CHARTA LEARNING · LESSON NOTES

Simulation: three ways to rehearse the future

CFA Level I · Quantitative Methods

You will learn: Explain and tell apart Monte Carlo simulation, historical simulation and bootstrap resampling, and say what each is used for and where each fails.

Running example: Examples use made-up figures, such as £10,000 invested for one year, and a $50m portfolio of $30m of shares and $20m of bonds.

1. Why simulate, and Monte Carlo

- When the rules are too tangled for a formula, act it out many times and look at the results. A single number (a point estimate) hides the range, like a river averaging 1 metre deep with a 3 metre hole. A simulation gives many outcomes, whose pattern is a distribution.
- Monte Carlo simulation takes a model (inputs in, an outcome out). Each uncertain input gets an assumed distribution; the computer draws random values from them. The futures are invented from rules a person chose.
- Recipe: 1. choose a distribution for each uncertain input (set the correlation if inputs move together); 2. draw one random value for each; 3. run the model: one trial, one outcome; 4. repeat thousands of times and read the results.
- Use it when there is no closed-form solution (no neat formula): path-dependent options, retirement plans with variable spending, portfolios with many interacting risks.

Estimate = average of trial outcomes. Noise shrinks like 1 ÷ √(number of trials)

Watch out: It works out the consequences of your assumptions precisely, but cannot say if they are right. Wrong correlation or thin tails make disasters look too rare.

2. Trials and noise

- Every trial is a random draw, so the answer has random error. More trials give a steadier answer, but with slowly fading returns: to halve the noise you need four times as many trials.
- More trials remove random noise only. They make a wrong model more precise, not more accurate. The fix is better inputs, such as realistic correlation and fatter tails.

£10,000 FOR ONE YEAR; END VALUE = 10,000 × (1 + RETURN); ASSUMED AVERAGE RETURN 4%

Five trial returns	+12%, −8%, +5%, −15%, +20%	
End values	11,200; 9,200; 10,500; 8,500; 12,000	**sum £51,400**
Average end value	51,400 ÷ 5	**£10,280 (+2.8%)**
Chance of a loss	trials 2 and 4 below 10,000: 2 ÷ 5	**40%**
Worst 1% tail with 100 / 10,000 / 1,000,000 trials		**1 / 100 / 10,000 trials**
100 to 10,000 trials, then to 1,000,000	√100 = 10 each time	**noise ÷ 10 each**

Watch out: Thinking more trials fix a wrong model. A 2% underfunding risk from 10,000 trials stays too low at 1,000,000 trials if correlation and fat tails were ignored.

3. Historical simulation

- It uses what actually happened, with no invented draws. Pick a look-back period (say the last 500 days). Each past day becomes one scenario: apply that day's market moves to today's positions. Together the results form an empirical distribution, built from observed data rather than a formula.
- Strengths: no assumption about the shape of returns. Fat tails come for free, and co-movement comes for free because each day is taken whole, so moves that happened together arrive together. No correlation estimate is needed.
- Weakness: it can only replay what happened. An event that never happened cannot appear, and a calm window gives a calm answer. The window drives the result: too short misses rare disasters, very long may describe a vanished world. It gives only as many scenarios as history has periods.

Watch out: Thinking historical simulation needs a distribution. It does not, but it can only show what the window contains.

4. Value at risk by historical simulation

- Scenario profit or loss is the sum, over all positions, of position size (money held today) times that past day's percentage move. Collect all scenarios and sort worst to best.
- Value at risk (VaR) is the loss you expect to beat only on the worst few per cent of days. To read it, take the k-th worst result where k = tail % × number of days. At 5% with 100 scenarios, that is the 5th worst.
- The window is the model: a calm window and a crisis window give very different answers for the same portfolio. A common defence is to run a deliberately stressed window alongside the normal one.

> **Scenario P&L = Σ (position × that day's move). Value at risk = k-th worst result**

SHARES $30M, BONDS $20M, 100-DAY WINDOW, 5% VAR (ILLUSTRATIVE)

Day 1: shares −3.0%, bonds +0.8%	30 × (−3.0%) + 20 × (+0.8%) = −0.90 + 0.16	**−$0.74m**
Day 2: shares −1.5%, bonds −0.4%	−0.45 − 0.08	**−$0.53m**
Five worst of 100	−1.12, −0.96, −0.88, −0.81, −0.74	**$m**
5% VaR	5th worst (5% of 100 = 5)	**$0.74m**
Calm window A / crisis window B	5th worst loss	**$0.42m / $1.35m**
B ÷ A	1.35 ÷ 0.42	**about 3.2×**

5. Bootstrap resampling

- It starts with real data, like historical simulation, but draws returns at random with replacement (each draw is put back, so the same return can appear many times) instead of replaying days in order, once each.
- Because you can keep drawing, you can build as many paths as you like, such as 10,000 paths of 30 years from 20 years of data. The price is that every number in every path really happened, so no resample can contain a value outside the data.

DATA: +4%, −2%, +7%, +1%, +10% (AVERAGE 4.0%)

Original sample	20 ÷ 5	**4.0%**
Resample A: 7, 7, 4, 10, −2	26 ÷ 5	**5.2%**
Resample B: 1, 1, 1, −2, 10	11 ÷ 5	**2.2%**
Biggest value any resample holds	never above 10	**10%**

Watch out: Mixing up historical simulation (each period once, in order) and bootstrap (random, with replacement).

6. Telling the three apart

- Ask one question: where do the scenarios come from? Historical simulation: the real past, each period once, so only as many scenarios as periods. Bootstrap: the real past, redrawn at random, so unlimited paths from values that happened. Monte Carlo: a distribution somebody chose.
- Only Monte Carlo can produce a result worse than anything in the record, because its scenarios come from a rule, not a list. That freedom is also its risk: the distribution and its settings are the analyst's choices, and if any is wrong every trial repeats the mistake.

20 YEARS OF ANNUAL RETURNS, WORST YEAR −37%, 30-YEAR HORIZON

Historical simulation	one-year scenarios only	**20, worst −37%**
Bootstrap	paths of 30 draws	**10,000, worst −37%**
Monte Carlo (normal, mean 7%, sd 16%)	a −50% year: z = (−50 − 7) ÷ 16	**z = −3.56, rare but possible**

Watch out: Thinking bootstrap or historical simulation can show something worse than any past outcome.

Key takeaways

- Simulation acts out many possible futures to show the range of outcomes, not one number.
- Monte Carlo draws from assumed distributions and is used when there is no closed-form solution.
- More trials cut random noise (by 1 ÷ √trials) but never fix a wrong model: precise is not accurate.
- Historical simulation applies real past changes to today's portfolio; the window is the model.
- Bootstrap redraws real returns at random with replacement: unlimited paths, only seen values.
- Only Monte Carlo can go beyond past extremes. Ask where the scenarios come from.

Quick check

1. 100,000 trials give a very precise loss estimate. What tempers your confidence?
   Answer: The inputs may be wrong. Precise is not accurate.
2. Which method can show a loss worse than any in the data?
   Answer: Monte Carlo: scenarios come from a chosen distribution.
3. A calm look-back window is used in historical simulation. Effect on VaR?
   Answer: It looks too small: no crisis days in the sample.