14 Monte Carlo Risk Simulation

After this chapter you will be able to:

1. explain why a three-column scenario table overstates the corners and leaves the middle empty
2. choose an input distribution for each uncertain input: triangular, uniform, or lognormal
3. compute the mean of a triangular distribution from $(a + m + b)/3$, and distinguish it from the mode that the base case prices
4. recognize that an adverse correlation between two inputs widens the output spread and deepens the bad tail
5. read a simulation output as center and risk: the mean, the P10, P50, and P90, the probability of loss, and the expected shortfall
6. compute a contingency as the chosen cost percentile minus the base estimate

chapter 12 ended with a verdict sentence that carried three numbers: an expected present worth, a cushion, and the probability of one losing state. Behind that sentence stood three hand-priced stories, pessimistic, most likely, and optimistic. Three stories are a start, but the future does not arrive in three versions. Revenues can come in 4% low, or 9% low, or 2% high, and each of those futures has its own present worth. This chapter replaces the handful of stories with thousands of them. Each uncertain input is written down as a *distribution* instead of a point, a computer draws one value from each distribution, prices that future with the same PW model as always, and repeats the draw-and-price loop thousands of times. The output is no longer one number, or three, but a distribution of present worths, and from it the analyst reads what no scenario table can show: the full probability of loss, the percentile spread, and the size of the bad tail. The method is standard practice in project risk work: cost engineering's recommended practices build contingency from exactly this kind of output, and published techno-economic studies, including the rare-earth recovery study this book read in chapter 10, report their headline risk findings as simulation percentiles and frequencies. The chapter builds the method on the crushing unit of chapter 12, first with ten draws a student can follow by hand, then with ten thousand, and closes with the discipline the method demands: a simulation is only as honest as the distributions fed into it.

Why three scenarios are not enough.

Q1. The scenario table of chapter 12 gives three present worths for the crushing unit: −22,811, +8,757, and +32,744 SAR. Can the probability that the project loses money be read from these three numbers?

Q2. An annual saving is described by a triangular distribution with minimum 20,000, most likely value 30,000, and maximum 46,000 SAR. Compute the mean. Is the most likely value a fair single-number summary?

Q3. A simulation of 10,000 iterations produces 2,400 present worths below zero. Estimate the probability of loss. Is the P50 of the output the same as its mean?

Answers.

Q1. Only crudely. If probabilities are attached to the stories, the probability of the one negative story (25% in Example 12.6) is a rough stand-in. But the table says nothing about the futures *between* the stories: a revenue slip of 8% is neither the pessimistic column nor the most-likely one, yet it already produces a loss. Reading $\Pr(\mathrm{PW} < 0)$ properly needs a distribution over the inputs, which is what this chapter builds.

Q2. The triangular mean is the average of the three parameters: $(20{,}000 + 30{,}000 + 46{,}000)/3 = 32{,}000$ SAR. It sits 2,000 SAR *above* the most likely value because the distribution is skewed to the right: the upside stretches 16,000 above the mode, the downside only 10,000 below it. Using the mode as the single estimate understates the average outcome, and using the mean hides the mode. A skewed input needs more than one number.

Q3. The loss probability is estimated by the relative frequency, $2{,}400/10{,}000 = 24\%$. The P50 (the median, the value half the iterations fall below) equals the mean only when the output distribution is symmetric. Project PW distributions usually are not, because the inputs are skewed, so a complete report gives both, plus the loss probability.

Reference. AACE International's recommended practices on cost risk analysis (118R-21, 119R-21) frame contingency setting around exactly the percentile outputs this chapter produces. Full references at the end of the chapter.

Monte Carlo simulation. Repeating a deterministic calculation thousands of times, each time with inputs drawn at random from stated distributions, and reading the collection of outputs as a distribution.

14.1 *From Three Scenarios to a Distribution*

The scenario analysis of Section 12.5 bundles the inputs into three coherent stories and prices each one. Its weakness is not the stories it tells but the futures it leaves out. A three-column table places all of its probability on three exact input combinations and none anywhere else, and this misleads in two directions at once.

1. **The corners are too extreme.** The pessimistic column of Example 12.5 puts the revenue at its floor, the cost at its ceiling, and the salvage at its floor, all in the same future. Every input landing on its worst value simultaneously is possible but very unlikely, and weighting that corner at 25% prices a future that will almost never happen.

2. **The middle is too empty.** Between the most-likely story and the pessimistic one lie all the moderately bad years, revenue 6% low, cost 3% high, and chapter 12 showed the crushing unit's verdict flips at a revenue slip of only 5.5%. Those in-between futures are where most of the loss probability actually lives, and the three-column table assigns them nothing.

Monte Carlo simulation repairs both defects with one move. Instead of three input combinations, it prices thousands of them, drawn at random from distributions the analyst writes down for each input. Futures near the most likely values are drawn often, corner futures rarely, in exact proportion to the probabilities the distributions assign. The method needs no new economics: each iteration is one scenario, priced with the same present-worth model as chapter 12. What is new is only the bookkeeping of thousands of answers.

Insight. Simulation is scenario analysis done ten thousand times by a tireless clerk. Nothing in the PW arithmetic changes; a draw of inputs goes in, one present worth comes out. The gain is that the collection of outputs is a *distribution*, so questions the scenario table cannot answer, the probability of loss, the P10, the size of the bad tail, become counting exercises.

Input distribution. The probability distribution written down for one uncertain input; it states which values are possible and how likely each region of values is.

Triangular distribution. A distribution specified by a minimum a, a most likely value m, and a maximum b. Its density rises straight from a to a peak at m and falls straight to b.

Uniform distribution. A distribution that is flat between a lower bound and an upper bound: every value in the interval is equally likely. Used when only the bounds are known.

Lognormal distribution. A right-skewed distribution for a quantity that is the product of many uncertain factors; it has a hard floor near zero and a long upper tail. A common shape for cost growth.

Reference. AACE RP 119R-21 (2022), Section 2.1, notes the lognormal fit of actual-to-estimate cost growth and builds its p10–p90 range tables from it.

In practice. The data pick the shape. If the maintenance log shows the crusher's availability between 88% and 96% with most months near 93%, that is a triangular; if the only defensible statement about a salvage value is "somewhere between 3,000 and 7,000 SAR," that is a uniform, and pretending to know a peak would be invented information. The bounds are physical before they are statistical: revenue cannot exceed nameplate capacity times price at full uptime, and a cost distribution cannot dip below the input prices already under contract.

14.2 *Choosing the Input Distributions*

The analyst's real work happens before any random number is drawn: turning each point estimate into a distribution. Three shapes cover most engineering-economy practice.

Triangular $(a; m; b)$ is the workhorse, because its three parameters are exactly the three numbers a planning team already produces for a scenario table: the pessimistic value, the most likely value, and the optimistic value. Its mean is the plain average of the three,

$$E[X] = \frac{a+m+b}{3}, \tag{14.1}$$

which matters because the mean and the mode differ whenever the triangle is lopsided.

Uniform $(a; b)$ says less and should be used when less is known: only the bounds can be defended, and no value inside them is more credible than another. Its mean is the midpoint $(a+b)/2$.

Lognormal suits quantities built by multiplication, and cost overruns behave that way: a quantity growth of 10% compounded with a rate growth of 10% gives 21%, not 20%. Cost-risk data bear the shape out; AACE's parametric risk practice notes that the ratio of actual to estimated cost is well fitted by a lognormal, with its long tail on the overrun side.

Example 14.1 — Writing down the crushing unit's input distributions

Turn the scenario estimates of Example 12.5 for the Wadi Phosphate crushing unit into input distributions, using the scenario values as the triangular parameters for revenue and operating cost and a uniform for the salvage. The investment of 55,000 SAR is a firm vendor quote and stays fixed. Then compute the expected present worth of the model and compare it with the base-case PW of +8,757 SAR.

Solution.

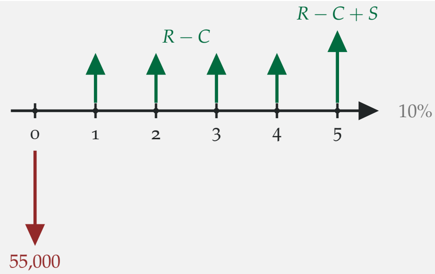

Figure 14.1. The crushing unit with its inputs left as symbols. Each iteration of the simulation fills in one draw of R, C, and S and prices this diagram.

Step 1: one distribution per input.

The scenario columns already contain the parameters:

input	distribution	parameters (SAR)	mean (SAR)
annual revenue R	triangular	(36,000; 42,000; 46,000)	41,333
operating cost C	triangular	(24,000; 26,000; 28,000)	26,000
salvage S	uniform	(3,000; 7,000)	5,000
investment I	fixed	55,000	55,000

Table 14.1. The crushing unit's input distributions. The triangular parameters are the pessimistic, most-likely, and optimistic scenario values of Example 12.5; the salvage keeps only its bounds; the firm quote stays a point.

Step 2: read the shapes.

The cost triangle is symmetric, so its mean equals its mode, 26,000 SAR. The revenue triangle is not: by Equation (14.1),

$$E[R] = \frac{36{,}000 + 42{,}000 + 46{,}000}{3} = 41{,}333 \text{ SAR},$$

which is 667 SAR *below* the most likely value, because the downside reaches 6,000 below the mode while the upside reaches only 4,000 above it. The planning team's own scenario bounds already said the bad side is longer; the distribution simply keeps that statement.

Step 3: the expected present worth.

The PW is linear in each input, so its expectation is the PW of the expected inputs:

$$E[\text{PW}] = -55{,}000 + (41{,}333.33 - 26{,}000)\,(3.7908) + 5{,}000\,(0.6209)$$
$$= -55{,}000 + 58{,}125.60 + 3{,}104.50 = \boxed{+6{,}230 \text{ SAR}}.$$

The base case of Example 12.1 said +8,757 SAR. The 2,527 SAR gap is not an arithmetic slip; it is the revenue skew priced in. Plugging in most likely values is not the same as taking the expected value, and for this project the difference is 29% of the base answer.

Insight. The base case computes the PW of the *modes*; the expected present worth prices the *means*. The two agree only when every input distribution is symmetric. Whenever a team's own scenario bounds are lopsided, the honest center of the project already sits away from the base case, before any simulation is run. Equation (14.1) finds that center in one line.

Correlation. A statement that two inputs tend to move together (positive) or in opposite directions (negative) across futures, rather than independently.

14.3 *Correlation Between Inputs*

Distributions describe each input alone. A simulation also needs to know how the inputs move *together*, because Section 12.5 made exactly this point against one-way sensitivity: in a bad year the inputs do not fail one at a time. If low revenue years tend to be high cost-per-tonne years, drawing R and C independently manufactures futures in which one input is bad and the other conveniently good, and those invented offsets shrink the spread of the output. The demonstration needs nothing more than two inputs with two values each.

Example 14.2 — What correlation does to the spread

Simplify the crushing unit to two inputs: revenue is 46,000 or 36,000 SAR per year with probability 0.5 each, operating cost is 24,000 or 28,000 SAR per year with probability 0.5 each, and the salvage is fixed at 5,000 SAR. Price all combinations under (i) independent draws, (ii) adverse pairing (low revenue always arrives with high cost), and (iii) benign pairing (low revenue always arrives with low cost). Compare the three risk pictures.

Solution.

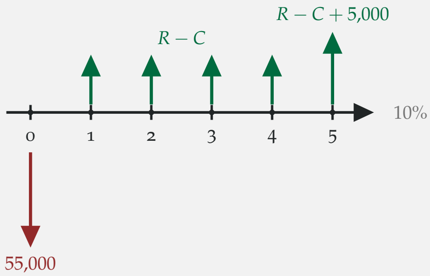

Figure 14.2. The two-input model. Correlation decides which (R, C) pairs can occur together, and with what probability.

Step 1: price the four corners once.

With $\text{PW} = -55{,}000 + (R - C)(3.7908) + 5{,}000\,(0.6209)$:

R (SAR)	C (SAR)	net (SAR/yr)	PW (SAR)
46,000	24,000	22,000	+31,502.10
46,000	28,000	18,000	+16,338.90
36,000	24,000	12,000	−6,405.90
36,000	28,000	8,000	−21,569.10

Table 14.2. The four possible futures of the two-input crushing unit. Correlation does not change these worths; it changes their probabilities.

Step 2: assign probabilities three ways.

Independence gives each corner $0.5 \times 0.5 = 0.25$. Adverse pairing allows only {high R, low C} and {low R, high C}, each with probability 0.5. Benign pairing allows only {high, high} and {low, low}, each 0.5.

Step 3: compare the risk pictures.

Every case has the same mean, +4,966.50 SAR, because correlation does not move the center of a sum; it moves the spread:

	independent	adverse	benign
mean (SAR)	+4,966.50	+4,966.50	+4,966.50
standard deviation (SAR)	20,414	26,536	11,372
$\Pr(\text{PW} < 0)$	0.50	0.50	0.50
$\Pr(\text{PW} < -20{,}000)$	0.25	**0.50**	**0**

Under adverse pairing the worst future, a loss of 21,569 SAR, is twice as likely as independence claims, and the standard deviation is 30% larger. An analyst who lets the software default to independence when the process actually pairs its failures reports a spread that is too narrow and a deep tail that is too thin.

Insight. Ignoring an adverse correlation narrows the output spread and flatters the tail; ignoring a benign one does the reverse. The default in every simulation tool is independence, so the burden is on the analyst to declare the pairings. The direction is a process fact, not a statistical choice, and the coherent-story rule of Section 12.5 is the same rule wearing scenario clothes.

Iteration. One pass of the loop: draw one value from each input distribution, compute the PW of that future, store it.

14.4 *The Simulation Loop*

With distributions and correlations declared, the machine's job is a loop of three lines, repeated thousands of times.

1: **for** $k = 1, 2, \ldots, 10{,}000$ **do**
2: draw $R_k \sim \text{tri}(36{,}000; 42{,}000; 46{,}000)$, $C_k \sim \text{tri}(24{,}000; 26{,}000; 28{,}000)$, $S_k \sim \text{unif}(3{,}000; 7{,}000)$
3: $\text{PW}_k \leftarrow -55{,}000 + (R_k - C_k)\,(P/A, 10\%, 5) + S_k\,(P/F, 10\%, 5)$
4: store PW_k
5: sort the stored PW_k; report percentiles, the mean, and the fraction below zero

Line 2 is where the randomness lives: a uniform random number is shaped into a draw from the target distribution by the inverse-CDF formula, the same one Section 21.7 types into a spreadsheet cell. Line 3 is chapter 12 unchanged, one scenario priced. Everything the chapter reports from here on is a summary of the stored list.

Example 14.3 — Ten draws by hand through the crushing unit

Before trusting a computer with ten thousand iterations, run ten by hand. Table 14.3 lists ten preset input draws, produced once from the distributions of Example 14.1 and rounded to the nearest 100 SAR. Price each draw, then summarize the ten present worths.

Solution.

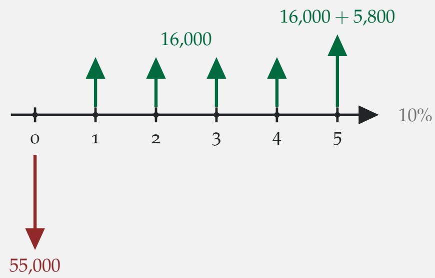

Figure 14.3. Draw 1 of Table 14.3 as a cash flow diagram: net 16,000 SAR per year with a 5,800 SAR salvage. Every iteration is one such diagram with its own arrow heights.

Step 1: price the first draw in full.

Draw 1 gives $R = 42{,}200$, $C = 26{,}200$, $S = 5{,}800$ SAR, so the net is 16,000 SAR per year and

$$\begin{aligned} \mathrm{PW}_1 &= -55{,}000 + 16{,}000\,(3.7908) + 5{,}800\,(0.6209) \\ &= -55{,}000 + 60{,}652.80 + 3{,}601.22 = +9{,}254.02 \text{ SAR.} \end{aligned}$$

Step 2: the other nine, same arithmetic.

k	R_k (SAR)	C_k (SAR)	S_k (SAR)	net (SAR/yr)	PW_k (SAR)
1	42,200	26,200	5,800	16,000	+9,254.02
2	42,400	27,100	3,300	15,300	+5,048.21
3	42,300	25,900	5,800	16,400	+10,770.34
4	44,000	24,900	5,600	19,100	+20,881.32
5	39,700	26,800	3,500	12,900	−3,925.53
6	38,700	27,300	4,500	11,400	−8,990.83
7	41,900	26,100	6,800	15,800	+9,116.76
8	43,400	26,400	4,900	17,000	+12,486.01
9	39,900	26,900	3,900	13,000	−3,298.09
10	37,100	25,100	5,500	12,000	−6,095.45

Table 14.3. Ten preset draws from the crushing unit's input distributions, priced one by one with $(P/A, 10\%, 5) = 3.7908$ and $(P/F, 10\%, 5) = 0.6209$. Each row is one iteration of the loop.

Step 3: summarize the ten.

Sorted, the ten worths run from $-8{,}990.83$ to $+20{,}881.32$ SAR. The median is the average of the fifth and sixth sorted values, $\boxed{+7{,}082 \text{ SAR}}$; the mean is $+4{,}525$ SAR; and four of the ten iterations lose money, a loss frequency of $\boxed{0.40}$.

Step 4: read the ten plainly.

Ten draws show the machinery, not the answer. The next section's ten thousand iterations put the true loss probability near 24%, and this sample of ten says 40%; that is not a contradiction but sampling noise, and the insight below sizes it. Note also what even ten draws already reveal that the base case hid: all four losing futures (draws 5, 6, 9, and 10) lose money without any input touching its pessimistic bound.

Percentile (P10, P50, P90). The value that a stated fraction of iterations falls at or below. P10 is a bad-tail marker, P50 the median, P90 a good-tail marker.

Probability of loss. The fraction of iterations with $\text{PW} < 0$: the simulation's estimate of $\Pr(\text{PW} < 0)$.

Expected shortfall. The average outcome in the bad tail, here reported as the mean of the worst 10% of iterations: if things go badly, how badly on average.

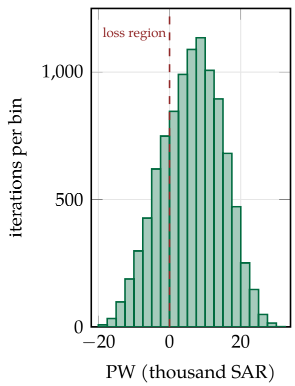

Figure 14.4. Histogram of 10,000 simulated present worths of the crushing unit (bin width 2,500 SAR). The mass left of the dashed line is the probability of loss; the left tail is longer than the right.

Insight. How many iterations are enough? A loss-probability estimate based on n iterations has a standard error of about $\sqrt{p(1-p)/n}$. At $n = 10$ and $p \approx 0.24$ that is 13 percentage points, useless. At $n = 10{,}000$ it is 0.4 of a point, tight enough for any project decision. Iterations are cheap; run thousands, and quote the estimate no finer than its standard error deserves.

14.5 Reading the Output

Running the loop of Section 14.4 for 10,000 iterations (independent draws, a fixed random seed so the run is reproducible) turns the crushing unit into a list of 10,000 present worths. Two pictures and a handful of numbers summarize the list.

Example 14.4 — Reading the crushing unit's simulated distribution

The 10,000-iteration run of the crushing unit produced the summary in Table 14.4. Read the output: report the center, the spread, the loss risk, and the verdict sentence; then compare the simulation's answers with the three-scenario analysis of chapter 12; and finally state what changes if the revenue and cost draws are adversely correlated instead of independent.

Solution.

statistic	independent	correlated ($\rho_{RC} = -0.5$)
mean (SAR)	+6,299	+6,302
standard deviation (SAR)	8,564	9,915
P10 (SAR)	−5,242	−7,177
P50 (SAR)	+6,641	+6,730
P90 (SAR)	+17,206	+19,022
probability of loss	0.242	0.271
expected shortfall, worst 10% (SAR)	−9,019	−11,415

Table 14.4. Summary of 10,000 iterations of the crushing unit model, drawn independently and then with an adverse revenue–cost correlation. Same distributions, same model; only the pairing changed.

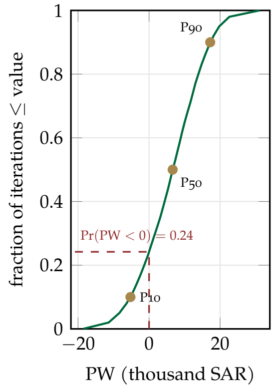

Figure 14.5. The cumulative curve of the same 10,000 iterations. Reading up from PW = 0 gives the loss probability; reading across from 0.10, 0.50, and 0.90 gives the percentiles.

Step 1: center and spread.

The simulated mean, +6,299 SAR, sits next to the exact expectation of +6,230 SAR from Example 14.1; the 69 SAR gap is sampling noise (one standard error of the mean is about 86 SAR here). The middle 80% of futures runs from the P10 of −5,242 to the P90 of +17,206 SAR, a spread of about 22,400 SAR on a 55,000 SAR investment.

Step 2: the loss risk.

2,421 of the 10,000 iterations lost money, so

$$\Pr(\mathrm{PW} < 0) \approx \boxed{24\%},$$

and when the project does lose, the losses average about 5,200 SAR, with the worst tenth of all futures averaging −9,019 SAR. Note the skew the histogram shows: the median (+6,641) sits above the mean (+6,299) because the long tail points down.

Step 3: the verdict sentence.

Extend the rule of chapter 12, center and risk together: *accept the crushing unit; the expected present worth is +6,300 SAR and the median future earns +6,600 SAR, but 24% of futures lose money, the P10 is −5,200 SAR, and the worst tenth averages a loss near 9,000 SAR.*

Step 4: compare with the scenario analysis.

The three-scenario expectation of Example 12.6 was +6,862 SAR with a "25% chance of losing about 22,800 SAR." The simulation refines both halves. Its center is a little lower (+6,230 exact) because the triangular shapes weight the long revenue downside properly, and its tail statement is very different in structure: losses are more frequent in total (24%) but almost never as deep as the pessimistic corner, whose all-bad-at-once combination (−22,811 SAR) was not reached even once in 10,000 draws; the worst draw was −18,575 SAR. Scenarios overstate the corners and understate the middle, exactly as Section 14.1 predicted.

Step 5: the correlated rerun.

With low-revenue futures drawn together with high-cost futures ($\rho_{RC} = -0.5$), the mean barely moves, but every risk number worsens: the spread grows 16%, the P10 deepens to −7,177 SAR, the loss probability rises to 27%, and the worst tenth deepens to −11,415 SAR. If the process pairs its failures, independence was flattery.

Contingency. An amount added to the base cost estimate to cover outcomes the base did not price; set from the cost distribution as a chosen percentile minus the base.

Reference. AACE RP 118R-21 (2022) covers contingency from estimate ranging with Monte Carlo simulation; RP 119R-21 (2022) covers the p10–p90 accuracy ranges and contingency it must reproduce. Full references at the end of the chapter.

14.6 Contingency from the Simulation

When the simulated quantity is a project *cost* rather than a present worth, the same output distribution answers the estimator's question from chapter 9: how much must be added to the base estimate so that the budget survives the futures the base ignored? The recipe is direct. Simulate the total cost, pick the confidence level the owner wants to fund, and set

$$\text{contingency} \;=\; \text{cost percentile at the chosen confidence} \;-\; \text{base estimate}. \tag{14.2}$$

Practice varies on the confidence level, for good reason. An owner funding a portfolio of many projects can budget each near P50 and let overruns and underruns cancel across the portfolio; an owner with one project that must not overrun funds higher, commonly P80, with cost-risk practice noting reserves as high as P90 for large single investments. The accuracy ranges attached to the estimate classes of chapter 9 are statements about this same distribution: a class range quoted as $-20\%/+30\%$ is a p10–p90 interval around the funded amount, and the simulation is how a specific project earns its own range instead of inheriting the class average.

Example 14.5 — Contingency for a crusher installation estimate

The installation estimate for a second crushing unit has three line items: the equipment package, 40,000 SAR, on a firm fixed-price quote; site construction, estimated triangular (12,000; 15,000; 22,000) SAR; and commissioning, known only to lie between 3,000 and 6,000 SAR (uniform). The base estimate sums the point values: $40{,}000 + 15{,}000 + 4{,}500 = 59{,}500$ SAR. Simulating the total 10,000 times gives a mean of 60,818 SAR, a P50 of 60,633 SAR, and a P80 of 62,849 SAR. Set the contingency at the P80 policy level and read the result.

Solution.

Step 1: the contingency.

By Equation (14.2),

$$\text{contingency} = 62{,}849 - 59{,}500 = \boxed{3{,}349 \text{ SAR } (5.6\% \text{ of base})},$$

so the funded budget is about 62,850 SAR, an amount the simulated cost stays within in 80% of futures.

Step 2: why the base needed protecting.

The simulated total exceeded the 59,500 SAR base in 68% of iterations. That is not bad luck; it is the construction triangle's shape. Its mode is 15,000 but its mean is $(12{,}000 + 15{,}000 + 22{,}000)/3 = 16{,}333$ SAR, because the overrun side stretches 7,000 above the mode and the underrun side only 3,000 below. A base built from most-likely values is a roughly P30 estimate here, and the contingency is the bridge from the P30 story the estimate tells to the P80 certainty the owner wants to fund.

Step 3: where the risk lives.

The firm quote contributes nothing to the spread; all of the 3,349 SAR of contingency is earned by the two estimated items, almost all of it by the construction triangle. The estimator's follow-up effort, tightening drawings and trenching surveys, belongs to that line, the same effort-follows-the-swing rule as Example 12.3.

In practice. Split the estimate by the quality of its sources before simulating. A signed fixed-price quote is a constant; a unit-rate estimate from preliminary drawings is a distribution; an allowance for unmapped ground is a wide one. Giving every line the same ± percentage, a common shortcut, erases exactly the information the contingency is supposed to price.

14.7 *When Not to Simulate*

A simulation output looks authoritative: four-digit percentiles, a smooth histogram, a loss probability to one decimal. None of that precision is evidence. The output is a deterministic consequence of the input distributions, and if those were guessed casually, the histogram is the same casual guess wearing a laboratory coat.

> **Insight.** Garbage distributions in, authoritative-looking garbage out. A Monte Carlo run adds no information to the inputs; it only translates the stated input uncertainty into output form, at the stated correlations. The step that deserves the review meeting is Table 14.1, the distribution choices and their sources, not the histogram. If the team cannot defend the bounds, the modes, and the pairings, the honest tools are the simpler ones of chapter 12, which at least display their guesses openly.

Cost-risk practice draws the same boundary. AACE's recommended practice for estimate ranging with Monte Carlo (118R-21) restricts the method to projects whose scope is well defined, Class 3 or better in the classes of chapter 9, with no new technology and minimal complexity, and warns that below that definition the method underestimates contingency, often by a lot: team-supplied ranges cannot see risks that have no line item yet. Three checks before trusting any simulation:

1. **Is the scope defined enough to range?** If the flowsheet itself may change, no distribution on today's line items covers tomorrow's line items.
2. **Do the distributions have sources?** Process limits, vendor quotes, historical logs, or at least a scenario workshop, for every input that matters in the tornado ranking.
3. **Are the pairings declared?** An independence default should be a decision, not an accident (Example 14.2).

14.8 *Case Study from the Literature*

chapter 10 read the rare-earth recovery study of Smerigan and Shi (2025) for its tax machinery. The same open-access paper carries a full probabilistic risk section, which is why it returns here: the authors wrap their techno-economic model in exactly this chapter's loop, three thousand Monte Carlo samples over uncertain prices, costs, and process performance, and report their headline finding as a probability, not a number: in 87% of baseline simulations, the internal rate of return exceeds the 15% hurdle.

Reference. Smerigan & Shi (2025), arXiv:2504.10495. Open access; full reference at the end of the chapter.

Example 14.6 — Case study: reading a published Monte Carlo risk section

Smerigan and Shi (2025) report a baseline NPV of $570 million at a 15% discount rate and, from 3,000 Monte Carlo samples, an IRR above 15% in 87% of baseline simulations; amounts are in US dollars, the paper's currency. (a) Translate the 87% into this chapter's loss probability. (b) Judge whether 3,000 iterations are enough for the reported precision. (c) State what this chapter's checklist would ask of the paper next.

Solution.

Step 1: from an IRR frequency to a loss probability.

The plant is a simple investment, one outlay followed by returns, so by the relation of chapter 7, IRR $>$ MARR exactly when PW > 0 at the MARR. The paper's statement is therefore this chapter's statement:

$$\Pr(\text{PW} > 0 \text{ at } 15\%) \approx 87\% \qquad \Longleftrightarrow \qquad \Pr(\text{loss}) \approx \boxed{13\%}.$$

A $570 million base NPV with a 13% chance of failing the hurdle is the center-and-risk sentence of Example 14.4 at industrial scale.

Step 2: is 3,000 enough?

The standard error of a proportion estimated from 3,000 iterations is

$$\sqrt{\frac{0.87\,(1-0.87)}{3{,}000}} = 0.006,$$

about six-tenths of a percentage point. Quoting "87%" to the nearest point is fully supported; quoting 87.3% would not have been. The paper's iteration count and its reporting precision match, which is one of the quiet marks of careful simulation work.

Step 3: the checklist.

This chapter would ask three things. First, are the input distributions stated with sources? The paper samples prices, costs, and performance over stated ranges, with the rare-earth basket price the dominant uncertainty, its analog of Table 14.1. Second, are correlations declared? Commodity prices and reagent costs can move together, and Section 14.3 showed what the independence default does to the tail. Third, where does the probability mass cross zero? The paper's scenario analysis answers in engineering coordinates: the system stays profitable only at capacities above 100,000 kg of phosphogypsum per hour and feed grades above 0.5% rare-earth content, the physical break-even frontier behind the 13%.

14.9 Chapter Summary

- **Monte Carlo simulation** repeats the ordinary PW calculation thousands of times with inputs drawn from stated distributions. Each iteration is one coherent scenario; the output is a distribution of present worths.

- Three discrete scenarios **overstate the corners and empty the middle**: all-bad-at-once futures almost never happen, while the moderately bad futures that carry most of the loss probability are invisible between the columns.

- **Distribution choice is the analysis.** Triangular $(a; m; b)$ when a team can defend three points, with mean $(a + m + b)/3$; uniform when only bounds are known; lognormal for multiplicative quantities such as cost growth. Bounds are process limits before they are statistics.

- The base case prices the modes; the expectation prices the means. For skewed inputs the two differ before any simulation is run (crushing unit: +8,757 vs. +6,230 SAR).

- **Correlation moves the spread, not the center.** Adverse pairings (revenue down with cost up) widen the output distribution and deepen the tail; the independence default narrows them. The pairing is a process fact.

- Read the output as **center and risk**: mean and P50; P10 and P90; the probability of loss; and the expected shortfall (the average of the bad tail). The crushing unit at 10,000 iterations: mean +6,300, P10 −5,200, P50 +6,600, P90 +17,200 SAR, loss probability 24%, worst tenth near −9,000 SAR.

- A proportion estimated from n iterations carries a standard error of $\sqrt{p(1-p)/n}$: 13 points at $n = 10$, 0.4 of a point at $n = 10{,}000$. Quote no finer than the noise.
- **Contingency** = chosen cost percentile − base estimate. Portfolio owners fund near P50; single must-not-overrun projects fund near P80. A base built from most-likely values on right-skewed items sits near P30 and is exceeded in most futures, which is why the contingency exists.
- **When not to simulate**: undefined scope, sourceless distributions, undeclared correlations. The output adds no information to the inputs; it only dresses them in percentiles.

14.10 *Exercises*

These problems build the simulation from its parts: first choosing and parameterizing input distributions, then hand-running single iterations, then reading percentiles and loss probabilities off simulation output, and finally setting contingency from a simulated cost distribution.

Distribution warm-ups and lecture practice

Exercise 14.1. An annual revenue is described as triangular with minimum 30,000, most likely 36,000, and maximum 48,000 SAR. (a) Compute the mean. (b) Compute the probability that the revenue lands at or below the most likely value, using the triangular CDF fact $F(m) = (m-a)/(b-a)$. (c) In one sentence: what would an analyst who replaced this input with the single number 36,000 SAR be assuming, and why is it wrong here?

Exercise 14.2. A project requires 30,000 SAR today and returns a uniform net amount A per year for 4 years, no salvage, MARR = 10%. Five preset draws of A from its distribution are: 8,600; 10,400; 9,100; 11,800; 7,900 SAR per year. (a) Price each draw. (b) Estimate the probability of loss from the five iterations, and check the count against the break-even value $A^* = 9{,}464$ SAR per year from Exercise 12.1. (c) How trustworthy is a loss probability estimated from five draws?

Quick checks (multiple choice)

Exercise 14.3. A simulation of a project reports P10 $= -4{,}000$ SAR, P50 $= +5{,}000$ SAR, P90 $= +15{,}000$ SAR. Which statement is **correct**?

(A) The project loses money in exactly 10% of iterations.
(B) The probability of loss is between 10% and 50%.
(C) The expected present worth is $+5{,}000$ SAR.
(D) 90% of iterations exceed $+15{,}000$ SAR.

Exercise 14.4. A simulated PW distribution is left-skewed, with mean $+4{,}000$ SAR, median $+5{,}500$ SAR, and a loss probability of 20%. Which statement is **correct**?

(A) In a large enough simulation the mean and median must agree.
(B) The long tail is on the loss side, and reporting the mean alone hides both the tail and the 20% loss chance.
(C) A 20% loss probability means the project loses 20% of its value.
(D) A median above the mean shows the simulation is faulty.

Contingency practice

Exercise 14.5. A conveyor extension estimate has three line items: drive package, 30,000 SAR on a firm quote; structure, triangular (38,000; 45,000; 60,000) SAR; erection, uniform between 8,000 and 12,000 SAR. The base estimate is $30{,}000 + 45{,}000 + 10{,}000 = 85{,}000$ SAR. A 10,000-iteration simulation of the total gives: mean 87,701 SAR, P50 87,279 SAR, P80 92,072 SAR, and the total exceeds the base in 69% of iterations. (a) Set the contingency at P80 and express it as a percentage of base. (b) Explain the 69% in one sentence. (c) Which line item earns the contingency, and what estimating action would shrink it?

Exam-style problems

Exercise 14.6. *Najd Copper* restudies the portable ore-sorting unit of Exercise 12.8 with simulation: investment 80,000 SAR; life 5 years; salvage fixed at 8,000 SAR; MARR $= 10\%$; annual net cash flow $A \sim$ triangular (16,000; 22,000; 27,000) SAR per year. Eight preset draws of A are: 21,300; 18,200; 24,500; 16,900; 22,800; 20,100;

25,900; 19,400. (a) (*8 pts*) Price all eight draws (factor notation first). (b) (*5 pts*) Estimate the loss probability from the eight iterations and verify the loss count against the break-even net $A^* = 19{,}793$ SAR per year. (c) (*6 pts*) Compute the expected present worth of the model from the triangular mean, and compare it with the most-likely-value PW of +8,365 SAR. (d) (*6 pts*) The full run of 10,000 iterations is to be quoted in a report. State the standard error of a loss probability near 0.30 at $n = 8$ and at $n = 10{,}000$, and write the one-sentence rule this implies for quoting simulation output.

Exercise 14.7. The projects office at *Wadi Phosphate* must fund a pipeline tie-in. Line items: engineering, 60,000 SAR under a fixed-fee contract; construction, triangular (100,000; 120,000; 165,000) SAR; commissioning, uniform between 15,000 and 25,000 SAR. The base estimate is $60{,}000 + 120{,}000 + 20{,}000 = 200{,}000$ SAR. A 10,000-iteration simulation gives:

	mean	P50	P80	P90
total cost (SAR)	208,235	206,961	220,784	228,081

and the total exceeds the base in 69% of iterations. (a) (*6 pts*) Company policy funds single critical projects at P80. Set the contingency and express it as a percentage of base. (b) (*6 pts*) Explain, with the relevant triangular mean, why the base is exceeded in 69% of futures even though every line item was estimated at its most likely value. (c) (*5 pts*) If the office funded at P50 instead, what overrun probability would it be accepting, and for what kind of owner is that policy defensible? (d) (*8 pts*) Which line item drives the spread, and what does the gap between P90 and P50 (about 21,100 SAR) represent in the language of chapter 9's accuracy ranges?

Assessing outside recommendations

Real recommendations arrive finished: a memo, a sheet, a verdict. The two problems below are audits. Decide whether to accept the recommendation, name every flaw, and correct the numbers. The flawed work is quoted exactly as received.

Exercise 14.8. An AI assistant was asked to simulate the crushing unit of Example 14.1 and returned:

> *"Inputs: revenue* $\sim$ *triangular* (36,000; 46,000; 46,000)*; operating cost* $\sim$ *triangular* (24,000; 26,000; 28,000)*; salvage* $\sim$ *uniform* (3,000; 7,000)*; investment fixed at 55,000 SAR; 10,000 iterations. Results: mean PW* +11,048 *SAR, P50* +12,458 *SAR, probability of loss 15%. The distribution of outcomes confirms the project is safely acceptable."*

The planning team's estimates for the revenue are those of Example 12.5: pessimistic 36,000, most likely 42,000, optimistic 46,000 SAR. Accept or reject the summary? Name each flaw and correct the numbers.

Exercise 14.9. A memo to the *Najd Copper* investment committee reads:

> *"We simulated the ore-sorting unit with 10,000 iterations. The average present worth is* +6,900 *SAR, so the project creates value and the committee can approve it. Detailed percentile tables are omitted for brevity."*

The full run behind the memo shows a left-skewed distribution with P10 = −4,736 SAR, P50 = +7,184 SAR, P90 = +18,181 SAR, and a loss probability of 23%. Accept or reject the memo? Name each flaw and write the sentence the committee should have received.

References

AACE International (2022). *Cost Risk Analysis and Contingency Determination Using Estimate Ranging for Inherent Risks with Monte Carlo Simulation*. Recommended Practice No. 118R-21, AACE International, Morgantown, WV. `https://source.aacei.org/2022/11/08/new-aace-recommended-practice-fills-gap-in-quantitative-risk-analysis-maturity-modelnew-rp-118r-21-cost-risk-analysis-and-contingency-determination-using-estimate-ranging-for-inherent-risks/`

AACE International (2022). *Cost Estimate Accuracy Range and Contingency Determination Using Tables Derived from Parametric Risk Models*. Recommended Practice No. 119R-21, AACE International, Morgantown, WV. `https://web.aacei.org/docs/default-source/toc/toc_119r-21.pdf`

Park, C. S. (2012). *Fundamentals of Engineering Economics* (3rd ed.). Pearson. [Cited throughout as "Park, 3rd ed."]

Smerigan, A., & Shi, R. (2025). Advancing sustainable rare earth element recovery from phosphogypsum. *arXiv preprint*, arXiv:2504.10495. `https://arxiv.org/abs/2504.10495`