Good)	State 3 (Fair)	State 4 (Poor)
Zone 1	1.41 years	4.38 years	26.74 years	100 years
Zone 2	2.02 years	9.29 years	3.72 years	100 years

Note: The 100-year values for state 4 are not empirical operational lifespans. They follow from $E[T_4] = 1/(1 - P_{44})$ with the published $P_{44} = 0.99$ and mark a semi-absorbing transition boundary: the MPM allows only a 1% annual transition from state 4 to state 5 and no unaided return to a superior state, so state 4 is almost absorbing in the model.

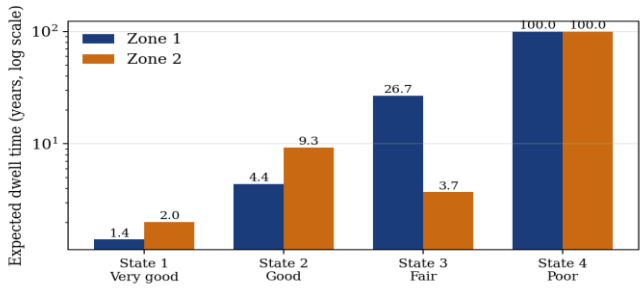


Figure 4. Expected dwell time per condition state by network zone (logarithmic scale).

Note: The 100-year bars for state 4 represent a semi-absorbing transition boundary implied by $P_{44} = 0.99$, not an empirical operational lifespan (see the note to Table 5).

No single fleet-wide inspection interval can serve both profiles. In zone 2, state 3 is a short warning window: an asset detected there reaches state 4 within roughly four years, so detection must be followed quickly by action. In zone 1 the same state is a long plateau during which assets accumulate, permitting less frequent inspection but producing a large standing cohort. The 100-year figure for state 4 in both zones follows from the published $P_{44} = 0.99$ and should be read as indicating that the MPM does not allow assets to leave state 4 unaided, not as a physical service-life expectation.

MANAGERIAL AND ECONOMIC IMPLICATIONS

Setting the Maintenance Repair Rate

The square-root law turns budget setting from a judgement into a calculation. Because HI improvement per unit of repair rate falls to roughly a fifth of its initial level by the third decile, incremental expenditure cases should be argued against marginal rather than cumulative benefit. The operational rule that follows is to fund the first decile of repair rate across as much of the fleet as possible before funding a deeper rate anywhere: the first unit applied to an untreated cohort outperforms the third unit applied to a treated one by a factor of about 4.5.

Directing the Budget

The dominance of placement over rate is the most actionable finding here. A utility can get the equivalent of a threefold rise in intervention volume simply by redirecting existing effort from the state 2 cohort to the state 4 cohort. This cuts against a familiar operational instinct — protect the assets that are still healthy, where intervention is cheaper and less disruptive. At population level that instinct is allocatively inefficient. It may still be right for individual assets of exceptional criticality, which a population-level model does not represent.

Differentiating Policy by Zone

Dwell times give each HI state an explicit planning horizon and show that the horizon is not uniform. Inspection intervals, crew deployment and the speed of escalation from detection to work order should differ between zones exhibiting a 7.2-fold difference in the persistence of state 3. A uniform policy necessarily over-inspects one zone and under-responds in the other, and both errors carry cost. Because the degraded cohort also accumulates faster in the fast-degrading zone, the placement advantage reported above is larger there, giving targeting a geographical as well as a condition dimension.

Methodological Transferability

The procedure used here is not specific to transformers, nor to this HI formulation. Any condition index coupled to a Markov transition model that can be evaluated at multiple repair rates yields a response surface from which the benefit exponent, the placement ratio and a parametric optimum can be recovered. For organizations that have invested in condition modelling but lack defensible cost data, this offers a route to allocation guidance that does not wait for a costing exercise that may never be completed to the required standard.

Boundary of the Economic Claims

It is worth being precise about what this study does not establish. The source dataset contains condition-monitoring records and modelled state distributions; it contains no maintenance costs, outage costs, replacement prices or customer-interruption values. The analysis establishes the shape of the benefit function and the relative productivity of intervention points, but not the position of the optimum in monetary terms, which requires the cost function the benefit must be set against. Equation (8) parameterizes rather than removes that requirement. No cost saving, payback period or return on investment is claimed, and none should be inferred.

LIMITATIONS AND FUTURE RESEARCH

Five limitations qualify what we report. The analysis inherits the assumptions of the source HI-MPM, monotonic deterioration and time-invariant transition probabilities; neither is strictly true of ageing plant and both would be worth relaxing. The transition probabilities derive from a single 373-unit population in one network, so the exponent estimated here should be treated as indicative for this asset class rather than universal until replicated. The initial state distribution used in the source performance-curve calculations is unidentified, which is why robustness rather than reproduction is the basis of the robustness claims; the Monte Carlo design mitigates but does not eliminate this. The repair rates examined were, as the source study states, selected to probe model behavior rather than drawn from an observed maintenance programmed, so the response surface describes the model rather than a realized intervention record. Finally, the linear cost assumption underlying equation (7) is plausible for labor- and materials-bound intervention but would not hold where intervention requires lumpy capital.

Four extensions follow. Replicating the exponent estimate on other asset classes and networks would establish whether the square-root law generalizes or is specific to this population. Pairing the response surface with observed maintenance costs would locate the optimum rather than parameterize it. Relaxing time-invariance to allow age-dependent transition probabilities would test whether the marginal structure is stable over the asset life cycle. And identifying whether the zone divergence in dwell time reflects environmental loading, operating regime or asset vintage would convert an observed difference into a manageable cause.

CONCLUSION

This paper set out to recover the economic structure of condition-based maintenance from the physical response of a Health Index coupled to a Markov Prediction Model, without the cost data that has long

impeded the practical adoption of maintenance optimization. Re-implementing and verifying the published HI-MPM of a 373-unit distribution transformer population, we reported four results.

Improvement in the HI performance condition curve scales as an approximate square root of the maintenance repair rate, $I(\rho) \propto \rho^{0.44}$ on published outcomes and $\rho^{0.51}$ on the simulated response surface. The implied marginal structure is steep: the second decile of repair rate returns 36% of the effectiveness of the first and the third 18%. Monte Carlo analyses over 4,000 randomized baseline HI state distributions confirm diminishing returns in every case with narrow dispersion, establishing the result as a property of the MPM transition structure rather than of a particular fleet.

Placement dominates rate. Repair directed at HI state 4 rather than state 2 yields a median benefit ratio of 8.8, and redirecting effort between states buys almost the same gain as tripling intervention volume at no additional cost. The power-law form admits a closed-form efficient repair rate expressed through the cost-to-benefit ratio alone, requiring no monetary data. And expected dwell times recovered from the MPM transition probabilities revealed a 7.2-fold divergence between network zones in the persistence of state 3, which no uniform inspection policy can serve efficiently.

For practice these combine into a single allocation rule for a constrained maintenance budget: spread a modest repair rate widely rather than concentrating a high rate narrowly, direct it at the state 4 cohort, and differentiate by zone. For research the wider contribution is methodological. The managerial value of an HI-based Markov model lies less in the condition trajectory it produces than in the marginal structure that trajectory implies, and that structure is recoverable from models already in service without new fieldwork and without costing data that organizations are frequently unable to supply.

ACKNOWLEDGMENTS

This study presents a secondary analysis of data and model outputs originally published in Yahaya et al. (2018), Sustainability, 10(10), 3399. These materials are reused under the Creative Commons Attribution 4.0 license, with appropriate attribution. The authors acknowledge the participating utility for providing the original condition-monitoring records and thank Universiti Teknikal Malaysia Melaka (UTeM) and the Ministry of Higher Education Malaysia for their support. The authors also acknowledge the utility company for its technical assistance. All data inputs used in this analysis are available in the cited open-access publication. The re-implementation code and Monte Carlo simulation procedure are available from the corresponding author upon reasonable request. The authors declare no competing interests.

REFERENCES

1. Alsumaidaee, Y.A.M., Chong, T.Y., Koh, S.P., Tiong, S.K., Chen, C.P., Ali, K., 2022. Review of medium voltage switchgear fault detection in a condition-based monitoring system by using deep learning. Energies 15 (18), 6762.
2. Azmi, A., Jasni, J., Azis, N., Kadir, M.Z.A., 2017. Evolution of transformer health index in the form of mathematical equation. Renewable and Sustainable Energy Reviews 76, 687–700.
3. Barlow, R., Hunter, L., 1960. Optimum preventive maintenance policies. Operations Research 8 (1), 90–100.
4. Brown, R.E., Humphrey, B.G., 2005. Asset management for transmission and distribution. IEEE Power and Energy Magazine 3 (3), 39–45.
5. Cesare, M.A., Santamarina, C., Turkstra, C., Vanmarcke, E.H., 1993. Modeling bridge deterioration with Markov chains. Journal of Transportation Engineering 118 (6), 820–833.
6. Dekker, R., 1996. Applications of maintenance optimization models: A review and analysis. Reliability Engineering & System Safety 51 (3), 229–240.
7. Haema, J., Phadungthin, R., 2012. Condition assessment of the health index for power transformers. In: Proceedings of the IEEE Power Engineering and Automation Conference (PEAM), pp. 1–4.
8. International Organization for Standardization, 2014. ISO 55000: Asset management — Overview, principles and terminology. ISO, Geneva.

9. Jahromi, A., Piercy, R., Cress, S., Service, J., Fan, W., 2009. An approach to power transformer asset management using health index. IEEE Electrical Insulation Magazine 25 (2), 20–34.
10. Jardine, A.K.S., Lin, D., Banjevic, D., 2006. A review on machinery diagnostics and prognostics implementing condition-based maintenance. Mechanical Systems and Signal Processing 20 (7), 1483–1510.
11. Jiang, Y., Saito, M., Sinha, K., 1988. Bridge performance prediction model using the Markov chain. Transportation Research Record 1180, 25–32.
12. Joskow, P.L., 2008. Incentive regulation and its application to electricity networks. Review of Network Economics 7 (4), 547–560.
13. Kulkarni, S., Salunke, M., Tatti, S., Patil, K., Joshi, A., 2021. Real time Ring Main Unit (RMU) monitoring system for distribution network. In: Proceedings of the 6th International Conference on Convergence in Technology (I2CT), pp. 1–5.
14. Li, L., Feng, L., Zhang, C., Lijun, S., 2016. Use of Markov chain model based on actual repair status to predict bridge deterioration in Shanghai, China. Transportation Research Record 2550, 106–114.
15. Micevski, T., Kuczera, G., Coombes, P., 2002. Markov model for storm water pipe deterioration. Journal of Infrastructure Systems 8 (2), 49–56.
16. Naderian, A., Cress, S., Piercy, R., Wang, F., Service, J., 2008. An approach to determine the health index of power transformers. In: Conference Record of the IEEE International Symposium on Electrical Insulation, pp. 192–196.
17. Panmala, N., Suwanasri, T., Suwanasri, C., 2022. Condition assessment of gas insulated switchgear using health index and conditional factor method. Energies 15, 9393.
18. Scatiggio, F., Pompili, M., 2013. Health index: the TERNA's practical approach for transformers fleet management. In: Proceedings of the IEEE Electrical Insulation Conference (EIC), pp. 178–182.
19. Schneider, J., Gaul, A.J., Neumann, C., Hogräfer, J., Wellssow, W., Schwan, M., Schnettler, A., 2006. Asset management techniques. International Journal of Electrical Power & Energy Systems 28 (9), 643–654.
20. Van Horenbeek, A., Pintelon, L., Muchiri, P., 2010. Maintenance optimization models and criteria. International Journal of System Assurance Engineering and Management 1 (3), 189–200.
21. Wang, H., 2002. A survey of maintenance policies of deteriorating systems. European Journal of Operational Research 139 (3), 469–489.
22. Yahaya, M.S., Azis, N., Ab Kadir, M.Z.A., Jasni, J., Hairi, M.H., Talib, M.A., 2017. Estimation of transformers health index based on the Markov chain. Energies 10 (11), 1824.
23. Yahaya, M.S., Azis, N., Selva, A.M., Ab Kadir, M.Z.A., Jasni, J., Hairi, M.H., Yang Ghazali, Y.Z., Talib, M.A., 2018. Effect of pre-determined maintenance repair rates on the health index state distribution and performance condition curve based on the Markov prediction model for sustainable transformers asset management strategies. Sustainability 10 (10), 3399.