Cox regression — the hazard ratio of each condition, with no life distribution assumed

Menu: Life Data ▸ Cox Regression…

Part 1 — What the function does

What it calculates. The question is whether a condition of a unit - its supplier, site, design version or load - raises or lowers its risk of failure, without assuming a life distribution. From one row per unit - its time to failure or its running time without failure, a censor flag, and up to three covariates (numbers, or labels such as site names) - ReliaEn fits the proportional hazards model by maximising the partial likelihood: at each failure time it compares the unit that failed with every unit still at risk. For each covariate it gives a hazard ratio: how many times the risk of failure at any moment is multiplied per unit of that covariate, or for one level against a reference level.

What ReliaEn prints	What it tells you
Partial log-likelihood with the null value and the global LR	How much the covariates add to a model without them
Hazard ratios: β, SE, HR with its interval, LR / p, verdict, PH check	Each condition's effect on the risk, whether the data support it, and whether the ratio stays steady over time
Residual diagnostics	Any unit the model fits poorly
A survival plot and a Cox–Snell residual plot	The survival at the covariate mean, and the model's overall fit

From which data. One row per unit: the time, a censor flag (0 = failed, 1 = still running), and one to three covariate columns, each continuous or categorical.

When an engineer uses it. When each unit has one life, its conditions are fixed for that life, and relative risks are wanted without choosing a distribution.

When not to use it.

Situation	Use instead
A covariate changes while the unit runs, units enter late, or one unit fails several times	Life Data ▸ Cox Regression (Counting Process)… (chapter 2.24)
A predicted life (B10, MTTF) at a new condition is needed	Life Data ▸ Regression with Life Data… (chapter 2.21)
Only the survival curves of groups are to be compared	Life Data ▸ Survival Plots… (chapter 2.15)

Part 2 — The example's input data

Setting. Seven cooling fans come from two suppliers, coded 0 for supplier A and 1 for supplier B. Six failed; one fan of supplier B was still running at 8 hours when the record was taken. The reliability engineer wants to know whether supplier B's fans fail sooner.

Fan	1	2	3	4	5	6	7
Time, hours	3	5	7	8	10	12	15
Flag (1 = still running)	0	0	0	1	0	0	0
Supplier (1 = B)	1	0	1	1	0	1	0

Input	Value
X1 type / Ties method / Confidence %	Continuous / Breslow ties / 95
X1* (comparison point)	1 - a supplier B fan
File in the worked-example zip	2-Cox-Regression/Example-1.csv — columns Time, Flag, X1 (and X1 label, the same supplier as A and B), 7 rows

	① A	② B	③ C	D
1	3	0	1	
2	5	0	0	
3	7	0	1	
4	8	1	1	
5	10	0	0	
6	12	0	1	
7	15	0	0	
8				

The example data in the worksheet

- ① Column A: the 7 times.
- ② Column B: the censor flags.
- ③ Column C: the supplier code.

Part 3 — How to use it, step by step

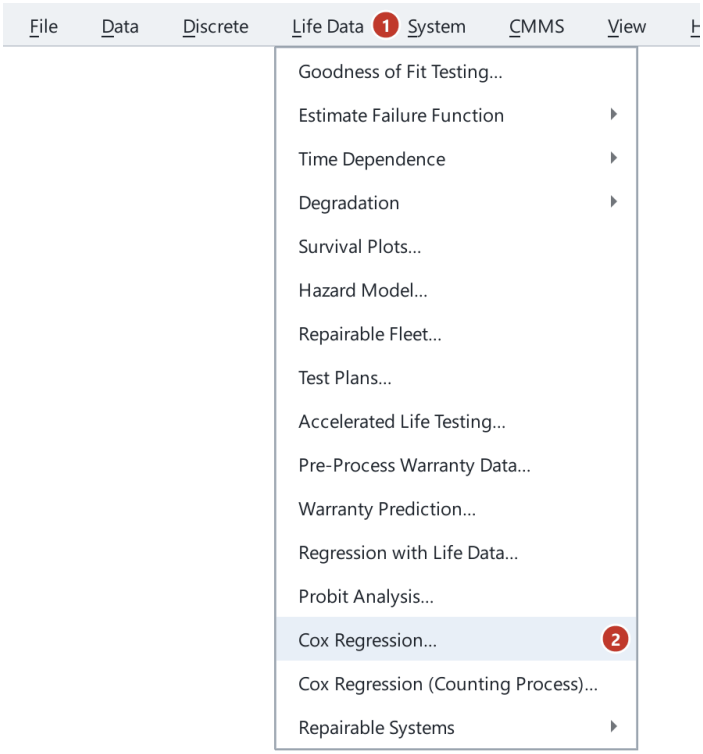
1. **Put the fan record in columns A to C.** Copy the 7 rows of Example-1.csv - its columns Time, Flag and X1 - into columns A, B and C of the worksheet, starting at row 1. The X1 label column stays out.

	① A	② B	③ C	D
1	3	0	1	
2	5	0	0	
3	7	0	1	
4	8	1	1	
5	10	0	0	
6	12	0	1	
7	15	0	0	
8				

Step 1 — the fan record in columns A to C

- ① Column A: the 7 times, 3 to 15 hours.
- ② Column B: the censor flags; the 1 is the fan still running at 8 hours.
- ③ Column C: the supplier code, 0 for A and 1 for B.

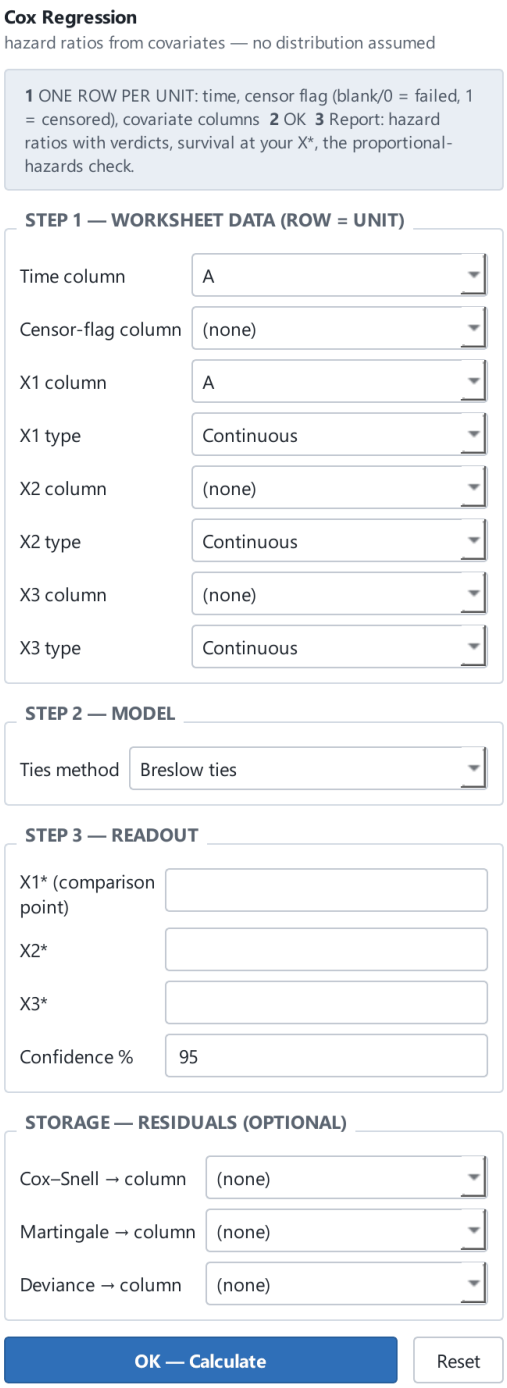
2. **Choose Life Data ▸ Cox Regression...** in the menu bar. The form Cox Regression opens in the left panel.



Step 2 — the menu path

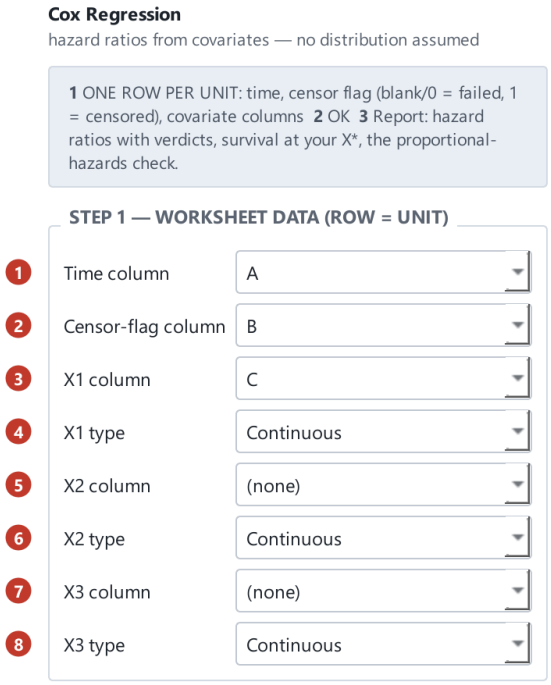
- ① Life Data: the menu.
- ② Cox Regression...: the entry to click.

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

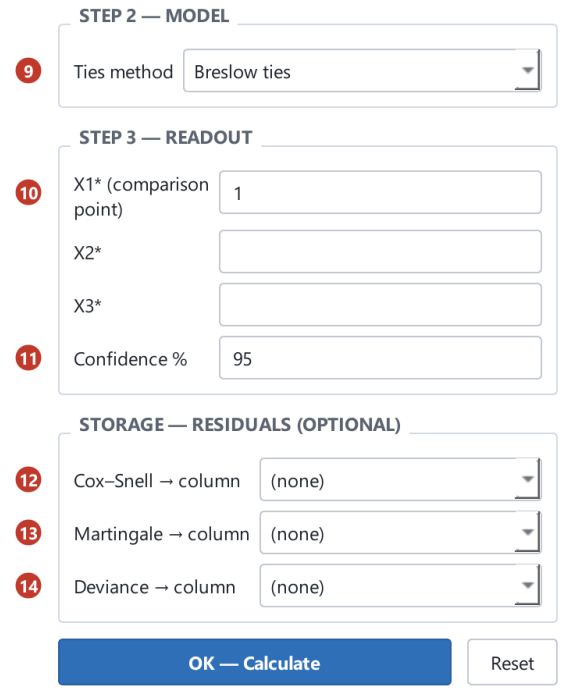


Step 2 — the form as it opens

Steps 3 to 16 set or check the form's fourteen boxes, shown in two pictures - ① to ⑧, then ⑨ to ⑭:



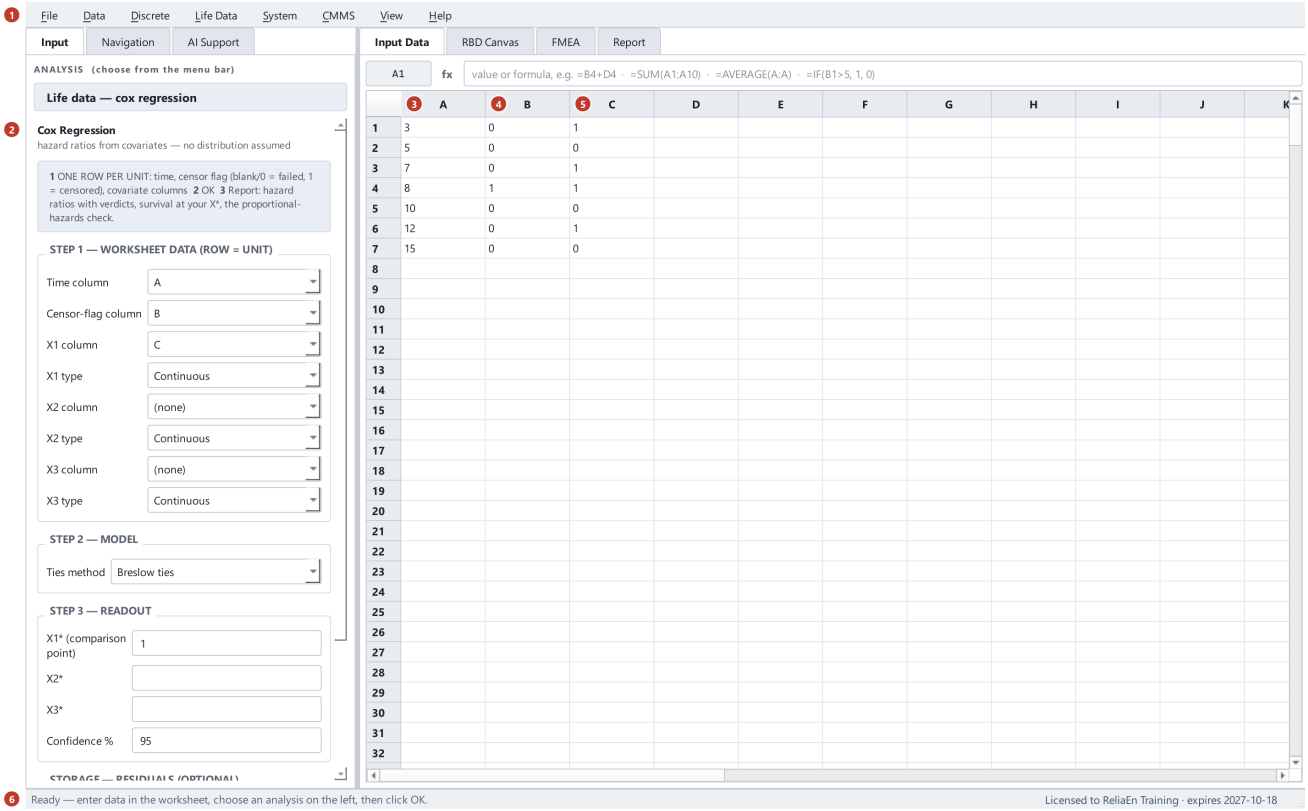
Steps 3 to 10 — the worksheet boxes set for this example



Steps 11 to 16 — the model, readout and storage boxes

3. ① **Check that Time column reads A.**
4. ② **In Censor-flag column, choose B.** Without it every row would count as a failure.
5. ③ **In X1 column, choose C.** The box opens on A.
6. ④ **Check that X1 type reads Continuous.** A 0/1 code is read as a number.
7. ⑤ **Check that X2 column reads (none).** One covariate only.
8. ⑥ **Check that X2 type reads Continuous.**
9. ⑦ **Check that X3 column reads (none).**
10. ⑧ **Check that X3 type reads Continuous.**
11. ⑨ **Check that Ties method reads Breslow ties.**
12. ⑩ **In X1* (comparison point), type 1.** The survival plot's second curve is a supplier B fan.
13. ⑪ **Check that Confidence % reads 95.**
14. ⑫ **Check that Cox–Snell → column reads (none).**
15. ⑬ **Check that Martingale → column reads (none).**
16. ⑭ **Check that Deviance → column reads (none).** These three boxes write residuals into worksheet columns; this example writes none.

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



Step 17 — 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 16; OK — Calculate at its foot.
- ③ Column A: the times.
- ④ Column B: the censor flags.
- ⑤ Column C: the supplier code.
- ⑥ 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 — Cox Regression

2026-11-04 10:42:13

partial likelihood over risk sets ; S(t|x) = S_0(t)^exp(β^T(x−$\bar{x}$))

The report's head

- The title Life Data — Cox Regression, the date and time of the run, and the model in one line: partial likelihood over risk sets; S(t|x) = S_0(t)^exp(β^T(x−$\bar{x}$)).

INPUT DATA

1	Data	this worksheet — 6 events / 1 censored, 1 covariate(s)
2	Ties method	Breslow
3	Model	semi-parametric proportional hazards — no life distribution assumed

Input data

- ① Data this worksheet — 6 events / 1 censored, 1 covariate(s).
- ② Ties method Breslow.
- ③ Model semi-parametric proportional hazards — no life distribution assumed.

RESULTS

1	Partial log-likelihood	-7.0365 (null -7.1389, global LR = 0.205)

Results

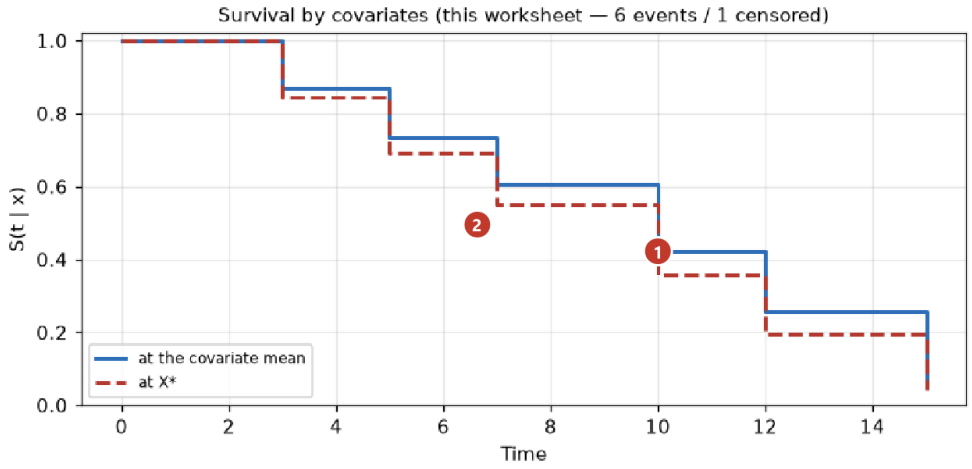
- ① Partial log-likelihood -7.0365 (null -7.1389, global LR = 0.205): the fit with the supplier, the fit without it, and the test between them.

INTERPRETATION

1. Each hazard ratio multiplies the CURRENT risk per unit of its covariate — no claim about the life distribution's shape (that is the Cox trade: no extrapolation beyond the data, no shape assumption). Believe a ratio only if its verdict column and the proportional-hazards check both pass.

Interpretation

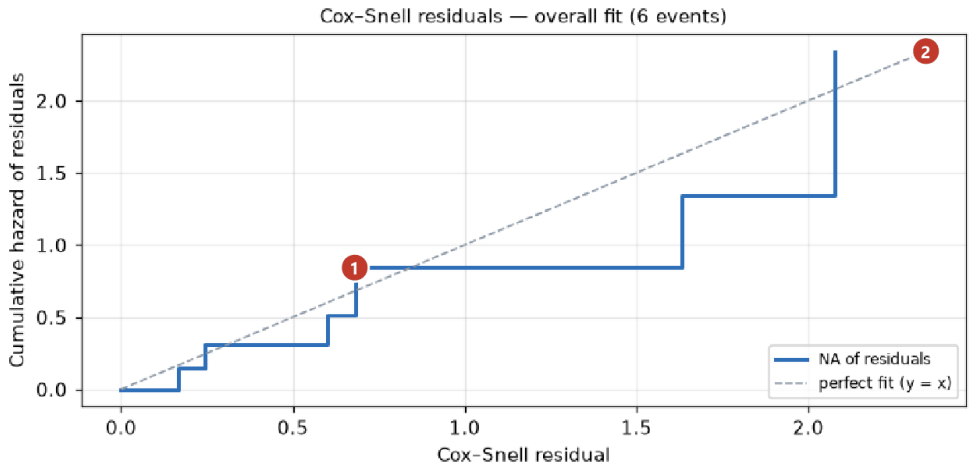
- ① ReliaEn's reading: Each hazard ratio multiplies the CURRENT risk per unit of its covariate — no claim about the life distribution's shape (that is the Cox trade: no extrapolation beyond the data, no shape assumption). Believe a ratio only if its verdict column and the proportional-hazards check both pass.



Interactive Plot

Survival by covariates (this worksheet — 6 events / 1 censored)

- ① at the covariate mean: S(t | x) for the average fan, stepping down at each failure.
- ② at X*: the same steps for a supplier B fan (X1 = 1), a little lower at every time.



Interactive Plot

Cox–Snell residuals — overall fit (6 events)

- ① NA of residuals: the cumulative hazard of the residuals.
- ② perfect fit (y = x): the line the residuals follow when the model fits.

HAZARD RATIOS (THIS WORKSHEET — 6 EVENTS / 1 CENSORED, 1 COVARIATE(S))

1 HR = e^β per covariate unit. The PH column is the assumption the whole model rides on: a drifting ratio means one number cannot summarize that covariate — stratify or split the follow-up.

2 Term	3 β	4 SE	5 Hazard ratio [CI]	6 LR / p	7 Verdict	8 PH check
9 X1	0.4162	0.9294	1.5162 [0.2453, 9.3721]	0.205 / 0.6510	NOT significant	stable

Hazard ratios (this worksheet — 6 events / 1 censored, 1 covariate(s))

- ① HR = e^β per covariate unit. The PH column is the assumption the whole model rides on: a drifting ratio means one number cannot summarize that covariate — stratify or split the follow-up.
- ② to ⑧ The columns Term, β, SE, Hazard ratio [CI], LR / p, Verdict, PH check.
- ⑨ The row for X1: β 0.4162, SE 0.9294, hazard ratio 1.5162 [0.2453, 9.3721], LR 0.205 / p 0.6510, NOT significant, stable.

RESIDUAL DIAGNOSTICS

1 The Cox–Snell plot carries the overall check: residuals on the unit line = the PH model fits.

2 No units with |deviance residual| > 2.5 — no outlier signal.

Residual diagnostics

- ① The Cox–Snell plot carries the overall check: residuals on the unit line = the PH model fits.
- ② No units with |deviance residual| > 2.5 — no outlier signal.

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 this report was made under.

Part 5 — How to interpret the results

Does the supplier matter? Supplier B's fans run at a hazard 1.5162 times supplier A's (β = 0.4162, SE 0.9294). But the 95 % interval runs from 0.2453 to 9.3721: it contains 1, so the data cannot tell more risk from less. The

likelihood-ratio test agrees - the partial log-likelihood rises only from −7.1389 without the supplier to −7.0365 with it, LR = 0.205, p = 0.6510 - and the verdict is NOT significant.

What to decide. Do not choose between the suppliers on these seven fans; collect more failures first. The PH check reads stable, so the single ratio would be a fair summary - once the data support it.

Why the interval is so wide. Six failures give the partial likelihood very little curvature: the standard error, 0.9294, is more than twice β itself. Each further failure narrows it.

A label column gives the same answer. The file's X1 label column holds the same suppliers as the text A and B. Read as Categorical, it gives the term X1 [B vs A] with exactly the same β, ratio and test: a 0/1 code and a two-level label are the same model.