Gamma distribution — life as a count of damage events, and the spares a mission needs

Menu: Life Data ▸ Time Dependence ▸ Gamma distribution…

Part 1 — What the function does

What it calculates. The gamma model reads a life as the time it takes to collect k random damage events that arrive, on average, one every θ time units: T ~ Gamma(k, θ). ReliaEn fits the shape k (the number of shocks a unit survives) and the scale θ (the mean spacing between shocks) by maximum likelihood, with bounds, and gives the mean life MTTF = k·θ, the median, the B-lives and the service intervals. It then tells you whether the units wear out, which engineering lever moves B10 more, how much a larger test would narrow the bound, and - when you give a mission - how many cold-standby spares reach your target.

What ReliaEn prints	What it tells you
$\hat{k}$ and $\hat{\theta}$ with bounds	How many damage events a unit survives, and how often they arrive
MTTF = k·θ, **median**, **B10**	The life figures of the fitted model
Shape verdict	Wear-out, early failure, or not proven - read from the bounds on k
Competitor models	Whether a Weibull or a lognormal explains the data as well, by AIC
The shock story and **value of more data**	Which fix pays more - tolerating one more hit or slowing the hits - and what a larger test would buy
Standby design	The smallest number of cold spares that reaches a mission's target reliability

From which data. The failure ages in one worksheet column, one row per failed unit, at least two, each above 0; survivor ages, if any, in a different column.

When an engineer uses it. When failure comes from accumulated damage - load cycles, particle loads, cold-standby units used one after another - and the question is how long units last, which lever extends their life, or how many spares a mission needs.

When not to use it.

Situation	Use instead
Failure comes from the weakest of many flaws	Life Data ▸ Time Dependence ▸ Weibull (2P) distribution…
Failure comes from growth that multiplies (a fatigue crack)	Life Data ▸ Time Dependence ▸ Lognormal distribution…
The failure rate is known to be constant	Life Data ▸ Estimate Failure Function ▸ Complete data…

Part 2 — The example's input data

Setting. A water plant runs 25 filter cartridges until each one clogs; each cartridge takes a series of particle loads before it fails. The reliability engineer wants their life, whether they wear out, and how many spare cartridges keep a 100-hour mission at 90% reliability when one runs at a time.

Input	Value
Failure ages, operating hours (25)	34.2, 37.3, 50.1, … (the 25 values of the file)
File in the worked-example zip	2-Gamma-Distribution/Example-1.csv — one column, heading Failure times, 25 rows
Mission time (standby)	100 — typed in the form
Standby target reliability %	90 — typed in the form
Confidence % / B-life level %	95 / 10 — the values the form opens with

	① A	B
1	34.2	
2	37.3	
3	50.1	
4	58.3	
5	65.5	
6	66.0	
7	77.0	
8	77.9	
9	79.6	
10	107.5	
11	115.9	
12	121.7	
13	129.5	
14	139.9	
15	143.8	
16	149.3	
17	153.8	
18	164.8	
19	168.2	
20	178.4	
21	180.2	
22	186.9	
23	191.5	
24	304.7	
25	402.8	
26		

The example data in the worksheet

- ① Column A: the 25 failure ages, one per row, in operating hours.

Part 3 — How to use it, step by step

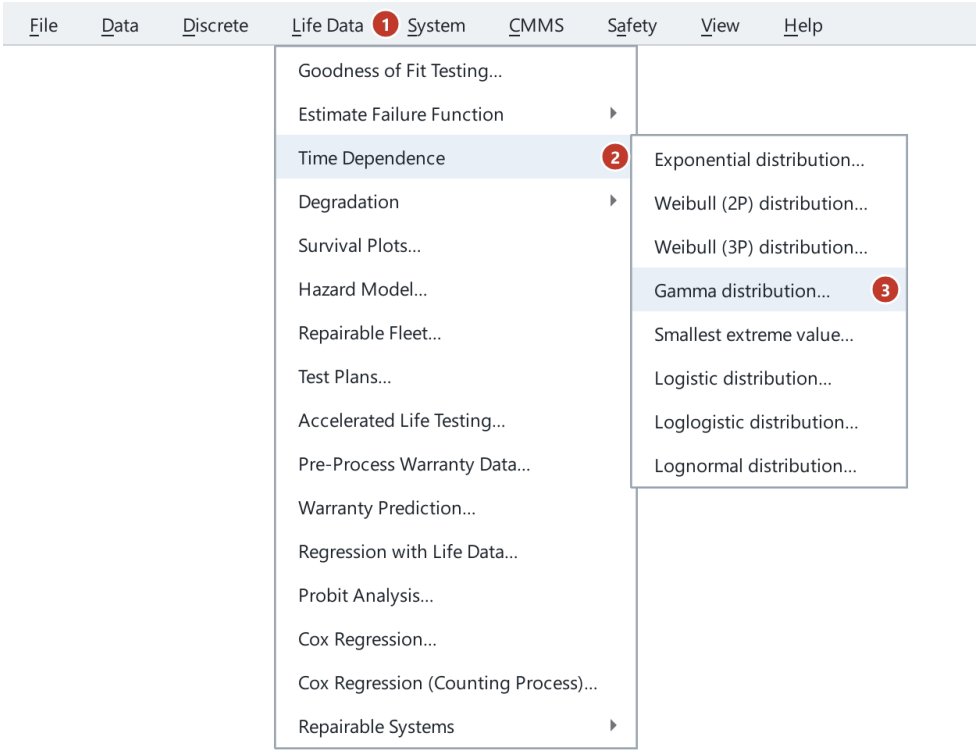
1. **Put the failure ages in column A.** Copy the 25 values of Example-1.csv that stand under its heading Failure times into column A of the worksheet, starting at row 1.

	① A	B
1	34.2	
2	37.3	
3	50.1	
4	58.3	
5	65.5	
6	66.0	
7	77.0	
8	77.9	
9	79.6	
10	107.5	
11	115.9	
12	121.7	
13	129.5	
14	139.9	
15	143.8	
16	149.3	
17	153.8	
18	164.8	
19	168.2	
20	178.4	
21	180.2	
22	186.9	
23	191.5	
24	304.7	
25	402.8	
26		

Step 1 — the failure ages in column A

- ① Column A: the 25 failure ages, 34.2 to 402.8.

2. **Choose Life Data ▸ Time Dependence ▸ Gamma distribution...** in the menu bar. The form Gamma Analysis (MLE) opens in the left panel.



Step 2 — the menu path

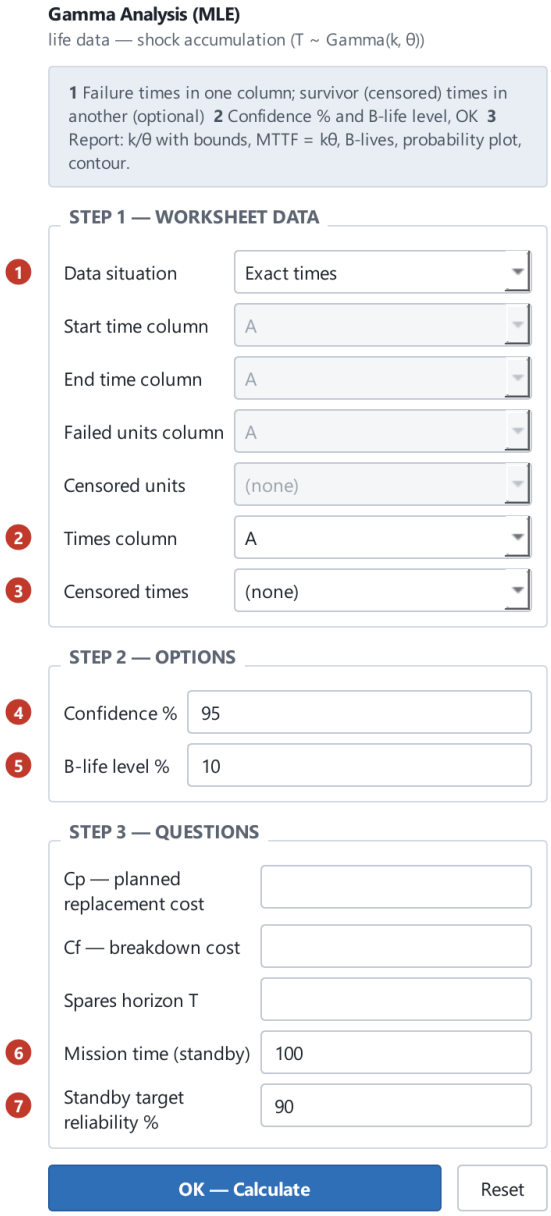
- ① Life Data: the menu.
- ② Time Dependence: the submenu.
- ③ Gamma distribution...: the entry to click.

The form as it opens, before any box is set:



Step 2 — the form as it opens

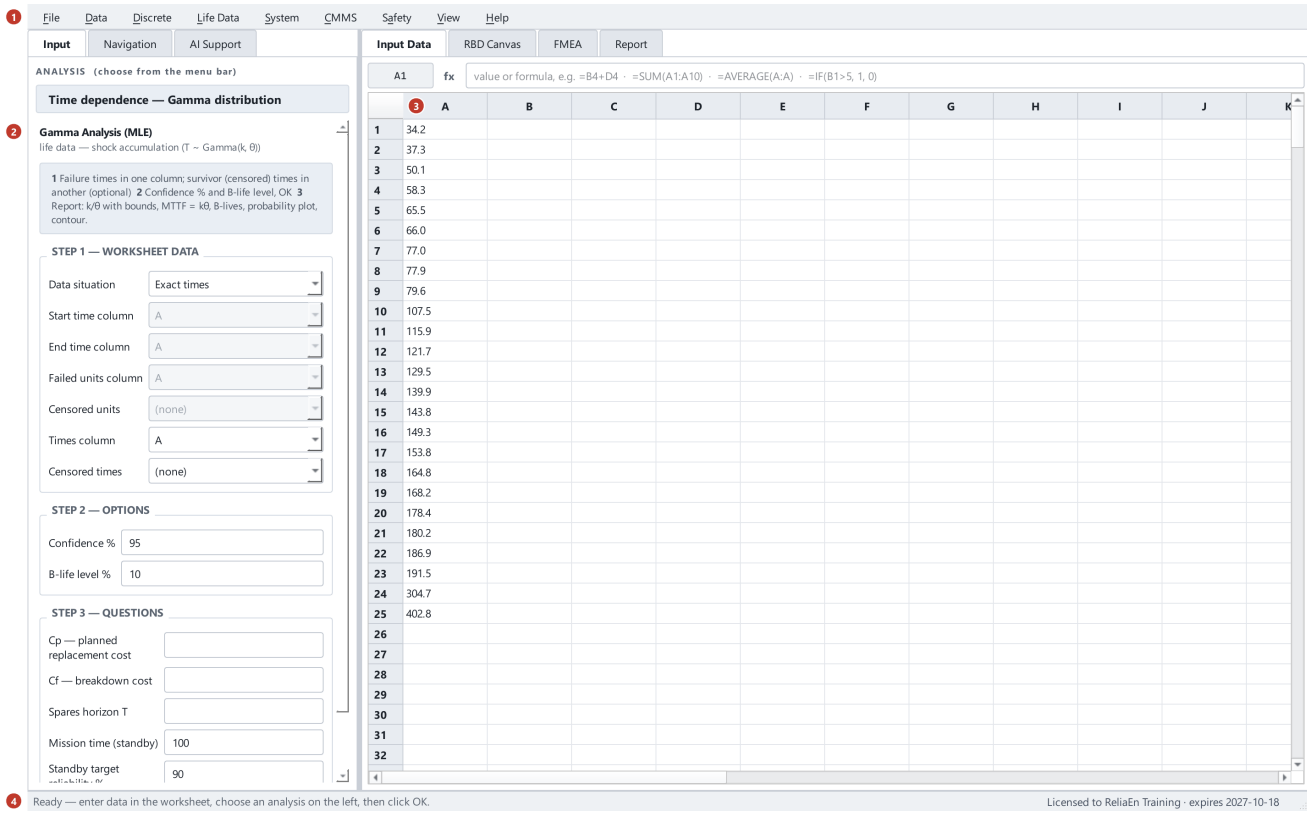
Steps 3 to 9 set the form's seven boxes, marked ① to ⑦ in the picture:



Steps 3 to 9 — the form set for this example

3. ① **Check that Data situation reads Exact times.** One age per unit. The four boxes under it, from Start time column to Censored units, are for inspection intervals and stay greyed out.
4. ② **In Times column, choose A.** The column that holds the failure ages.
5. ③ **Leave Censored times at (none).** Every cartridge was run until it clogged.
6. ④ **Check that Confidence % reads 95.** It sets every bound on the page.
7. ⑤ **Check that B-life level % reads 10.** The B-life printed in the results.
8. ⑥ **In Mission time (standby), type 100.** The mission the standby designer sizes the spares for. Cp — planned replacement cost, Cf — breakdown cost and Spares horizon T stay empty in this example.
9. ⑦ **In Standby target reliability %, type 90.** The reliability the mission must reach, as a percent.

10. **Click OK — Calculate** at the foot of the form (② in the picture below). ReliaEn calculates and shows the report.



Step 10 — the window just before OK

- ① The menu bar: File, Data, Discrete, Life Data, System, CMMS, View, Help.
- ② The form, set as in steps 3 to 9; OK — Calculate at its foot.
- ③ Column A: the 25 failure ages.
- ④ The status bar: Ready — enter data in the worksheet, choose an analysis on the left, then click OK. and Licensed to ReliaEn Training · expires 2027-10-18.

Part 4 — The report, card by card

Life Data — Gamma (2P) — Time Dependence

2026-11-13 15:09:07

T ~ Gamma(k, θ) ; R(t) = Q(k, t/θ) ; B_q = θ·P^{-1}(k, q) ; MTTF = k·θ

The report's head

- The title Life Data — Gamma (2P) — Time Dependence, the date and time of the run, and the model in one line: T ~ Gamma(k, θ); R(t) = Q(k, t/θ); B_q = θ·P^{-1}(k, q); MTTF = k·θ.

INPUT DATA		
1	n — units on test	25
2	r — failures	25
3	censored (survivors)	0
4	Estimation	profile MLE (complete data)

Input data

- ① n — units on test 25.
- ② r — failures 25. Check r first: a blank cell is skipped.
- ③ censored (survivors) 0.
- ④ Estimation profile MLE (complete data): the fit keeps $\hat{k}\cdot\hat{\theta}$ equal to the average failure age and searches k alone.

RESULTS		
1	$\hat{k}$ — shape (number of shocks)	3.08544 [1.82136, 5.22683]
2	$\hat{\theta}$ — scale (mean shock spacing)	43.881 [24.7563, 77.7796]
3	MTTF = k·θ	135.39
4	Median life	121.08
5	B10 life	50.67 [33.61, 76.38]

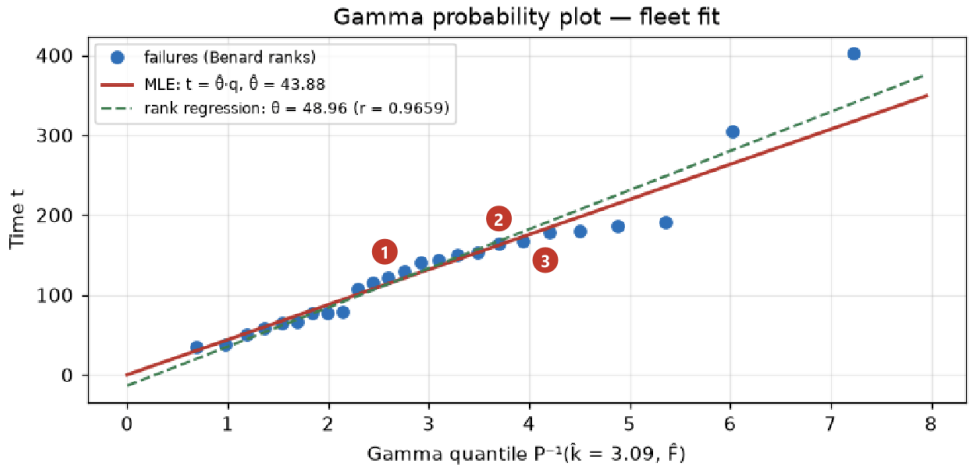
Results

- ① $\hat{k}$ — shape (number of shocks) 3.08544 [1.82136, 5.22683]: a cartridge survives about three particle loads.
- ② $\hat{\theta}$ — scale (mean shock spacing) 43.881 [24.7563, 77.7796]: one load about every 44 hours.
- ③ MTTF = k·θ 135.39: exactly the average of the 25 ages.
- ④ Median life 121.08: half the cartridges clog by this age.
- ⑤ B10 life 50.67 [33.61, 76.38]: one cartridge in ten has clogged by this age.

INTERPRETATION	
1	Shock-accumulation model: the unit fails after k ≈ 3.09 random damage events arriving about one per 43.88 time units — T ~ Gamma(k, θ), R(t) = Q(k, t/θ). MTTF = kθ = 135.39; 10% of units fail by B10 = 50.67. Physics slot: shocks/standby (gamma) vs weakest link (Weibull) vs multiplicative growth (lognormal).

Interpretation

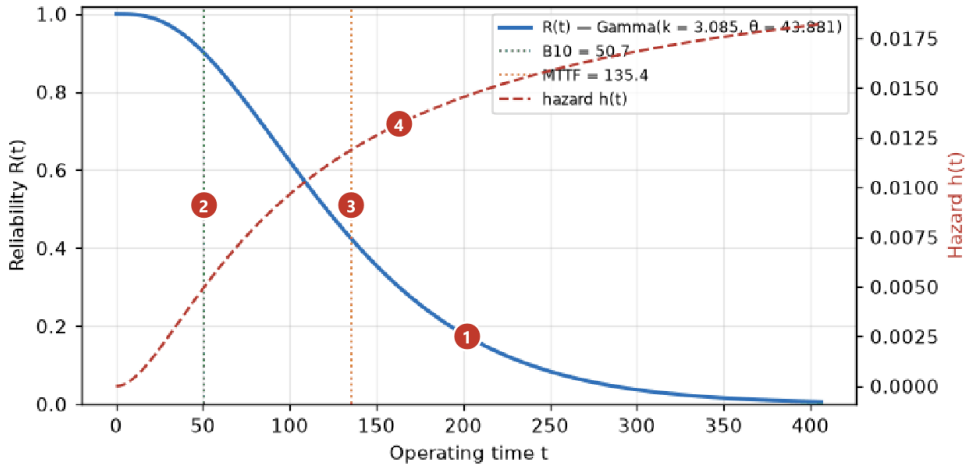
- ① ReliaEn's reading: Shock-accumulation model: the unit fails after k ≈ 3.09 random damage events arriving about one per 43.88 time units — T ~ Gamma(k, θ), R(t) = Q(k, t/θ). MTTF = kθ = 135.39; 10% of units fail by B10 = 50.67. Physics slot: shocks/standby (gamma) vs weakest link (Weibull) vs multiplicative growth (lognormal).



Interactive Plot

The gamma probability plot

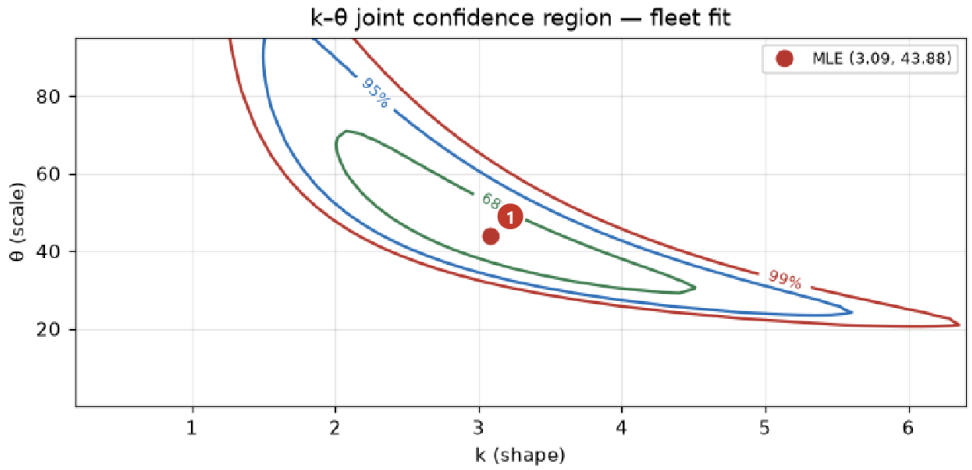
- ① failures (Benard ranks): the 25 ages on the gamma probability plot at the fitted shape.
- ② MLE: t = θ̂·q, θ̂ = 43.88: the fitted line, through the origin.
- ③ rank regression: θ = 48.96 (r = 0.9659): the plot's own straight-line fit, for comparison.



Interactive Plot

Reliability and hazard

- ① R(t) — Gamma(k = 3.085, θ = 43.881): the share of cartridges still running at each age.
- ② B10 = 50.7: where R crosses 90%.
- ③ MTTF = 135.4: the mean life on the curve.
- ④ hazard h(t): rising with age toward a ceiling of 1/θ̂.



Interactive Plot

The joint confidence region of k and θ

- ① MLE (3.09, 43.88): the estimate, inside its 68%, 95% and 99% regions; k and θ trade against each other along the region's long axis.

SHAPE VERDICT — WHAT k SAYS (FLEET FIT)

1 WEAR-OUT with confidence: k̂ = 3.085 and even its lower bound (1.821) exceeds 1 — failure risk grows with age. NOTE the gamma nuance: the hazard rises toward a CEILING of 1/θ̂ = 0.02279, unlike the Weibull hazard which grows without bound — old units approach a constant (worst) failure rate, not an ever-increasing one.

Shape verdict — what k says (fleet fit)

- ① WEAR-OUT with confidence: k̂ = 3.085 and even its lower bound (1.821) exceeds 1 — failure risk grows with age. NOTE the gamma nuance: the hazard rises toward a CEILING of 1/θ̂ = 0.02279, unlike the Weibull hazard which grows without bound — old units approach a constant (worst) failure rate, not an ever-increasing one.

ESTIMATION METHOD COMPARISON — FLEET FIT

1 MLE is the reported model; the RR line is the probability plot's own fit at the MLE shape — shown for cross-checking against rank-regression-based software. A large intercept or slope gap signals model misfit.

2 Method	3 θ (scale)	4 Intercept	5 r
6 MLE (reported)	43.8810	0 (by construction)	–
7 Rank regression (plot line)	48.9558	-13.71	0.9659

Estimation method comparison — fleet fit

- ① MLE is the reported model; the RR line is the probability plot's own fit at the MLE shape — shown for cross-checking against rank-regression-based software. A large intercept or slope gap signals model misfit.
- ② to ⑤ The columns Method, θ (scale), Intercept, r.
- ⑥ MLE (reported): θ 43.8810, intercept 0 (by construction).
- ⑦ Rank regression (plot line): θ 48.9558, intercept -13.71, r 0.9659.

PROBABILITY-PLOT FIT CHECK (PPCC, 5% LEVEL) — FLEET FIT

1 Straightness NOT rejected (r = 0.9659 ≥ critical 0.9482).

2 Seeded simulation at the FITTED shape (treated as known — an approximation, stated rather than hidden).

Probability-plot fit check (PPCC, 5% level) — fleet fit

- ① Straightness NOT rejected (r = 0.9659 ≥ critical 0.9482). The gamma model is not rejected by the plot.
- ② Seeded simulation at the FITTED shape (treated as known — an approximation, stated rather than hidden).

COMPETITOR MODELS — GAMMA VS WEIBULL VS LOGNORMAL (SAME DATA)

1 Lognormal(μ = 4.737, σ = 0.597) explains this data best by AIC.

2 Model	3 ln-likelihood	4 AIC	
5 Gamma (this analysis)	-141.168	286.3	
6 Weibull(β = 1.782, η = 153.014)	-142.052	288.1	
7 Lognormal(μ = 4.737, σ = 0.597)	-141.014	286.0	◄ best

8 With small samples the three are often nearly indistinguishable by fit statistics — the MECHANISM decides: damage accumulation/standby → Gamma; weakest link of many flaws → Weibull; multiplicative growth (fatigue crack, repair times) → Lognormal.

Competitor models — Gamma vs Weibull vs Lognormal (same data)

- ① Lognormal(μ = 4.737, σ = 0.597) explains this data best by AIC.
- ② to ④ The columns Model, ln-likelihood, AIC.
- ⑤ Gamma (this analysis): -141.168, AIC 286.3.
- ⑥ Weibull(β = 1.782, η = 153.014): -142.052, AIC 288.1.
- ⑦ Lognormal(μ = 4.737, σ = 0.597): -141.014, AIC 286.0, ◄ best.
- ⑧ With small samples the three are often nearly indistinguishable by fit statistics — the MECHANISM decides: damage accumulation/standby → Gamma; weakest link of many flaws → Weibull; multiplicative growth (fatigue crack, repair times) → Lognormal.

THE SHOCK STORY — WHICH FIX PAYS MORE? (FLEET FIT)

1 The fit reads as: a unit survives ≈ 3.09 damage events, one about every 43.88 time units. Two different engineering levers move B10: raising the damage TOLERANCE (material, design margin → k) and slowing the damage ARRIVAL (operating stress, environment → θ).

2 Scenario	3 Parameters	4 B10
5 Today	k = 3.09, θ = 43.88	50.7
6 Tolerate ONE more hit (design margin)	k + 1 = 4.09	79.1
7 Slow the shock rate 10%	θ × 1.1	55.7
8 Slow the shock rate 20%	θ × 1.2	60.8
9 Slow the shock rate 30%	θ × 1.3	65.9

10 B10 scales exactly with θ; the k-lever's value depends on where k already is (diminishing per-hit gains as k grows).

The shock story — which fix pays more? (fleet fit)

- ① The fit reads as: a unit survives ≈ 3.09 damage events, one about every 43.88 time units. Two different engineering levers move B10: raising the damage TOLERANCE (material, design margin → k) and slowing the damage ARRIVAL (operating stress, environment → θ).
- ② to ④ The columns Scenario, Parameters, B10.
- ⑤ Today: k = 3.09, θ = 43.88, B10 50.7.
- ⑥ Tolerate ONE more hit (design margin): k + 1 = 4.09, B10 79.1.
- ⑦ Slow the shock rate 10%: θ × 1.1, B10 55.7.
- ⑧ Slow the shock rate 20%: θ × 1.2, B10 60.8.
- ⑨ Slow the shock rate 30%: θ × 1.3, B10 65.9.
- ⑩ B10 scales exactly with θ; the k-lever's value depends on where k already is (diminishing per-hit gains as k grows).

VALUE OF MORE DATA — DEFENSIBLE B10 (FLEET FIT)

1 Point estimate today: 50.7. The bound narrows with √n — prices the next test campaign before it is bought.

Test campaign	Defensible B10 (lower bound)
today	33.6
n = 30	34.8
n = 35	35.8
n = 40	36.6
n = 50	37.9

2 3 4 5 6 7

7 The projection assumes the estimates hold.

Value of more data — defensible B10 (fleet fit)

- ① Point estimate today: 50.7. The bound narrows with √n — prices the next test campaign before it is bought.
- ② today: the lower B10 bound 33.6.
- ③ n = 30: 34.8.
- ④ n = 35: 35.8.
- ⑤ n = 40: 36.6.
- ⑥ n = 50: 37.9.
- ⑦ The projection assumes the estimates hold.

STANDBY DESIGN — SPARES NEEDED FOR THE MISSION (FROM k̂ AND θ̂)

1 n* = 1 cold spares (+ the operating unit — system life Gamma(2·k̂, θ̂)) reach R(mission 100) ≥ 90%.

n — cold spares	R_sys(mission)	Marginal gain
0	0.6222	+0.6222
1	0.9762	+0.3539
2	0.9996	+0.0234
3	1.0000	+0.0004

2 3 4 5 6 7 8

9 Cold standby: ONE unit operating + n spares, each with the fitted life Gamma(k̂, θ̂) → system life Gamma((n + 1)·k̂, θ̂) EXACTLY (independent gamma lives with one scale add their shapes); R = Q((n + 1)·k̂, t/θ̂). Closed form — no simulation.

Standby design — spares needed for the mission (from k̂ and θ̂)

- ① n* = 1 cold spares (+ the operating unit — system life Gamma(2·k̂, θ̂)) reach R(mission 100) ≥ 90%. One spare is enough.
- ② n — cold spares: the spares beside the operating cartridge.
- ③ R_sys(mission): the reliability of the group over the 100-hour mission.
- ④ Marginal gain: what each added spare buys.
- ⑤ No spare: 0.6222 - the running cartridge alone.
- ⑥ One spare: 0.9762, a gain of +0.3539 - the first row at or above 90%.
- ⑦ Two spares: 0.9996, a gain of only +0.0234.
- ⑧ Three spares: 1.0000, +0.0004.
- ⑨ Cold standby: ONE unit operating + n spares, each with the fitted life Gamma(k̂, θ̂) → system life Gamma((n + 1)·k̂, θ̂) EXACTLY (independent gamma lives with one scale add their shapes); R = Q((n + 1)·k̂, t/θ̂). Closed form — no simulation.

SERVICE INTERVALS — OPERATING TIME TO STAY AT OR ABOVE TARGET RELIABILITY

1 Model: Gamma(k = 3.085, θ = 43.881) — service before the unit accumulates this much operating time to keep reliability at or above the target.

2 Target R	3 Operating time t
95%	37.8
4 90%	50.7
85%	61.0
80%	70.1
75%	78.8
70%	87.1
65%	95.4
60%	103.7
55%	112.2
50%	121.1
45%	130.4
40%	140.3
35%	151.1
30%	163.0
25%	176.5
20%	192.5
15%	212.2
10%	238.7
5%	281.9

SERVICE INTERVALS — operating time to stay at or above target reliability

- ① Model: Gamma(k = 3.085, θ = 43.881) — service before the unit accumulates this much operating time to keep reliability at or above the target.
- ② Target R: the reliability to keep, from 95% down to 5%.
- ③ Operating time t: the longest run that keeps it.
- ④ The 90% row: 50.7 hours - B10.

PERCENTILE TABLE — B-LIVES WITH BOUNDS

1 B_p = the age by which p% of units fail. Bounds by each family's own interval machinery; a — cell means no bound machinery exists for that family — the point estimate stands alone.

2 Percentile	3 B-life	4 95% bounds
1%	20.44	[9.778, 42.73]
5%	37.82	[22.81, 62.69]
5 10%	50.67	[33.61, 76.38]
25%	78.77	[58.65, 105.8]
50%	121.1	[96.08, 152.6]
75%	176.5	[140.7, 221.5]
90%	238.7	[185.8, 306.7]
95%	281.9	[215.4, 368.9]
99%	375.2	[276.7, 508.8]

Percentile table — B-lives with bounds

- ① B_p = the age by which p% of units fail. Bounds by each family's own interval machinery; a — cell means no bound machinery exists for that family — the point estimate stands alone.
- ② to ④ The columns Percentile, B-life, 95% bounds.
- ⑤ The 10% row: 50.67 [33.61, 76.38] - the B10 of the results card.

1 Licensed to ReliaEn Training · RL-2026-0017

The report's foot

- ① The licence line Licensed to ReliaEn Training · RL-2026-0017: the licence the report was made under.

Part 5 — How to interpret the results

What the fit says. A cartridge survives about three particle loads ($\hat{k}$ = 3.08544), one about every 43.881 hours ($\hat{\theta}$). Even the lower bound on k, 1.82136, is above 1, so the cartridges wear out. The mean life 135.39 hours is exactly the average of the 25 ages; one cartridge in ten has clogged by 50.67 hours (B10), and the defensible figure for a commitment is its lower bound, 33.61.

Which fix pays more. Tolerating one more load lifts B10 from 50.7 to 79.1 hours; slowing the shock rate 30% ($\theta \times 1.3$) gives only 65.9. A tougher cartridge pays more than a gentler duty here. A larger test would move the defensible B10 only slowly - from 33.6 today to 37.9 at 50 cartridges.

Which family. The lognormal edges the gamma by AIC, 286.0 against 286.3. A gap that small proves nothing: choose the family by the mechanism, which here is load accumulation - the gamma.

The spares for the mission. One cold spare beside the running cartridge reaches the 100-hour mission at 90% (R = 0.9762); with none the mission reaches only 0.6222. Read the marginal gain before buying more: a second spare adds only +0.0234.