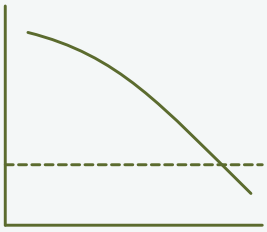
LESSON BOOK — VOLUME II



Asset Health and Condition Monitoring Volume II

Prognostics, remaining useful life, maintenance decisions and applications

Draft · sixteen lessons · 137 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Asset Health and Condition Monitoring

Volume II

Prognostics, remaining useful life, maintenance decisions and industrial applications

A lesson sequence, drafted as a book (skeleton draft)

Ricardo Marquez

November 12, 2026

How to read this book

This book is about one question: what the current condition of an engineering asset is, how that condition is changing, how deterioration can be detected, and how that information should influence maintenance and operational decisions. The answer is built in the order the discipline itself uses,

$$\text{measure} \rightarrow \text{assess} \rightarrow \text{diagnose} \rightarrow \text{predict} \rightarrow \text{decide}.$$

Volume I, Chapters 1 to 20, covered the first three steps: the foundations of condition monitoring, degradation mechanisms, measurements and signal analysis, health assessment, and fault detection and diagnostics. This second volume, Chapters 21 to 36, covers the last two: estimating degradation and predicting its course (Part VI), remaining useful life and condition-dependent risk (Part VII), maintenance policies and economic decisions (Part VIII), and industrial applications that bring the methods together (Part IX). Chapter numbers run continuously; a reference to a chapter numbered 20 or below points to Volume I, whose results this volume builds on and does not repeat.

The centrifugal pump and motor assembly of Volume I remains the teaching example. Part IX adds electrical and thermal equipment and closes with an integrated case study that runs from measurements to a maintenance decision. The numbers in the examples are illustrative.

State of this draft. This volume is a skeleton drafted to the plan of the lesson sequence: every lesson is in place with its sections, key equations, a worked example and exercises, ready to be expanded into book prose. Figures marked *to be drawn* are placeholders.

Contents

How to read this book ii

VI Degradation Estimation and Prognostics 1

21 Degradation Trend Estimation 2

21.1 The engineering question . . . 2
21.2 Observed trend and degradation trajectory . . . 2
21.3 The linear trend model . . . 3
21.4 Uncertainty of the estimated rate . . . 3
21.5 Residual checks . . . 4
21.6 Nonlinear trend models . . . 4
21.7 Detecting acceleration . . . 5
21.8 Model selection and overfitting . . . 5
21.9 Local trends and smoothing . . . 6
21.10 Confounding and maintenance resets . . . 6
21.11 Extrapolation to a limit . . . 6
21.12 Worked example — Trend of the pump bearing indicator . . . 7
21.13 What to retain . . . 8
21.14 Exercises . . . 8

22 Statistical and Stochastic Degradation Models 10

22.1 The engineering question . . . 10
22.2 Two kinds of variability . . . 10
22.3 Degradation path models . . . 11
22.4 Failure-time distribution from a random-rate model . . . 11
22.5 The Wiener process with drift . . . 11
22.6 First-passage time of the Wiener process . . . 12
22.7 The gamma process . . . 12
22.8 The inverse Gaussian process . . . 13
22.9 Estimating process parameters . . . 13
22.10 Covariates and random effects in process models . . . 13
22.11 Choosing a model . . . 14
22.12 Worked example — First-passage time of the bearing indicator . . . 14
22.13 What to retain . . . 15

22.14Exercises 16

23 Physics-Based and Data-Driven Prognostics 17
23.1 The engineering question 17
23.2 A classification of prognostic approaches 17
23.3 Physics-based prognostics 18
23.4 Example — Crack growth by the Paris law 18
23.5 Example — Bearing rating life 19
23.6 Data-driven prognostics 19
23.7 Run-to-failure data and censoring 20
23.8 Health-indicator construction for prognostics 20
23.9 Hybrid prognostics 21
23.10Comparing the approaches 21
23.11Selecting an approach 21
23.12Worked example — Three views of the pump bearing 22
23.13What to retain 23
23.14Exercises 23

24 Bayesian Updating and Sequential Estimation 25
24.1 The engineering question 25
24.2 The state-space model 25
24.3 The recursive Bayesian filter 26
24.4 Updating a degradation rate with a conjugate prior 26
24.5 The Kalman filter 27
24.6 A level-and-rate degradation model 27
24.7 Nonlinear models: the extended Kalman filter 27
24.8 The particle filter 28
24.9 Joint estimation of state and parameters 28
24.10From filtering to prediction 29
24.11Tuning and consistency checks 29
24.12Worked example — One filter step for the pump bearing 29
24.13What to retain 30
24.14Exercises 31

VII Remaining Useful Life and Risk 33

25 Remaining Useful Life: Definitions and Estimation 34
25.1 The engineering question 34
25.2 The formal definition 34
25.3 The end-of-life criterion 35
25.4 Remaining useful life as a distribution 35
25.5 Relation to mean residual life 36
25.6 Remaining useful life from a random-rate model 36

25.7 Remaining useful life from a stochastic process . . . 36
25.8 Monte Carlo estimation . . . 37
25.9 Future operating conditions . . . 37
25.10How remaining useful life evolves . . . 38
25.11What remaining useful life is not . . . 38
25.12Reporting remaining useful life . . . 38
25.13Worked example — Remaining useful life of the pump bearing . . . 38
25.14What to retain . . . 39
25.15Exercises . . . 40

26 Uncertainty and Prediction Intervals 41
26.1 The engineering question . . . 41
26.2 Sources of prognostic uncertainty . . . 41
26.3 Aleatory and epistemic uncertainty . . . 42
26.4 Confidence intervals and prediction intervals . . . 42
26.5 Methods of uncertainty propagation . . . 42
26.6 Variance of the remaining useful life . . . 43
26.7 Uncertainty in the limit . . . 43
26.8 One-sided bounds . . . 44
26.9 Coverage and calibration . . . 44
26.10How uncertainty should evolve . . . 44
26.11Model-form uncertainty . . . 45
26.12Communicating a prognosis . . . 45
26.13Worked example — Prediction interval and RUL uncertainty for the pump bearing 45
26.14What to retain . . . 47
26.15Exercises . . . 47

27 Condition-Dependent Failure Risk 49
27.1 The engineering question . . . 49
27.2 The hazard function . . . 49
27.3 Age-based conditional failure probability . . . 50
27.4 The proportional hazards model . . . 50
27.5 Interpreting the coefficients . . . 50
27.6 Time-varying covariates . . . 51
27.7 Estimation . . . 51
27.8 Two routes to a failure probability . . . 51
27.9 Health indices and failure probability . . . 52
27.10Risk . . . 52
27.11Probability of failure over a decision horizon . . . 53
27.12Worked example — Failure risk of a slurry pump bearing . . . 53
27.13What to retain . . . 54
27.14Exercises . . . 54

28 Prognostic Validation and Performance Metrics **56**
28.1 The engineering question . . . 56
28.2 How prognostic validation differs from diagnostic validation . . . 56
28.3 Ground truth for remaining useful life . . . 57
28.4 Prediction error . . . 57
28.5 Accuracy metrics . . . 57
28.6 Relative accuracy . . . 58
28.7 Prognostic horizon . . . 58
28.8 Alpha-lambda accuracy . . . 58
28.9 Convergence and stability . . . 59
28.10Evaluating probabilistic predictions . . . 59
28.11Asymmetric scoring . . . 59
28.12Validation design . . . 59
28.13Decision-oriented validation . . . 60
28.14Worked example — Scoring a bearing prognosis . . . 60
28.15What to retain . . . 61
28.16Exercises . . . 62

VIII Maintenance Strategies and Optimization **64**

29 Condition-Based Maintenance Policies **65**
29.1 The engineering question . . . 65
29.2 The elements of a maintenance decision . . . 65
29.3 The P-F curve and the P-F interval . . . 66
29.4 Control-limit policies . . . 66
29.5 The long-run cost rate . . . 66
29.6 Lead time . . . 67
29.7 Remaining-life-based policies . . . 67
29.8 Risk-based intervention . . . 68
29.9 Imperfect maintenance . . . 68
29.10Markov decision models . . . 68
29.11Imperfect monitoring . . . 69
29.12Opportunistic and grouped maintenance . . . 69
29.13Worked example — Choosing a preventive threshold for the pump bearing . . . 69
29.14What to retain . . . 70
29.15Exercises . . . 70

30 Inspection Scheduling and Monitoring Intervals **72**
30.1 The engineering question . . . 72
30.2 Continuous and periodic monitoring . . . 72
30.3 The inspection interval and the P-F interval . . . 73
30.4 Imperfect inspection . . . 73
30.5 A random P-F interval . . . 73

30.6 The delay-time model . . . 74
30.7 Probability that a defect causes a failure . . . 74
30.8 The cost model and the optimal interval . . . 74
30.9 Estimating the delay-time model . . . 75
30.10 Hidden failures and failure-finding intervals . . . 75
30.11 Condition-dependent inspection intervals . . . 76
30.12 Sampling for trend estimation . . . 76
30.13 Costs and risks of inspection . . . 76
30.14 Worked example — Vibration route interval for the pump fleet . . . 77
30.15 What to retain . . . 78
30.16 Exercises . . . 78

31 Repair, Replacement, and Economic Decisions 80
31.1 The engineering question . . . 80
31.2 Cost categories . . . 80
31.3 The time value of money . . . 81
31.4 Equivalent annual cost . . . 81
31.5 Age replacement . . . 81
31.6 Block replacement . . . 82
31.7 Minimal repair and periodic replacement . . . 82
31.8 Economic life . . . 83
31.9 Repair or replace with condition information . . . 83
31.10 Sensitivity analysis . . . 83
31.11 Non-economic considerations . . . 84
31.12 Worked example — Overhaul or replace the pump . . . 84
31.13 What to retain . . . 85
31.14 Exercises . . . 85

32 Risk-Based Asset Prioritization 87
32.1 The engineering question . . . 87
32.2 Risk as the basis for priority . . . 87
32.3 Probability of failure . . . 88
32.4 Consequence of failure . . . 88
32.5 Risk matrices . . . 88
32.6 Failure modes, effects, and criticality analysis . . . 89
32.7 Quantitative risk ranking . . . 89
32.8 Risk reduction and benefit-cost ratio . . . 89
32.9 Selection under a budget . . . 90
32.10 Redundancy and system effects . . . 90
32.11 Uncertainty in the ranking . . . 91
32.12 Relation to asset-management standards . . . 91
32.13 Worked example — Allocating a maintenance budget . . . 91
32.14 What to retain . . . 92
32.15 Exercises . . . 93

IX Industrial Applications and Integrated Case Studies 95

33 Rotating Equipment: Pumps, Motors, and Bearings 96
33.1 The engineering question . . . 96
33.2 The machine train and its failure modes . . . 96
33.3 Organizing the program . . . 97
33.4 Rolling-element bearings: stages of degradation . . . 97
33.5 Bearing indicators and prognosis . . . 98
33.6 Hydraulic degradation of the pump . . . 98
33.7 Cavitation and seals . . . 99
33.8 Shaft, coupling, and alignment . . . 99
33.9 The induction motor . . . 99
33.10From indicators to actions . . . 100
33.11Worked example — The pump bearing enters its accelerating phase . . . 100
33.12What to retain . . . 101
33.13Exercises . . . 102

34 Electrical Equipment: Transformers, Cables, and Batteries 104
34.1 The engineering question . . . 104
34.2 How electrical assets differ . . . 104
34.3 Transformer insulation thermal aging . . . 105
34.4 What thermal loss of life does and does not say . . . 105
34.5 Dissolved gas analysis . . . 105
34.6 Other transformer indicators . . . 106
34.7 Transformer health indices . . . 106
34.8 Power cables . . . 107
34.9 Batteries: state of health . . . 107
34.10Estimating state of health . . . 108
34.11Battery degradation models and remaining life . . . 108
34.12Worked example — Transformer insulation loss of life . . . 109
34.13What to retain . . . 110
34.14Exercises . . . 110

35 Thermal and Process Equipment: Chillers and Heat Exchangers 112
35.1 The engineering question . . . 112
35.2 Performance-based monitoring . . . 112
35.3 Heat-exchanger performance . . . 113
35.4 Data validation and uncertainty . . . 113
35.5 Fouling models . . . 113
35.6 Cleaning as a renewal . . . 114
35.7 Hydraulic effects and tube integrity . . . 114
35.8 Chiller performance indicators . . . 115
35.9 Chiller faults and signatures . . . 115
35.10Prognosis for reversible degradation . . . 116

35.11Worked example — Condenser fouling and cleaning interval 116
35.12What to retain 117
35.13Exercises 118

36 Integrated Case Study: From Measurements to Maintenance Decision 120
36.1 The engineering question 120
36.2 The chain of reasoning 120
36.3 Function, failure modes, and criticality 121
36.4 Measurements and data quality 121
36.5 Indicators and reference 121
36.6 Detection 121
36.7 Diagnosis 122
36.8 Trend and state estimation 122
36.9 Remaining useful life at month 6 123
36.10The first decision 123
36.11New evidence 124
36.12The revised decision 124
36.13Outcome and review 125
36.14What the case shows 125
36.15What to retain 125
36.16Exercises 126
36.17End of Volume II 127

Part VI

Degradation Estimation and Prognostics

Chapter 21

Degradation Trend Estimation

21.1 The engineering question

Volume I ended with a monitoring system that can detect abnormal behavior, diagnose its likely cause, and demonstrate that its conclusions are trustworthy.

Volume II asks what happens next.

Consider the centrifugal pump that has served as the running example. Its drive-end bearing has been diagnosed with a developing defect (Lesson 19). The normalized vibration indicator (Lesson 14) has been recorded once a month for six months, and each value is higher than the last.

The maintenance planner now asks:

1. How fast is the bearing deteriorating?
2. How well is that rate known?
3. Is the deterioration steady or accelerating?
4. When might the indicator reach its limit?

The central engineering question is:

> **Given a sequence of imperfect condition measurements, how can engineers estimate the underlying degradation trend, quantify the uncertainty of the estimated rate, and determine whether deterioration is accelerating?**

This lesson introduces regression-based trend estimation, the standard error of an estimated rate, nonlinear trend models, tests for acceleration, and the distinction between an observed trend and the physical degradation trajectory.

21.2 Observed trend and degradation trajectory

Lesson 8 described degradation by a state $d(t)$. Lesson 9 described a measurement as the state plus error.

For a degradation indicator recorded at times $t_1, \ldots, t_n$, the two ideas combine into

$$\boxed{y_k = x(t_k) + \epsilon_k,}$$

where:

- $x(t)$: the underlying degradation trajectory, expressed in the units of the indicator;
- y_k: the recorded indicator at time t_k;
- ϵ_k: measurement and normalization error.

Three objects must be kept apart:

- The **observations** y_k are what was recorded.
- The **fitted trend** is a curve estimated from the observations.
- The **degradation trajectory** $x(t)$ is the physical quantity the curve is meant to represent.

A fitted trend is an estimate. It inherits error from the measurements, from the choice of curve, and from any operating-condition effect that normalization did not remove.

21.3 The linear trend model

The simplest trend model is the linear degradation model of Lesson 8, now with unknown coefficients:

$$x(t) = a + bt.$$

Here a is the indicator level at $t = 0$ and b is the degradation rate.
The ordinary least-squares estimates minimize the sum of squared residuals

$$S(a, b) = \sum_{k=1}^{n} (y_k - a - bt_k)^2 .$$

Setting the derivatives to zero gives

$$\boxed{\hat{b} = \frac{S_{ty}}{S_{tt}}, \qquad \hat{a} = \bar{y} - \hat{b}\,\bar{t},}$$

with

$$S_{tt} = \sum_k (t_k - \bar{t})^2, \qquad S_{ty} = \sum_k (t_k - \bar{t})(y_k - \bar{y}).$$

The estimate $\hat{b}$ is a weighted average of the observed changes. Observations far from the mean time $\bar{t}$ carry the most weight.

21.4 Uncertainty of the estimated rate

The estimated rate is a random quantity because the observations contain error.
Assume independent errors with common variance σ^2. The residual variance is estimated by

$$s^2 = \frac{1}{n-2} \sum_{k=1}^{n} \left(y_k - \hat{a} - \hat{b}t_k \right)^2 .$$

The standard error of the rate is

$$\boxed{\mathrm{se}(\hat{b}) = \frac{s}{\sqrt{S_{tt}}}.}$$

A $100(1-\alpha)\%$ confidence interval for the rate is

$$\hat{b} \pm t_{1-\alpha/2,\,n-2}\ \mathrm{se}(\hat{b}),$$

where $t_{1-\alpha/2,\,n-2}$ is a Student-t quantile. The interpretation is the frequentist one given in Lesson 16.

Three practical consequences follow from the formula:

1. A longer observation window increases S_{tt} and reduces the standard error.
2. More observations in the same window help, but less than extending the window.
3. Noisier indicators (larger s) give less certain rates.

For equally spaced observations with spacing Δ, $S_{tt} = \Delta^2 n(n^2-1)/12$. Doubling the window length at fixed n halves the standard error.

21.5 Residual checks

The standard error is valid only if the model assumptions hold. The residuals

$$e_k = y_k - \hat{a} - \hat{b}t_k$$

should be examined for:

- **Curvature.** A run of negative residuals in the middle and positive residuals at both ends indicates that a straight line is the wrong shape.
- **Changing scatter.** Scatter that grows with the indicator level suggests a multiplicative error, which a logarithmic transformation may stabilize.
- **Serial correlation.** Correlated residuals (Lesson 18) make the standard error too small.
- **Outliers.** A single corrupted record (Lesson 10) can dominate a short series.

  Figure: Two panels. Left: indicator versus time with the fitted straight line. Right: residuals versus time with a zero line, showing a curved pattern for a misfit and random scatter for an adequate fit.

21.6 Nonlinear trend models

Lesson 8 introduced power-law and exponential degradation. Both can be estimated by regression after transformation.

Exponential trend. If

$$x(t) = x_0 e^{kt},$$

then

$$\ln x(t) = \ln x_0 + kt.$$

A straight-line fit of $\ln y_k$ against t_k estimates k as the slope. The quantity $\ln 2/k$ is the doubling time of the indicator.

Power-law trend. If

$$x(t) = at^b,$$

then

$$\ln x(t) = \ln a + b \ln t,$$

and a straight-line fit of $\ln y_k$ against $\ln t_k$ estimates the exponent b.

Direct nonlinear least squares. When the model includes an offset, such as $x(t) = c + at^b$, no transformation makes it linear. The coefficients are then found by numerical minimization of the sum of squared residuals.

The transformation changes the error assumption. Fitting $\ln y$ assumes errors proportional to the level; fitting y assumes errors of constant size. The residual plot indicates which is closer to the data.

21.7 Detecting acceleration

Whether deterioration is accelerating matters more than the current rate, because acceleration shortens the time to any limit.

Three standard checks are used.

Quadratic term. Fit

$$x(t) = a + bt + ct^2$$

and test whether c differs from zero using $\hat{c}/\operatorname{se}(\hat{c})$ against a Student-t distribution with $n - 3$ degrees of freedom. A positive, statistically significant $\hat{c}$ indicates acceleration.

Split-window slopes. Estimate the rate separately on the first and second halves of the record. With independent halves, the difference $\hat{b}_2 - \hat{b}_1$ has standard error $\sqrt{\operatorname{se}(\hat{b}_1)^2 + \operatorname{se}(\hat{b}_2)^2}$.

Model comparison. Compare a linear fit with an exponential or power-law fit on held-out later observations (Section 21.8).

A short, noisy record rarely supports a firm conclusion about acceleration. The honest statement is then that the data are consistent with both a constant and an increasing rate.

21.8 Model selection and overfitting

A model with more coefficients always fits past observations at least as well. That does not make it a better predictor.

Two safeguards are standard.

Information criteria. For least-squares fits with Gaussian errors, the Akaike information criterion is

$$\mathrm{AIC} = n \ln\!\left(\frac{\mathrm{SSE}}{n}\right) + 2p,$$

where SSE is the sum of squared residuals and p is the number of estimated coefficients. Lower is better. The term $2p$ penalizes added coefficients.

Out-of-sample prediction. Fit each candidate to the earlier observations and compare predictions with the later ones. This is the time-ordered validation recommended in Lesson 20.

Physical plausibility comes first. A fitted trend that predicts recovery of irreversible damage should be rejected whatever its fit statistics (Lesson 8).

21.9 Local trends and smoothing

Degradation rates change. A single line through the entire history can hide a recent change. Common local estimators are:

- **Moving-window regression.** Fit the line to the most recent m observations only. Small m responds quickly but gives a noisy rate.
- **Weighted regression.** Give recent observations larger weights, in the spirit of the exponentially weighted moving average of Lesson 18.
- **Piecewise-linear models.** Fit separate lines before and after a change point (Lesson 18).

The window length trades responsiveness against variance, the same tradeoff that governed detection in Lesson 18.

21.10 Confounding and maintenance resets

A rising indicator is evidence of degradation only if other explanations have been removed. Before estimating a trend, check that:

1. The indicator has been normalized for operating point (Lesson 14). A gradual seasonal change in load or temperature produces a trend with no degradation behind it.
2. The sensor has not drifted (Lesson 9).
3. The measurement procedure has not changed (Lesson 10).
4. Maintenance events are marked. Lubrication, alignment, or cleaning can reset the indicator. A trend must be fitted within one segment between such events, not across them.

   Figure: Indicator versus time with two vertical lines marking maintenance events; three separate fitted segments, each restarting from a lower level.

21.11 Extrapolation to a limit

For a linear trend and a limit x_F (Lesson 8), the projected time at which the fitted line reaches the limit, measured from the current time t_n, is

$$\hat{\tau} = \frac{x_F - \hat{x}(t_n)}{\hat{b}}, \qquad \hat{x}(t_n) = \hat{a} + \hat{b}t_n.$$

Because $\hat{b}$ appears in the denominator, a modest relative error in the rate produces a similar relative error in $\hat{\tau}$. A rate known to within 10% gives a projected time known to no better than about 10%, before any other source of uncertainty is counted.

This projection is not yet a remaining useful life. It ignores future variability of the degradation process, the uncertainty of the current level, and the possibility that the model form is wrong. Lessons 22 to 26 add these in turn.

21.12 Worked example — Trend of the pump bearing indicator

The normalized bearing vibration indicator of the pump is recorded monthly. All values are illustrative.

Month t_k	Indicator y_k (mm/s)
0	2.0
1	2.3
2	2.5
3	2.9
4	3.1
5	3.4

The site uses an indicator limit of $x_F = 7.1$ mm/s as its end-of-useful-life criterion for this bearing.

21.12.1 Step 1 — Means and sums

$$\bar{t} = 2.5, \qquad \bar{y} = \frac{16.2}{6} = 2.7.$$

$$S_{tt} = 6.25 + 2.25 + 0.25 + 0.25 + 2.25 + 6.25 = 17.5.$$

$$S_{ty} = 1.75 + 0.60 + 0.10 + 0.10 + 0.60 + 1.75 = 4.9.$$

21.12.2 Step 2 — Rate and intercept

$$\hat{b} = \frac{4.9}{17.5} = 0.28 \text{ mm/s per month}, \qquad \hat{a} = 2.7 - 0.28(2.5) = 2.0 \text{ mm/s}.$$

21.12.3 Step 3 — Residual variance

The fitted values are 2.00, 2.28, 2.56, 2.84, 3.12, 3.40. The residuals are 0, 0.02, −0.06, 0.06, −0.02, 0.

$$\text{SSE} = 0.008, \qquad s^2 = \frac{0.008}{4} = 0.002, \qquad s = 0.0447 \text{ mm/s}.$$

21.12.4 Step 4 — Standard error and confidence interval

$$\text{se}(\hat{b}) = \frac{0.0447}{\sqrt{17.5}} = 0.0107.$$

With $t_{0.975,4} = 2.776$,

$$\boxed{\hat{b} = 0.28 \pm 0.030 \text{ mm/s per month} \quad (95\%).}$$

21.12.5 Step 5 — Projection to the limit

The fitted level at month 5 is 3.40 mm/s.

$$\hat{\tau} = \frac{7.1 - 3.40}{0.28} = 13.2 \text{ months.}$$

Using the ends of the rate interval, 0.250 and 0.310, the projection ranges from $3.7/0.310 = 11.9$ to $3.7/0.250 = 14.8$ months.

21.12.6 Step 6 — Interpret

The indicator is rising at about 0.28 mm/s per month, and the rate is known to roughly ±11%.

The residuals show no clear curvature, but six points cannot exclude acceleration.

The range of 11.9 to 14.8 months reflects rate uncertainty only, under a straight-line model. It is a projection of the fitted line, not a statement about when the bearing will fail.

21.13 What to retain

1. Observations, the fitted trend, and the physical degradation trajectory are three different things.
2. Least squares gives the rate as $\hat{b} = S_{ty}/S_{tt}$.
3. The standard error of the rate is $s/\sqrt{S_{tt}}$; a longer window reduces it faster than denser sampling.
4. Residuals must be checked for curvature, changing scatter, serial correlation, and outliers.
5. Exponential and power-law trends become linear after logarithmic transformation.
6. Acceleration is assessed with a quadratic term, split-window slopes, or out-of-sample comparison.
7. Extra coefficients improve fit, not necessarily prediction.
8. Trends must be fitted on normalized indicators and within maintenance segments.
9. A projected time to a limit inherits the relative uncertainty of the rate and is not yet a remaining useful life.

21.14 Exercises

Exercise 21.1 — Least-squares rate

An indicator is recorded at $t = 0, 1, 2, 3, 4$ months with values 1.0, 1.3, 1.5, 1.8, 2.0 mm/s.

1. Calculate $\bar{t}$, $\bar{y}$, S_{tt}, and S_{ty}.
2. Calculate the rate $\hat{b}$ and intercept $\hat{a}$.
3. Calculate the residuals and the residual standard deviation s.

Exercise 21.2 — Rate uncertainty

A linear trend fitted to $n = 8$ observations gives $\hat{b} = 0.40$ units per month, $s = 0.12$ units, and $S_{tt} = 28$ month^2.

1. Calculate $\text{se}(\hat{b})$.
2. Calculate the 95% confidence interval for the rate, using $t_{0.975,6} = 2.447$.
3. The current fitted level is 3.0 units and the limit is 7.0 units. Calculate the projected time to the limit at the point estimate and at both ends of the interval.

Exercise 21.3 — Exponential trend

An indicator is 1.5 mm/s at $t = 0$ and 2.25 mm/s at $t = 4$ months. Assume $x(t) = x_0 e^{kt}$.

1. Calculate k.
2. Calculate the doubling time.
3. Calculate the time, from $t = 0$, at which the indicator reaches 6.0 mm/s.
4. Calculate the same time under a straight line through the two points, and explain the difference.

Exercise 21.4 — Trend or artifact

A pump's raw bearing vibration rises steadily from March to August and falls from September to February. Pump speed is higher in summer.

1. Explain why a line fitted to the March-to-August data is not evidence of degradation.
2. State which Volume I method should be applied before trend estimation.
3. Describe what the normalized indicator would show if the bearing were also degrading.

21.14.1 Where Lesson 22 takes us

Lesson 22 — Statistical and Stochastic Degradation Models

A fitted line describes one unit's past. It does not describe how much the future path can wander, or how units of the same design differ from one another.

The central question will be:

How can degradation be modeled as a random process, so that the time to reach a limit has a probability distribution?

We will study degradation path models with random effects, the Wiener and gamma processes, and first-passage times.

Chapter 22

Statistical and Stochastic Degradation Models

22.1 The engineering question

Lesson 21 fitted a straight line to six monthly observations of the pump bearing indicator and projected that the line reaches the limit in about 13 months.

Two things were missing from that projection.

First, the bearing will not follow the fitted line exactly. Its future path will wander around any trend.

Second, the plant operates other pumps of the same design. Their bearings degrade at different rates, and their histories contain information about this one.

The central engineering question is:

> **How can degradation be represented as a random process, so that differences between units and variability along one unit's path are both described, and the time to reach a limit has a probability distribution?**

This lesson introduces degradation path models with random effects, the Wiener process with drift, the gamma process, the inverse Gaussian process, and the first-passage time.

22.2 Two kinds of variability

Lesson 8 listed sources of uncertainty in degradation. For modeling, they fall into two groups.

Unit-to-unit variability. Nominally identical units degrade at different rates because of differences in material, assembly, alignment, and service. One unit's rate is fixed but unknown.

Temporal variability. Along one unit's life, the increments of degradation fluctuate because load, lubrication, and environment fluctuate.

A model with only unit-to-unit variability produces smooth paths that fan out from the origin. A model with only temporal variability produces irregular paths around a common trend. Real populations show both.

> **Figure:** Two panels of degradation versus time for ten units with a horizontal limit line. Left: straight lines with different slopes (random effects). Right: irregular paths around one common trend (stochastic process).

Measurement error (Lesson 9) is a third, separate contribution. It affects what is recorded, not the degradation itself.

22.3 Degradation path models

The general degradation path model writes the observation on unit i at time t_{ik} as

$$\boxed{y_{ik} = g(t_{ik}; \phi, \theta_i) + \epsilon_{ik},}$$

where:

- g: the path function, such as a linear, power-law, or exponential form (Lesson 8);
- ϕ: fixed effects, common to all units;
- θ_i: random effects, specific to unit i and drawn from a population distribution;
- ϵ_{ik}: measurement error.

This is a random-effects (mixed-effects) regression model. It is the standard statistical model for repeated degradation measurements on several units.

The simplest case is a linear path through the origin with a random rate:

$$x_i(t) = \theta_i t.$$

22.4 Failure-time distribution from a random-rate model

Let the limit be x_F. Under $x_i(t) = \theta_i t$, unit i reaches the limit at

$$T_i = \frac{x_F}{\theta_i}.$$

The failure-time distribution follows directly from the distribution of the rate:

$$\boxed{F_T(t) = P(T \le t) = P\left(\theta \ge \frac{x_F}{t}\right).}$$

If $\theta \sim \mathcal{N}(\mu_\theta, \sigma_\theta^2)$ with negligible probability of a negative rate,

$$F_T(t) = 1 - \Phi\left(\frac{x_F/t - \mu_\theta}{\sigma_\theta}\right).$$

If θ is lognormal, T is also lognormal, because $\ln T = \ln x_F - \ln \theta$.

This construction links degradation data to the lifetime distributions of reliability engineering. A lifetime distribution can be derived from degradation measurements before any unit has failed.

22.5 The Wiener process with drift

Lesson 8 introduced the model $d(t) = d_0 + rt + \sigma W(t)$. In the standard notation of stochastic degradation modeling it is written

$$\boxed{X(t) = x_0 + \mu t + \sigma B(t),}$$

where $B(t)$ is standard Brownian motion, μ is the drift, and σ is the diffusion coefficient.
Its defining properties are:

- Increments over disjoint intervals are independent.
- The increment over an interval of length Δt is Gaussian:

$$X(t+\Delta t) - X(t) \sim \mathcal{N}(\mu\,\Delta t,\ \sigma^2 \Delta t).$$

- Paths are continuous but not monotone.

The non-monotone paths make the Wiener process suitable for indicators that fluctuate, such as vibration levels, and questionable for strictly irreversible quantities such as wear depth.

22.6 First-passage time of the Wiener process

The first-passage time to the limit is

$$T = \inf\{t \geq 0 : X(t) \geq x_F\}.$$

Let $d = x_F - x_0 > 0$ be the distance to the limit and $\mu > 0$. Then T has an inverse Gaussian distribution with density

$$\boxed{f_T(t) = \frac{d}{\sigma\sqrt{2\pi t^3}} \exp\left[-\frac{(d-\mu t)^2}{2\sigma^2 t}\right], \qquad t > 0.}$$

Its mean and variance are

$$\boxed{E[T] = \frac{d}{\mu}, \qquad \mathrm{Var}(T) = \frac{d\,\sigma^2}{\mu^3}.}$$

Its cumulative distribution function is

$$F_T(t) = \Phi\left(\frac{\mu t - d}{\sigma\sqrt{t}}\right) + \exp\left(\frac{2\mu d}{\sigma^2}\right)\Phi\left(-\frac{\mu t + d}{\sigma\sqrt{t}}\right).$$

The mean equals the deterministic time to threshold of Lesson 8. The variance is new: it grows with the distance to the limit and falls steeply as the drift increases.
Lesson 8 distinguished $P(X(t) \geq x_F)$ from $P(T \leq t)$. The first term of F_T is the former; the second term is the correction for paths that crossed the limit and came back.

22.7 The gamma process

For monotone degradation such as wear, corrosion, and crack growth, the gamma process is the standard choice.
A gamma process with shape rate α and scale β has independent increments with

$$\boxed{X(t+\Delta t) - X(t) \sim \mathrm{Gamma}(\mathrm{shape} = \alpha\,\Delta t,\ \mathrm{scale} = \beta).}$$

Increments are never negative, so paths never decrease.

Starting from $X(0) = 0$,

$$E[X(t)] = \alpha\beta t, \qquad \mathrm{Var}(X(t)) = \alpha\beta^2 t.$$

Because paths are monotone, the two probabilities of Lesson 8 coincide:

$$\boxed{P(T \le t) = P(X(t) \ge x_F).}$$

The right-hand side is a gamma tail probability with shape αt and scale β, available in any statistical table or library.

A nonlinear mean trend is obtained by replacing αt with a shape function $\alpha(t)$, for example αt^b.

22.8 The inverse Gaussian process

The inverse Gaussian process is a second monotone model. Its increments over an interval Δt follow an inverse Gaussian distribution with mean proportional to Δt.

It is used when gamma increments fit poorly, and it accommodates random effects and covariates conveniently. Its use parallels that of the gamma process: paths are monotone and $P(T \le t) = P(X(t) \ge x_F)$.

The name is shared with the first-passage distribution of Section 22.6, but the roles differ. There the inverse Gaussian distribution described a time. Here it describes degradation increments.

22.9 Estimating process parameters

Process parameters are estimated from observed increments.

For a Wiener process observed at times $t_0 < t_1 < \cdots < t_n$ without measurement error, let $\Delta x_k = x_k - x_{k-1}$ and $\Delta t_k = t_k - t_{k-1}$. The maximum-likelihood estimates are

$$\boxed{\hat{\mu} = \frac{\sum_k \Delta x_k}{\sum_k \Delta t_k} = \frac{x_n - x_0}{t_n - t_0}, \qquad \hat{\sigma}^2 = \frac{1}{n}\sum_{k=1}^{n} \frac{(\Delta x_k - \hat{\mu}\,\Delta t_k)^2}{\Delta t_k}.}$$

The drift estimate uses only the first and last observations. The diffusion estimate uses the scatter of the increments.

For a gamma process, the method of moments matches the mean and variance of the increments, and maximum likelihood refines the result.

When measurement error is not negligible, the recorded increments mix process variability with measurement noise. Ignoring this inflates $\hat{\sigma}^2$. Separating the two requires a state-space formulation, which Lesson 24 provides.

22.10 Covariates and random effects in process models

Lesson 7 showed that degradation rate depends on operating stress. Stochastic models carry this dependence in the drift or shape:

$$\mu(u) = \mu_0 \exp(\gamma^{\mathsf{T}} u),$$

where u is a vector of operating covariates such as load or temperature.

Unit-to-unit variability enters by making the drift random:

$$\mu_i \sim \mathcal{N}(\mu_\mu, \sigma_\mu^2).$$

A Wiener process with random drift combines both kinds of variability from Section 22.2. Its drift for a particular unit can be updated as that unit's data arrive, which is the subject of Lesson 24.

22.11 Choosing a model

Model	Paths	Typical use
Random-rate path	Smooth, differ by unit	Wear, capacity fade across a fleet
Wiener process	Fluctuating	Vibration and other noisy indicators
Gamma process	Monotone, jumps	Wear, corrosion, crack growth
Inverse Gaussian process	Monotone	Alternative to gamma

The choice follows the questions of Lesson 8:

1. Is the physical quantity monotone?
2. Is the dominant variability between units or along the path?
3. Are there enough units to estimate a population distribution?
4. Does the model reproduce the scatter of observed increments?

A stochastic model does not remove model-form uncertainty. A Wiener process with constant drift is still a linear-trend model, with the limitations discussed in Lesson 21.

22.12 Worked example — First-passage time of the bearing indicator

The pump bearing indicator is modeled as a Wiener process. All values are illustrative.

- Current level: $x_0 = 3.4$ mm/s (month 5 of Lesson 21).
- Limit: $x_F = 7.1$ mm/s.
- Drift: $\mu = 0.28$ mm/s per month.
- Diffusion coefficient: $\sigma = 0.15$ mm/s per $\sqrt{\text{month}}$, estimated from the plant's earlier bearing histories.

22.12.1 Step 1 — Distance to the limit

$$d = 7.1 - 3.4 = 3.7 \text{ mm/s}.$$

22.12.2 Step 2 — Mean first-passage time

$$E[T] = \frac{3.7}{0.28} = 13.2 \text{ months}.$$

22.12.3 Step 3 — Variance and standard deviation

$$\mathrm{Var}(T) = \frac{3.7(0.15)^2}{(0.28)^3} = \frac{0.08325}{0.021952} = 3.79 \text{ months}^2,$$

$$\boxed{\mathrm{sd}(T) = 1.95 \text{ months.}}$$

22.12.4 Step 4 — Probability of crossing within 10 months

The first term of $F_T(10)$ is

$$\Phi\left(\frac{0.28(10) - 3.7}{0.15\sqrt{10}}\right) = \Phi(-1.90) = 0.029.$$

The second term evaluates to about 0.005. Therefore

$$\boxed{P(T \le 10 \text{ months}) \approx 0.034.}$$

22.12.5 Step 5 — A monotone comparison

Lesson 8 used an impeller wear rate of 0.10 mm/year. A gamma process with $\alpha = 4$ per year and $\beta = 0.025$ mm has the same mean rate, $\alpha\beta = 0.10$ mm/year.

After four years,

$$E[X(4)] = 0.40 \text{ mm}, \qquad \mathrm{sd}(X(4)) = \sqrt{4(4)(0.025)^2} = 0.10 \text{ mm}.$$

22.12.6 Step 6 — Interpret

The mean first-passage time reproduces the 13.2 months of Lesson 21. The model adds a spread of about two months caused by path variability alone.

The drift was treated as known. Its uncertainty, shown in Lesson 21 to be about ±11%, is not yet included. Lesson 26 combines the two.

22.13 What to retain

1. Degradation variability has two parts: differences between units and fluctuation along one path.
2. Degradation path models describe unit differences through random effects.
3. A random-rate model gives the failure-time distribution as $P(\theta \ge x_F/t)$.
4. The Wiener process has independent Gaussian increments and non-monotone paths.
5. Its first-passage time to a fixed limit is inverse Gaussian, with mean d/μ and variance $d\sigma^2/\mu^3$.
6. The gamma and inverse Gaussian processes are monotone, so $P(T \le t) = P(X(t) \ge x_F)$.
7. Drift is estimated from the total change; diffusion from the scatter of increments.
8. Measurement error inflates the estimated process variability unless it is modeled separately.
9. A stochastic model with constant drift still assumes a linear mean trend.

22.14 Exercises

Exercise 22.1 — Wiener first-passage time

A degradation indicator follows a Wiener process with $\mu = 0.25$ units per month and $\sigma = 0.20$ units per $\sqrt{\text{month}}$. The distance to the limit is $d = 2.0$ units.

1. Calculate the mean first-passage time.
2. Calculate its variance and standard deviation.
3. Recalculate the standard deviation if the drift doubles, and explain the result.

Exercise 22.2 — Gamma process

Wear depth follows a gamma process with $\alpha = 5$ per year and $\beta = 0.04$ mm.

1. Calculate the mean wear rate.
2. Calculate the mean and standard deviation of wear after 5 years.
3. Explain why $P(T \le t) = P(X(t) \ge x_F)$ holds for this model but not for a Wiener process.

Exercise 22.3 — Parameter estimation

A Wiener-type indicator is observed at unit time steps. Five successive increments are 0.3, 0.2, 0.4, 0.1, and 0.5 units.

1. Calculate $\hat{\mu}$.
2. Calculate $\hat{\sigma}^2$ and $\hat{\sigma}$.
3. Explain how unmodeled measurement error would affect $\hat{\sigma}^2$.

Exercise 22.4 — Random-rate model

Wear follows $x_i(t) = \theta_i t$ with $\theta \sim \mathcal{N}(0.10,\, 0.02^2)$ mm/year. The limit is $x_F = 1.0$ mm.

1. Calculate the time to the limit for a unit with the mean rate.
2. Calculate $P(T \le 8 \text{ years})$, using $\Phi(1.25) = 0.894$.
3. A particular unit has measured wear of 0.45 mm after 3 years. State how this information should change the assessment for that unit.

22.14.1 Where Lesson 23 takes us

Lesson 23 — Physics-Based and Data-Driven Prognostics

The models of this lesson are statistical. They describe how degradation varies, but not why.

The central question will be:

When should a prognostic model be built from the physics of the degradation mechanism, when from recorded data, and when from both?

We will compare physics-based, data-driven, and hybrid approaches and the evidence each requires.

Chapter 23

Physics-Based and Data-Driven Prognostics

23.1 The engineering question

Lessons 21 and 22 predicted the bearing indicator by extending what has been observed. Neither used any knowledge of why a bearing degrades.

That knowledge exists. Lessons 5 to 7 described wear laws, fatigue, crack growth, and thermal aging. Bearing manufacturers publish rating-life formulas. These describe how degradation depends on load, speed, and temperature.

On the other hand, a physical model needs inputs that may be unknown for a machine in service, while the plant may hold years of recorded histories from similar machines.

The central engineering question is:

> **When should a prognostic model be built from the physics of the degradation mechanism, when from recorded data, and when from a combination of the two?**

This lesson defines the three families of prognostic methods, gives representative examples of each, and states what evidence each requires.

23.2 A classification of prognostic approaches

Prognostic methods are conventionally grouped by the source of their predictive structure.

Approach	Predictive structure from	Main requirement
Reliability-based	Lifetime statistics of a population	Failure records
Physics-based	A model of the mechanism	Loads, material data
Data-driven	Recorded condition histories	Representative histories
Hybrid	Both	Both, in smaller amounts

Reliability-based methods use age alone. They predict for a population, not for a particular unit, and are the subject of Lesson 27.

The other three use the condition of the individual unit. They are the methods of prognostics and health management (PHM).

All of them fit the framework of the book's central objects. A prognostic model is a state equation

$$x_{t+1} = f(x_t, u_t, w_t)$$

together with a measurement equation $y_t = h(x_t, u_t) + \epsilon_t$. The approaches differ in where f comes from.

23.3 Physics-based prognostics

A physics-based model writes f from the degradation mechanism. Volume I introduced several such laws:

- Archard's wear law (Lessons 5 and 6): wear volume proportional to load and sliding distance.
- The Arrhenius relationship (Lesson 7): reaction-driven aging accelerating with temperature.
- Miner's rule (Lesson 6): fatigue damage accumulating as cycle fractions.
- Crack-growth laws (Lesson 6).
- Fouling deposition-removal balances (Lesson 6).

The strengths of physics-based models are:

1. They predict under operating conditions that have not been observed, within the validity range of the law.
2. They need little or no failure data.
3. Their parameters have physical meaning and can be checked.

Their limitations are:

1. The mechanism must be known and must remain the dominant one.
2. Loads, material properties, and initial condition must be known or estimated.
3. A detailed model can be expensive to build and validate for each asset type.

23.4 Example — Crack growth by the Paris law

The Paris law relates crack growth per load cycle to the stress-intensity range:

$$\boxed{\frac{da}{dN} = C(\Delta K)^m, \qquad \Delta K = Y\,\Delta\sigma\sqrt{\pi a},}$$

where:

- a: crack length;
- N: number of load cycles;
- $\Delta\sigma$: stress range;
- Y: geometry factor;
- C, m: material constants.

For constant Y, constant $\Delta\sigma$, and $m \neq 2$, integration from a_0 to a critical length a_c gives the number of cycles

$$\boxed{N = \frac{a_c^{\,1-m/2} - a_0^{\,1-m/2}}{(1 - m/2)\, C\, (Y \Delta\sigma \sqrt{\pi})^m}.}$$

As an illustration, take $m = 3$, $C = 1 \times 10^{-11}$ (crack length in m, ΔK in MPa$\sqrt{\text{m}}$), $Y = 1$, $\Delta\sigma = 100$ MPa, $a_0 = 1$ mm, and $a_c = 10$ mm.

$$(Y \Delta\sigma \sqrt{\pi})^3 = (177.2)^3 = 5.57 \times 10^6,$$

$$N = \frac{2\,(0.001^{-1/2} - 0.01^{-1/2})}{(1 \times 10^{-11})(5.57 \times 10^6)} = \frac{2(31.62 - 10.00)}{5.57 \times 10^{-5}} \approx 7.8 \times 10^5 \text{ cycles.}$$

Two features are typical of physics-based prognosis. The current condition a_0 enters directly, so an inspection result updates the prediction. And the result is very sensitive to the stress range: with $m = 3$, a 10% increase in $\Delta\sigma$ reduces N by about 25%.

23.5 Example — Bearing rating life

The basic rating life of a rolling-element bearing is

$$\boxed{L_{10} = \left(\frac{C}{P}\right)^p \text{ million revolutions,}}$$

where:

- C: basic dynamic load rating;
- P: equivalent dynamic bearing load;
- $p = 3$ for ball bearings and $p = 10/3$ for roller bearings.

L_{10} is the life that 90% of a population of identical bearings is expected to reach or exceed under the stated load. The formula is standardized in ISO 281.

In operating hours at rotational speed n in rpm,

$$L_{10h} = \frac{10^6}{60\, n}\, L_{10}.$$

The rating life is physics-informed, since it rests on rolling-contact fatigue, but it is a population statistic. It says nothing about whether a particular bearing is among the early or late failures. It contains no measurement from the bearing in service.

Its value in prognostics is the load dependence. It states how life scales when the load changes.

23.6 Data-driven prognostics

A data-driven model learns f, or a direct mapping from measurements to remaining life, from recorded histories.

Three families are in common use.

Trend extrapolation. A degradation indicator is fitted with a trend or stochastic process and extended to a limit. Lessons 21 and 22 are data-driven methods of this kind.

END OF THE PREVIEW

Asset Health and Condition Monitoring ~~Volume~~ II

These are the opening pages. The complete book runs 137 pages.

The full PDF opens with a token. Write to **rmarquez123@gmail.com**, name the book, and I will send you one.

One token opens every book on the site.

Ricardo Marquez, Ph.D.

ricardo-marquez.com