

Monte Carlo simulation: should we **build Plant 2?**

The single-number spreadsheet says a second plant is worth 45 million dollars. Simulating the same assumptions 10,000 times gives an expected value of 17 million, with a 29% chance of losing money. It also finds a better plan: wait a year, and build only if demand shows up.

Functions `fx.RiskGBMA` · `fx.RiskOUA` · `fx.RiskFitGBMA` · `fx.RiskFitOUA` · `fx.RiskPertA` · `fx.RiskBernoulliA` · `fx.RiskSPTMetaLogA` · `fx.RiskCopulaGaussA` · `fx.RiskVariablesTableA` · `fx.RiskMeasuresA` · chart data

In XL Edge XL Edge tab > Productivity Tools > Stat Tools > the Monte Carlo and Time Series menus

Example *XL Edge Example - Monte Carlo.xlsx* (same folder). Every number in this guide is read from it.

Needs Excel for Microsoft 365. The functions also work without XL Edge once they are in a workbook.

1 The decision

2 Why one spreadsheet gets it wrong

3 Using the ribbon

4 From history to assumptions

5 The assumptions

6 The calculation engine

7 Results: build now or wait

8 Choosing the wait rule

9 What drives the answer

10 Checking the model

11 Practical advice

12 Function reference

1 The decision

The manufacturer from the Copula guide sells everything Plant 1 can make: 430,000 units a year, and market demand has just reached it. Plant 2 would add 200,000 units of capacity for a capital cost of about 55 million dollars, take 2 years to build, and cost 4 million a year to run. Every extra unit sold earns the price (200 dollars) less conversion (45) and material (about 72), so roughly 83 dollars. Two plans are on the table:

- **Build now.** Start in 2027; Plant 2 runs from 2029 (2030 if construction runs late).

- **Wait a year.** Watch 2027's demand. Build in 2028 only if it reaches 450,000 units; Plant 2 then runs from 2030 (or 2031).

Both are valued the same way: ten years of after-tax cash flows from the extra units, discounted at a 9.0% cost of capital, plus a terminal value for the years after 2036.

2 Why one spreadsheet gets it wrong

Put the expected demand, the expected material cost, the average capital cost, no delay and the median ramp-up into one spreadsheet, and Plant 2's NPV is **45.3 million**. Sheet *Outputs* builds exactly that spreadsheet, year by year. It is not wrong arithmetic. It is the wrong question: it values the plant *at* the average future, not *over* the futures that might happen. The two differ whenever the payoff bends, and capacity makes it bend twice:

- **The upside is capped.** Once demand passes both plants' capacity, extra demand earns nothing. Good years are worth less than the average suggests.
- **The downside is not.** If demand stalls near Plant 1's capacity, Plant 2 sits idle, but the capital is spent and the fixed costs run anyway.

The simulated mean NPV of building now is **17.2 million**, about 62% below the spreadsheet. Sam Savage calls this the *flaw of averages*: plans based on average conditions are wrong on average.

A one-year example

Plant 2 sells only what Plant 1 cannot: demand above 430,000 units, up to 200,000 more. Suppose next year's demand is equally likely to be 530,000 or 730,000 units, so the average is 630,000:

Extra units Plant 2 sells

	Demand	Extra units for Plant 2
The plan: use the average demand	630,000	200,000 (full)
A low year	530,000	100,000
A high year	730,000	200,000 (capped: the last 100,000 earn nothing)
The average of the two years	630,000	150,000


The average demand is the same, but Plant 2 sells 25% fewer units than the plan says. The bad year costs the full amount; the good year cannot make it up. That asymmetry, repeated over ten years and thousands of futures, is the whole gap.

THE FLAW OF AVERAGES IN ONE SENTENCE

Because the plan assumes Plant 2 will be full, every shortfall in demand costs real sales while every surplus earns nothing: **the downside is counted in full and the upside is capped**. (There is a floor too: below Plant 1's capacity, Plant 2 sits idle and its sales cannot fall further. But by then the capital and the fixed costs are already committed.)

Where the 28.1 million goes

To find out which uncertainty causes the gap, let one input vary across the 10,000 trials while every other input stays at its plan value, then the next, and so on. Sheet *Outputs*, "From the Plan to the Simulated Mean", does this with one formula per row.



One bar does the work. Each step is one uncertainty on its own; the overlap is what is left when they all vary together.

Build now: from the plan to the simulated mean, \$ millions (sheet *Outputs*)

Step	Change	Why
Single-point plan	45.3	Every input at its expected value
Capital cost varies	0.0	A straight line: overruns and savings cancel out
Material cost varies	-0.1	A straight line: dear years and cheap years cancel out
First-year ramp-up varies	0.0	One year only, and roughly symmetric
Construction delay	-0.5	Not the flaw of averages: the plan simply assumed no delay
Demand varies, against fixed capacity	-27.4	The bend: Plant 2 is idle below 430,000 and full above 630,000
Risks overlapping	-0.1	For example, overruns arrive with delays (the copula)
Simulated mean	17.2	

Averaging is safe when the payoff is a straight line. NPV moves dollar for dollar with capital cost and material cost, so the high and low trials cancel and the plan's average gets the right answer. Demand is the only input that meets a bend, and it accounts for 27.4 of the 28.1 million. Two things make it so large:

- **The plan sits on the ceiling.** By 2034 the expected demand is above both plants' capacity, so the plan fills Plant 2 every year from then on. In the simulation, 41% of futures are still below that line in 2036 (see the year-by-year table at the end of this section).

- **The terminal value multiplies the miss.** It values every year after 2036 on 2036's cash flow. In the plan it is worth 61.3 million today, more than the whole NPV. Uncertain demand takes 18.7 million off it, and 8.7 million off the first ten years.

WHY THE GAP IS THIS LARGE HERE

Plant 2 is an illustrative case with deliberately visible uncertainty: demand volatility of 10% a year, compounding for ten years, against a plan that fills the new plant exactly. The gap grows with that uncertainty. With half the demand volatility (5%, the same random draws), demand's step in the bridge shrinks from 27.4 to 9.2 million. A steadier market would show a smaller gap and a more volatile one a larger gap, but the direction stays the same: whenever capacity caps the upside, the single-point plan overstates the value.

Build now: three futures, by demand in 2036 (sheet *Outputs*, "Three Futures")

Demand in 2036	Share of trials	Mean NPV (\$ millions)	Chance of a loss
Below Plant 1's capacity: Plant 2 idle	8%	-75.1	100.0%
Between: Plant 2 partly used	33%	-12.6	62.6%
Above both plants: Plant 2 full	59%	46.7	0.2%
All trials	100%	17.2	29.0%

The "full" row is worth about what the plan says. **The plan is not the average future; it is the future in which things go right.** Three numbers tell a decision-maker the rest: the median NPV is 33.0 million, the mean is 17.2 million (a few deep losses pull it below the median), and only 34% of the trials reach the plan's 45.3 million.

Build now: the plan against the simulation, extra units sold (sheet *Outputs*, "Year by Year")

Year	2029	2031	2033	2036
Plan (expected demand)	69,589	122,131	180,199	200,000
Simulated mean	50,017	107,664	130,640	153,686
Share of trials with demand above both plants	8%	24%	40%	59%

3 Using the ribbon

XL Edge tab > Productivity Tools > Stat Tools (arrow)

Piece of the model	Menu	Item
Fit demand and material cost to history	Insert Time Series	Fit geometric Brownian motion; Fit mean reversion (Ornstein-Uhlenbeck)
Ten-year demand and material paths	Insert Time Series	Geometric Brownian motion; Mean reversion (Ornstein-Uhlenbeck)
Capital cost, delay, ramp-up	Insert Monte Carlo Distribution	PERT; Bernoulli; SPT metalog
Link overruns to delays	Insert Monte Carlo Copula	Gaussian (Normal)
Summary tables	Insert Monte Carlo Statistics	Insert Monte Carlo Variables Table
VaR, CVaR, ES	Insert Monte Carlo Statistics	Insert Monte Carlo Risk Measures
Charts	Insert Monte Carlo Chart	Fan Chart (Time Series); Outcome Histogram (P10 / P50 / P90); Tornado (Sensitivity)

Each item asks a few questions, shows the size of the result, and writes one live formula. The income statement and cash flow arithmetic between them is ordinary Excel, written once per column, and it spills down all the trials.

4 From history to assumptions

Sheet *Data* holds fifteen years of market demand and sixty months of material cost (illustrative). Two fitting functions turn them into time-series parameters:

```

demand: growth and its volatility, a year
=fx.RiskFitGBMA(D10:D24)
material cost: monthly data (dt = 1/12), parameters a year
=fx.RiskFitOUA(G10:G69, 1/12)

```

Demand: geometric Brownian motion

Parameter	Fit	Used
Drift (growth a year)	5.1%	5.0%
Volatility	10.2%	10.0%

Material cost: Ornstein-Uhlenbeck

Parameter	Fit	Used
Long-run mean (\$)	71.37	72.00
Speed of reversion	0.66	0.60
Volatility (\$)	9.75	10.00

Two models, two kinds of uncertainty. Demand follows *geometric Brownian motion*: it grows by a random percentage each year, and the shocks compound, so the range of possible demand keeps widening. Material cost follows *Ornstein-Uhlenbeck* mean reversion: high prices draw new supply and low prices

cut it, so cost is pulled back toward its long-run mean. A gap halves in about 1.2 years, so the range stops widening. The fan charts in section 7 show both shapes.

FIT, THEN JUDGE

Fifteen yearly points pin the drift down only roughly: its standard error is about volatility $\div \sqrt{14}$, near three percentage points. Mean-reversion speed is harder still. The fits are a starting point, which is why the yellow cells on *Assumptions* hold rounded planning values, and the fits sit beside them in gray for comparison.

5 The assumptions

Sheet *Assumptions* (yellow cells are the inputs)

Input	Model	Values	Why
Market demand	GBM	starts at 430,000; drift 5.0%, volatility 10.0%	Fitted to fifteen years, rounded.
Material cost (\$ per unit)	Ornstein-Uhlenbeck	starts at 69.87; mean 72.00, speed 0.60, volatility 10.00	Fitted to sixty months.
Capital cost	PERT	50 / 55 / 75 million (mean 57.5)	The engineers' three-point estimate; overruns run long.
Construction delay	Bernoulli	25% chance of one extra year	Permits and equipment lead times.
Overruns and delays	Gaussian copula	$\rho = 0.60$	Projects that run late usually run over budget too.
First-year ramp-up	SPT metalog	P10 55%, P50 70%, P90 85% of capacity	The plant manager's low / likely / high.

Fixed inputs: price 200 and conversion cost 45 dollars a unit; Plant 1 430,000 and Plant 2 200,000 units a year; Plant 2's fixed cost 4,000,000 dollars a year; tax 21%; WACC 9.0%; terminal growth 2.0%; straight-line depreciation over 10 years. Tax is charged on the plant's own profit, and an early loss saves tax, because the company's other profits absorb it.

6 The calculation engine

Sheet *Simulation_Calculations* has one row per trial. The first nine columns hold each trial's project risks. Then come nine ten-year blocks (demand, material cost, unit margin, then capacity, extra units and cash flow for each strategy), and finally the NPVs. Every block is one spilled formula. Named inputs (*Iterations*, *NYears*, *Plant1Cap*, *WACC* and so on) keep the formulas readable.

```

R9: demand, 10,000 trials x 10 years, in one spill
=fx.RiskGBMA(DemandStart, DemandDrift, DemandVol, NYears, , Iterations, 21, SimSeed)
AC9: material cost
=fx.RiskOUA(MatStart, MatMean, MatSpeed, MatVol, NYears, , , Iterations, 22, SimSeed)
H9, J9, K9: correlated uniforms, then capital cost and delay from them
=fx.RiskCopulaGaussA(CapexDelayMatrix, Iterations, 30, SimSeed)
=fx.RiskPertA(Assumptions!$I$14, Assumptions!$I$15, Assumptions!$I$16, CHOOSECOLS(H9#, 1))
=fx.RiskBernoulliA(DelayProb, CHOOSECOLS(H9#, 2))
AY9: Plant 2's capacity each year (M9# = the year it starts; L9# = first-year ramp-up)
=LET(yr, SEQUENCE(1, NYears), Plant2Cap * (yr ≥ M9#) * IF(yr = M9#, L9#, 1))
BJ9: the extra units Plant 2 sells: demand above Plant 1, up to Plant 2's capacity
=IF(R9# < Plant1Cap + AY9#, R9#, Plant1Cap + AY9#) - IF(R9# < Plant1Cap, R9#, Plant1Cap)

```

A column of trials meets a row of years and Excel broadcasts the pair into a full block: `SEQUENCE(1, NYears)` is one row, `M9#` one column, and `yr ≥ M9#` is $10,000 \times 10$. The capacity line, `BJ9`, is where the flow of averages lives: `MIN(demand, capacity)`, written with `IF` so it works cell by cell.

Cash flow is after-tax operating profit, plus depreciation, less the capital spent. NPV discounts the ten years with `MMULT` and adds a terminal value: the last year's after-tax operating cash flow, growing at 2.0% for ever.


```

DM9: NPV, one value per trial
=LET(df, (1 + WACC) ^ (-SEQUENCE(NYears)),
    final, (TAKE(BJ9#, , -1) * TAKE(AN9#, , -1) - Plant2Fixed) * (1 - TaxRate),
    tv, final * (1 + TermGrowth) / (WACC - TermGrowth),
    MMULT(BU9#, df) + tv * INDEX(df, NYears))


```

The wait strategy repeats the blocks with the plant a year later, and the rule is applied at the end: `NPV wait = (2027 demand ≥ threshold) × NPV if built`. Both strategies use the **same** demand, cost and risk draws, so their difference is the decision's effect, not luck.

7 Results: build now or wait



Demand widens: P10 to P90 band, median, the expected path (Plan, dashed) and the two capacity lines.




Material cost reverts: the band stops widening after a year or two.


NPV, \$ millions (sheet *Outputs*)

	Build now	Wait a year
Single-point plan (one spreadsheet)	45.3	47.4
Mean	17.2	14.5
Median	33.0	0.0
P5 / P95	-75.5 / 68.0	-33.1 / 63.5
Probability NPV is below zero	29.0%	9.2%
Value at Risk, 95% (loss not exceeded in 19 trials of 20)	75.5	33.1
CVaR, 95% (average loss in the worst 5%)	79.4	57.0
Probability Plant 2 is built	100.0%	49.4%
Mean NPV when it is built	17.2	29.4


Against the model's 450,000 rule, **building now has the higher mean**, by 2.7 million, and **waiting has far less risk** (section 8 finds a rule that wins on both): the chance of losing money falls from 29.0% to 9.2%, and the 95% VaR from 75 to 33 million. Waiting works like an option: in the 51% of futures where 2027 demand disappoints, the company simply doesn't build (NPV zero), and when it does build, the plant is worth 29.4 million on average because it is built into stronger demand.



Build now: P5, P50 and P95 lines. A fat left tail: idle-plant futures.



Wait a year: the spike at zero is the futures where the plant is never built.



The two S-curves. At any NPV, read up to the chance of doing worse. Waiting (black) sits lower on the left (less chance of a loss) and slightly higher on the right (less chance of a big win).

8 Choosing the wait rule

"Build if 2027 demand reaches 450,000" is one rule among many. Because the model keeps each trial's NPV *if built*, any threshold can be tested on the same 10,000 futures with one formula per cell, for example `=AVERAGE((N9# ≥ threshold) * DN9#)` for its mean NPV:

Sheet *Outputs*, "The Wait Rule": \$ millions

Rule	Built	Mean NPV	P(NPV < 0)	P5 NPV
Build now	100%	17.2	29.0%	-75.5
Wait, then build regardless	100%	18.7	27.8%	-70.3
Wait, build if demand ≥ 420,000	75%	18.7	16.7%	-57.5
... if ≥ 440,000	58%	16.1	11.7%	-43.3
... if ≥ 450,000 (the model's rule)	49%	14.5	9.2%	-33.1
... if ≥ 470,000	33%	10.8	5.2%	-1.6
... if ≥ 500,000	14%	5.4	1.6%	0.0

Two findings. First, even *waiting a year and building regardless* beats building now on the mean (18.7 against 17.2 million): in the early years demand barely exceeds Plant 1, so an early Plant 2 mostly sits idle while its capital is already spent. The single spreadsheet sees this timing effect too (47.4 against 45.3 million). Second, and this one it could never show, a low bar of 420,000 units keeps nearly all of that value (18.7 million) while cutting the chance of a loss to 16.7%. It is **better than building now on both counts**. Raising the bar buys more safety at a price: at 500,000 the loss chance is only 1.6%, but the mean falls to 5.4 million. Where to sit on that trade-off is a question of risk appetite, and the table puts numbers on it.

CHANGE THE RULE, THE WHOLE MODEL FOLLOWS

The threshold is a yellow cell on *Assumptions* (name `WaitDemand`). Type 420,000 and the wait column, the histograms, the S-curves and the risk measures all recalculate in about a second.

RECOMMENDATION

Plant 2 is still worth building, and the better way to do it is to wait a year: build in 2028 if 2027 demand reaches 420,000 units. That rule's mean NPV is 18.7 million against 17.2 million for building now, and its chance of a loss is 16.7% against 29.0%. Waiting does not remove the risk: about one future in six still loses money, because demand can stall after the plant is committed.

What the model leaves out. Here waiting costs only time. If a competitor would fill the gap, or customers turned away while Plant 1 is full would not come back, building now is worth more than the model shows. Management's question is whether that is worth more than 1.5 million of mean NPV and a 29% rather than 17% chance of a loss.


What the simulation adds to the decision

The same assumptions, two ways of using them (\$ millions)

	One spreadsheet (expected values)	Monte Carlo (10,000 futures)
What is Plant 2 worth?	Build now 45.3; wait 47.4	Build now 17.2; wait, build if demand $\geq$ 420,000: 18.7
Can it lose money?	No sign of it	A 29.0% chance if built now; the worst 5% of futures lose 79.4 on average
Why wait?	Timing only: the capital is spent a year later	Timing <i>and</i> the right to walk away: the chance of a loss falls from 29.0% to 16.7%
Which wait rule?	Cannot tell: expected 2027 demand (452,047) clears every rule up to that level	The threshold table prices each rule's mean and risk
What to learn first?	Nothing: the future is one number	Demand: it swings the NPV by 127.7 million (section 9)

The point estimate is not wrong arithmetic; it answers a narrower question. "What is the plan worth if everything goes as expected?" gives 45.3 million, 2.6 times the expected value. Asked to compare timing, the same spreadsheet even prefers waiting, but for the wrong reason and at the wrong size, and with no sign of the 29.0% chance of a loss or of the value of keeping the option to walk away. Simulation answers the question the board is actually asking: what is this worth across the futures that might happen, and how do we protect ourselves against the bad ones?

9 What drives the answer



Build now: mean NPV, \$ millions, when each input is in its lowest or highest 10% of trials (`fx.RiskTornado!`, Method "swing", drawn by *Tornado (Sensitivity)*).

Demand in 2036 dominates: when 2036 demand is in its lowest tenth, building now averages -73.1 million; in its highest tenth, 54.5 . That is a swing of 127.7 million. The next two inputs (*Average material cost* and *Capital cost*) swing it by only 19.7 and 14.7 million. The engineers' capital estimate and the construction schedule matter far less than the market. So the most valuable thing to learn before committing is demand, and that is exactly what waiting a year buys.

10 Checking the model

Sheet *Outputs* ends with checks, each shown as simulated, expected and difference, with a blank *Manual sign-off* cell for your ✓:

Check	Simulated	Expected	3 standard errors
Demand 2036, mean (GBM: $\text{start} \times e^{\text{drift} \times 10}$)	707,969	708,950	6,881
Material cost 2036, mean (OU: reverts toward 72.00)	72.05	71.99	0.27
Material cost 2036, standard deviation (OU's exact formula)	9.07	9.13	0.19
Capital cost, mean (PERT: $(\text{min} + 4 \times \text{most likely} + \text{max}) \div 6$)	57.46M	57.50M	0.13M
Share of trials late	24.6%	25.0%	1.3%
Copula: normal-score correlation of capital cost and delay	0.600	0.600	0.019
Trial 1's NPV, rebuilt with Excel's <code>NPV()</code> from its cash flows	-73.95M	-73.95M	difference 0.00

The single-point plan has its own tie-out: Excel's `NPV()` of its cash flows, plus the discounted terminal value, less the table's total is 0.00.

11 Practical advice

- **Model the decision, not just the forecast.** The valuable output here is not a number for Plant 2 but a comparison of plans on the same futures. Keep each strategy's NPV in its own column, and keep the "if built" value so any rule can be tested afterwards.
- **Watch for kinks.** Capacity limits, minimums, thresholds, taxes on profits only, contract caps: wherever the payoff bends, the average input gives the wrong average output, and simulation is worth the trouble.
- **Choose the right time-series model.** Things that grow and compound (demand, revenue, asset values) suit GBM; things pulled back toward a norm (commodity costs, interest rates, margins) suit mean reversion. The wrong choice makes ten-year ranges far too wide or far too narrow.
- **The terminal value can be the whole answer.** In the single-point plan, the terminal value is worth 61 million today, more than the plan's entire NPV of 45 million. The first ten years alone do not pay back the plant. So test the terminal growth rate (a yellow cell), and ask whether the plant really will earn for decades.
- **Link risks that travel together.** Overruns and delays share causes, so a copula ties them ($\rho = 0.60$). Leaving them independent would understate the bad tail. The *Copula* guide covers copulas in depth.
- **Same draws for every strategy.** Give each random input its own VarID, and reuse the same draws across the plans you compare. Then 10,000 trials separate the plans far more precisely than their separate standard errors suggest.
- **Sampling error.** The mean NPV is good to about ± 1.3 million (three standard errors). Percentiles in the far tail need more trials.

12 Function reference

`fx.RiskGBMA(Start, Drift, Volatility, Periods, [dt], [Trials], [VarID], [Seed])`

Argument	Default	Meaning
Start	required	Today's value (> 0); not included in the output.
Drift, Volatility	required	Expected growth a unit of time ($E[S(t)] = \text{Start} \times e^{\text{Drift} \times t}$) and the SD of log changes.
Periods, dt	required, 1	Columns of output, and the length of each period.
Output: Trials × Periods. Trials may also be a Trials × Periods block of uniforms.		

`fx.RiskOUA(Start, LongRunMean, Speed, Volatility, Periods, [dt], [DegFree], [Trials], [VarID], [Seed])`

Argument	Default	Meaning
LongRunMean, Speed	required	The level the process reverts to, and how fast; half-life = $\text{LN}(2) \div \text{Speed}$.
Volatility	required	Instantaneous volatility a unit of time. Exact discretization, so any dt is exact.
DegFree	none	Student t shocks with this many degrees of freedom (> 2), for fatter tails.

`fx.RiskFitGBMA(Prices, [dt]) • fx.RiskFitOUA(Series, [dt])`

Argument	Default	Meaning
Prices / Series	required	History, oldest first.
dt	1	Time between observations in the unit you want the parameters in (1/12 for monthly data, yearly parameters).
Output: GBM 3 × 2 (Drift, Volatility); OU 5 × 2 (LongRunMean, Speed, Volatility, HalfLife).		

Distributions, summaries and charts in the example

Function	Arguments	Notes
<code>fx.RiskPertA</code>	Min, Mode, Max, [Trials]...	Pass a copula column as Trials.
<code>fx.RiskBernoulliA</code>	P, [Trials]...	1 when the uniform exceeds 1 – P, so a high copula uniform means "late".
<code>fx.RiskSPTMetalogA</code>	Low, Median, High, [LowerBound], [UpperBound], [Alpha]...	Three quantiles, P10 / P50 / P90 by default.
<code>fx.RiskCopulaGaussA</code>	Correlation, [Trials], [VarID], [Seed]	See the Copula guide.
<code>fx.RiskVariablesTableA</code>	Variables, [Percentiles], [Names]	Summary statistics, one column per variable.
<code>fx.RiskMeasuresA</code>	Trials, [Confidence], [LossesPositive], [Name]	VaR, CVaR and ES; P&L by default, losses reported positive.
<code>fx.RiskChartFanA</code>	Paths, [Lower], [Upper], [Name], [Start], [Plan], [Labels]	Fan-chart data for a Trials × Periods block.
<code>fx.RiskChartHistA</code>	Trials, [Bins], [Lower], [Upper], [Name]	Histogram data with percentile lines.
<code>fx.RiskChartCDFA</code>	Trials, [Bins], [Points], [Lower], [Upper]	S-curve points.
<code>fx.RiskTornadoA</code>	Output, Inputs, [Labels], [Method], [Tail], [Name]	"swing" (mean output in each input's tails) or "rank".

References: Savage (2009), *The Flaw of Averages*; Dixit and Pindyck (1994), *Investment under Uncertainty*; Glasserman (2003), *Monte Carlo Methods in Financial Engineering*, ch. 3; Keelin (2016), *Decision Analysis* 13(4) (metalog); McNeil, Frey and Embrechts (2015), *Quantitative Risk Management* (VaR, CVaR, copulas). The example's history and assumptions are illustrative, not any real company's. XL Edge and the fx.Risk... functions: © 2026 W Phillips, edgewisedata.com, MIT License.