Scenarios and simulation

Portfolio Management

1 A story: the river that is ankle deep

- An average can hide where the deep water is
- Cut history into regimes: recession, boom, rate shock
- Compare return, risk, drawdown and correlation in each
- Ask: does the strategy depend on one environment?

IN THE LESSON

- Here is an old story.
- A traveller reads that a river is on average only ankle deep, so he wades in.
- The river was shallow for most of its width and over his head in the middle.
- The average was true and also useless.
- Strategy returns work the same way.
- A strategy might earn a lovely average over twenty years, but if the profits came in calm years and the losses came in a crisis, the average hides the danger.
- Historical scenario analysis is the habit of checking the depth at every point, not just the average.
- We cut the history into chunks that share a situation, and look at each chunk on its own.

2 How a scenario analysis is done

From history to a verdict

Choose regimes → Split the months → Measure each piece

Compare the pieces → Count the episodes

IN THE LESSON

- Let us lay out the method as a chain.
- Step one, choose the regimes.
- These are stretches of history that share a situation, such as recessions, inflation shocks, rising rate periods or crises.
- Step two, split the data, so every month goes into exactly one bucket.
- Step three, measure each bucket separately: the average return, the volatility, the worst drawdown, and the correlation with other holdings.
- Step four, compare the buckets with each other and with the whole sample.
- And the last step is the one people skip, which is to ask how many independent episodes sit behind each bucket, because that tells you how much to trust it.

3 Heron Carry: the whole sample hides a split

Heron Carry, 240 months of backtested returns

Regime	Mean a month
Expansion, 192 months	+1.0%
Recession, 48 months	−1.5%
Correlation with equities, expansion	0.1
Correlation with equities, recession	0.7

IN THE LESSON

- Now our first case.
- Heron Carry is a made-up strategy with two hundred and forty months of backtested returns, that is twenty years.
- We split them by the economy.
- In one hundred and ninety two expansion months it averaged plus one per cent a month.
- In forty eight recession months it averaged minus one and a half per cent a month.
- Its correlation with equities was zero point one in expansions, but zero point seven in recessions.
- And those forty eight recession months came from only three separate episodes.
- Keep these figures in mind, because on the next slides we turn them into a verdict, and the verdict is not the flattering one.

4 The full sample is a weighted average of regimes

Full mean = Σ (weight × regime mean)

Regime weight	Months in regime ÷ total
Regime mean	Average return in the regime
Full-sample mean	The blended headline number

IN THE LESSON

- Before we read the regimes, let us see how the whole sample average is built.
- It is simply each regime's mean, weighted by the share of months spent in that regime.
- The weight is the number of months in the regime divided by the total months.
- Add up the weighted means and you get the full sample mean.
- This matters because it shows that the headline number is a blend.
- A good regime that lasts long can pull the average up and cover a bad regime that is short but painful.
- When you see a single average in a report, imagine this blend behind it and ask what the pieces were.

5 Example 1: the blended average

Weights 0.8 and 0.2; means +1.0% and −1.5%

Expansion part	0.8 × 1.0	**+0.8%**
Recession part	0.2 × −1.5	**−0.3%**
Full-sample mean	0.8 − 0.3	**0.5% a month**

IN THE LESSON

- Example one.
- The weights are one hundred and ninety two over two hundred and forty, which is zero point eight, and forty eight over two hundred and forty, which is zero point two.
- Expansion contributes zero point eight times one per cent, which is plus zero point eight.
- Recession contributes zero point two times minus one and a half, which is minus zero point three.
- Add them and the full sample mean is zero point five per cent a month.
- That sounds like a modest, steady return.
- But look at what it is made of: a good return four months out of five, and a heavy loss in the fifth.
- The steady look is only an effect of averaging.

6 Example 2: diversification when you need it

Half and half, each with volatility 10%

Correlation in expansion		**0.1**
Portfolio volatility	10 × √((1+0.1)/2)	**7.4%**
Correlation in recession		**0.7**
Portfolio volatility	10 × √((1+0.7)/2)	**9.2%**

IN THE LESSON

- Example two looks at correlation.
- Suppose Heron holds half of its money in Heron Carry and half in an equity fund, and each has a volatility of ten per cent.
- The portfolio volatility is ten times the square root of one plus the correlation, over two.
- With a correlation of zero point one, as in expansions, that is seven point four per cent, a good reduction from ten.
- But with a correlation of zero point seven, as in recessions, it is nine point two per cent, barely below ten.
- The diversification was real in the good times and mostly gone in the bad times, when it was needed.
- That is the second thing a scenario analysis reveals.

7 Risk reduction fades in stress

Portfolio volatility (%)

Alone	**10.0**
Expansion mix	**7.4**
Recession mix	**9.2**

Same two holdings, same weights; only the regime changes

IN THE LESSON

- Here are those volatilities as bars.
- The vertical axis shows portfolio volatility in per cent a year.
- The first bar is the stand alone risk of ten per cent, which is what you would have with no diversification at all.
- The second bar is the portfolio in expansion, at seven point four.
- Notice how much the mix helps there.
- The third bar is the same portfolio in recession, at nine point two.
- The gap between the second and third bars is diversification that disappeared.
- The picture explains why it matters that correlations often rise in stress: the protection was priced into the calm months and was not there on the day it was needed.

8 Reading the result with care

- Expansion: a diversifier; recession: a loser
- Forty eight months, but only three independent episodes
- The next episode may not look like any past one
- A conclusion should be conditional on the regime

IN THE LESSON

- So what is the verdict on Heron Carry?
- In expansions it looks like a diversifier, earning one per cent a month with almost no link to equities.
- In recessions it loses a per cent and a half a month and moves with equities.
- An investor wanting protection would find little.
- But we must be careful in the other direction too.
- Three recessions give us three observations of how the strategy behaves under stress, not forty eight.
- The months inside one recession are not independent evidence.
- The next downturn may have a different cause and hit the strategy differently.
- So the honest conclusion is conditional, and it is held with modest confidence.

9 Four mistakes in scenario analysis

- Trusting the average that hides the bad regime
- Counting months as if they were independent episodes
- Assuming average correlation holds in a crisis
- Picking regimes after seeing which ones look good

IN THE LESSON

- Let us stop and collect the traps.
- First, judging a strategy by its full sample average when it loses heavily in the one regime that matters most to you.
- Second, treating forty eight months as forty eight pieces of evidence, when they come from three episodes.
- Third, assuming the diversification you measured on average will turn up in a crisis, when correlations tend to rise exactly then.
- And fourth, choosing the regimes after looking at the data, so that they flatter the result.
- Pick the regimes first, from economic reasoning, and report all of them, including the ones you would rather not show.

10 A story: dice or old diaries

- Both build thousands of possible futures
- Historical: redraw days from the actual past
- Monte Carlo: draw from a model you specify
- The question is where the tails come from

IN THE LESSON

- Now to simulation.
- Imagine you want to know how bad a bad year could be, but you only have one history, and history happened only once.
- You can do one of two things.
- Option one: take the diaries of the past twenty years, cut them into single days, put the days in a hat, and draw them out at random to build thousands of new made up years.
- That is historical simulation, a bootstrap.
- Option two: study the diaries, work out the average, the typical swing and how things move together, then throw dice that follow those rules.
- That is Monte Carlo simulation.
- Both give thousands of possible futures.
- They differ in where the randomness comes from.

11 Two machines for making futures

Historical simulation

Take actual returns → Redraw whole dates → String into paths

Monte Carlo simulation

Choose a model → Set its numbers → Draw random returns

IN THE LESSON

- Here are the two machines side by side.
- The top row is historical simulation.
- First we take the actual returns we have.
- Next we pick whole dates at random, with replacement, which means a date can be picked more than once.
- Picking a whole date matters, because it carries every asset's return for that day together, so the real correlations come along for free.
- Then we string the picks into a path.
- The bottom row is Monte Carlo.
- We choose a statistical model, estimate its numbers from data or set them by judgement, and then draw random returns from it, path after path.
- Same output, a different engine.

12 Historical simulation: what you get

Historical simulation, in four lines

Fat tails and skew	Kept as they were
Correlations	Real ones kept
Outcomes never seen	Cannot appear
Hidden assumption	Past is typical

IN THE LESSON

- Let us weigh the historical, or bootstrap, approach.
- Its great strength is realism with no effort.
- Because it reuses real returns, the fat tails are exactly as fat as they were in the sample.
- Because it reuses whole dates, the correlations between assets are the real ones, and we never had to model them.
- But it has a ceiling.
- It can only combine outcomes that have already happened, so it can never produce a month worse than the worst month in the data.
- And it assumes that the period we sampled is representative of the future.
- If the history was unusually calm, or unusually wild, the simulation inherits that.

13 Monte Carlo: what you get

Monte Carlo simulation, in four lines

Fat tails and skew	Only if modelled
Correlations	Must be specified
Outcomes never seen	Can appear
Hidden assumption	The model is right

IN THE LESSON

- Now the Monte Carlo approach.
- Its strength is freedom.
- Because it draws from a model, it can produce outcomes that never happened, and it can test hypothetical conditions, for example what if volatility were double, or what if correlations all went to nought point nine.
- You can even build in regime shifts.
- The price is that the answer is only as good as the model.
- A normal distribution has thin tails, so it treats extreme losses as almost impossible.
- The correlations and volatilities must be estimated or guessed, and mistakes there flow straight into the result.
- So Monte Carlo is flexible but depends on the quality of the model chosen.

14 Example 3: the worst month, two ways

Worst month −18%; normal model mean 0.7%, volatility 4.5%

Bootstrap: worst possible month		**−18%**
Normal model: z of −18%	(−18 − 0.7) ÷ 4.5	**−4.16**
Chance in one month		**0.0016%**
Waiting time	1 ÷ chance ÷ 12	**about 5,100 years**

IN THE LESSON

- Example three shows the two machines going wrong in opposite directions.
- Heron Multi Asset has two hundred and forty months of returns, and its worst month lost eighteen per cent.
- A bootstrap of ten years draws one hundred and twenty months from those two hundred and forty.
- The worst month it can ever produce is minus eighteen per cent.
- Now a normal model with a mean of zero point seven and a volatility of four and a half per cent a month.
- The z value for minus eighteen is minus eighteen minus zero point seven, over four point five, which is minus four point one six.
- That is about sixteen chances in a million a month, roughly once in five thousand years, yet it did happen in the sample.

15 From few assumptions to many

- Bootstrap of dates
- Fat-tailed model
- Normal model

IN THE LESSON

- Let us place the methods on one line.
- On the left are approaches that lean on the data and assume little.
- The bootstrap of whole dates sits furthest left, because it assumes only that the sampled history is representative.
- In the middle is a fat tailed model, with parameters estimated from the data.
- It can reach past the worst month in the sample, but it has chosen a shape for the tails.
- On the right is a normal model, which gives you a bell shape whatever the data looked like, and so it has the strongest assumption about the tails.
- No point on the line is assumption free.
- You are only choosing which assumption you would rather make.

16 The trap: historical is not assumption free

- Historical simulation still assumes the past is typical
- No month worse than the worst on record
- Monte Carlo is only as good as the model chosen
- Describe each method by what it assumes

IN THE LESSON

- The classic exam trap is to say that historical simulation makes no assumptions.
- It does.
- It assumes that the future will look like the sampled past, that the sample is long enough to hold the bad events, and that each date can be redrawn independently of the others.
- If the sample missed a crisis, so will every simulated path.
- Monte Carlo is not safer for being fancier, either.
- A model that gets the tails wrong gives confident, precise, wrong numbers.
- When you describe either method, say what it assumes.
- Resampling assumes the future resembles the past, and Monte Carlo assumes its model is right.
- The sensible reply is to run both and compare.

17 A story: baking from a recipe

- A simulation is a recipe repeated thousands of times
- Decide the ingredients before pressing go
- Then judge the batch by how many cakes burn
- Not by the average cake

IN THE LESSON

- Say you are about to bake a cake from a recipe, thousands of times, to see how often it fails.
- Before you start you need to decide a few things.
- What cake are you making?
- What ingredients, and how much of each?
- How many cakes will you bake, and for how long do you leave each in the oven?
- A simulation is just the same.
- You must decide what to simulate, the distribution and its numbers, the number of trials and the horizon.
- Only then do you press go.
- And when the cakes come out, you do not judge the recipe by the average cake, you look at how many burned.
- Let us walk through those choices in order.

18 The choices a simulation needs

Inputs and choices

What to simulate → Distribution shape → Means, risks, correlations

Trials and horizon → Run and read the output

IN THE LESSON

- Here are the decisions in order.
- First, what to simulate: the returns of assets, of factors, or of a whole strategy.
- Second, the distribution: its shape, normal or fat tailed, and whether to allow regime shifts.
- Third, its parameters: the means, volatilities and correlations, estimated from data or set by judgement.
- Fourth, the number of trials and the horizon, that is how many paths to draw and how long each one runs.
- Then you run it.
- The last box is reading the output, and that is where the value is.
- Each earlier box is a place where an analyst's judgement enters, which is why a good report lists them all.

19 Example 4: how many trials?

Mean 6%, volatility 12%; error of the simulated mean

Error rule	volatility ÷ √trials	
100 trials	12 ÷ 10	**1.2%**
10,000 trials	12 ÷ 100	**0.12%**

IN THE LESSON

- Example four is about the number of trials.
- Heron simulates a portfolio with a mean of six per cent and a volatility of twelve per cent.
- Each trial is one random draw.
- The average of many trials is itself uncertain, with a standard error equal to the volatility divided by the square root of the trials.
- With one hundred trials, that is twelve over ten, or one point two per cent, which is a lot of noise.
- With ten thousand trials it is twelve over one hundred, which is zero point one two per cent.
- So a hundred times more trials shrinks the noise by a factor of ten.
- More trials do not make the model right.
- They only make the answer steadier.

20 Reading the output: VaR and conditional VaR

VaR (normal) = mean − z × volatility

Mean return	Average of the trials
z value	2.326 at the 1% level
Volatility	Spread of the trials
VaR	Loss exceeded in 1% of trials
Conditional VaR	Average loss beyond the VaR

IN THE LESSON

- The output of a simulation is not one number, it is a whole distribution of results.
- We summarise it with its mean, its percentiles, and two tail measures.
- Value at risk, or VaR, at one per cent is the loss that is exceeded in only one per cent of trials.
- Conditional VaR is the average loss in those worst one per cent of trials, so it tells you how bad things are once you are past the VaR line.
- Under a normal distribution, VaR is the mean minus the z value times the volatility, and the z value for one per cent is two point three two six.
- In a simulation you simply read these straight from the sorted trials.

21 Example 5: the normal run

Pension fund, 10,000 trials, mean 6%, volatility 12%

Mean return		**6.0%**
1% VaR	6 − 2.326 × 12	**loss 21.9%**
1% conditional VaR	average beyond VaR	**loss 26.0%**
Chance of losing over 25%		**0.5%**

IN THE LESSON

- Example five reads the normal run for Heron's pension fund, ten thousand trials, mean six per cent, volatility twelve per cent.
- The one per cent VaR is six minus two point three two six times twelve, which is minus twenty one point nine per cent, so a loss of twenty one point nine.
- The conditional VaR, the average loss beyond that, is twenty six point zero per cent.
- The chance of losing more than twenty five per cent is about half of one per cent.
- Notice the pattern: the mean is a pleasant six, and the interesting numbers are all in the tail.
- That is why the simulation was run in the first place.

22 Same mean, a fatter tail

Loss at the 1% level (%)

VaR normal	**21.9**
VaR fat	**25.8**
CVaR normal	**26.0**
CVaR fat	**38.3**

Same mean 6% and volatility 12% in both runs

IN THE LESSON

- Now repeat the run with a fat tailed distribution, a Student t with four degrees of freedom, scaled to the same volatility.
- The vertical axis shows the loss, in per cent.
- The first pair of bars is the one per cent VaR.
- Normal is twenty one point nine and fat tailed is twenty five point eight, so it moves by about four points.
- The second pair is the conditional VaR.
- Normal is twenty six point zero, fat tailed is thirty eight point three, a jump of twelve points.
- The mean was six per cent in both runs and never moved.
- The choice of distribution, made before a single trial, shows up almost entirely in the tail.

23 The trap: reading the mean as a forecast

- The simulated mean mostly echoes the input
- The value is in the spread and the tail percentiles
- Report VaR, conditional VaR and the chance of a miss
- Normal tails usually understate the loss

IN THE LESSON

- Here is the trap for this outcome.
- A trustee sees a simulation with an average return of six per cent and says, so we expect six per cent.
- That treats the simulation as a forecast.
- It is not.
- The average is the least interesting part, since it was largely put in by the analyst as an input.
- The reason to simulate is the spread: the percentiles, the VaR, the conditional VaR, and the probability of missing a target.
- In our example the chance of losing more than a quarter of the fund goes from about one year in two hundred to about one in ninety, depending on the distribution.
- That is the number a board can act on.

24 A story: testing a bridge

- Change one choice at a time, hold the rest fixed
- Typical dials: look-back, rebalancing, costs, distribution
- Does the conclusion survive reasonable changes?
- Robust results barely move; fragile ones collapse

IN THE LESSON

- Think about how engineers test a bridge.
- They do not only check it at the planned load.
- They ask what happens with heavier trucks, with stronger wind, or with a cheaper steel.
- If the bridge stands under every reasonable change, they trust it.
- If one small change brings it down, they worry.
- Sensitivity analysis is that habit applied to a backtest or a simulation.
- We change one choice at a time, such as the look back window, the rebalancing frequency, the transaction cost, the factor definition or the distribution, and watch whether the conclusion survives.
- Changing one thing at a time matters, because then you know which change did the damage.

25 Net return after costs

Net = gross − turnover × cost

Gross return	Before trading costs
Turnover	Share of fund traded a year
Cost per unit	Paid per unit traded
Net return	What the investor keeps
Break-even cost	Gross ÷ turnover

IN THE LESSON

- The most important sensitivity for many strategies is trading cost.
- Every time a strategy trades, it pays a cost.
- The yearly bill is turnover, the fraction of the portfolio traded in a year, times the cost per unit traded.
- So net return is gross return minus turnover times cost.
- A strategy with high turnover is therefore very sensitive to the cost assumption, because the cost is multiplied by a big number.
- We can also turn the formula round and ask at what cost the net return hits zero.
- That break even cost is the gross return divided by turnover, and it tells you how much room for error there is.

26 Example 6: doubling the cost

Momentum 3% gross, turnover 4; Value 2% gross, turnover 0.5

Momentum, cost 0.50%	3.0 − 4 × 0.50	**1.0%**
Momentum, cost 1.00%	3.0 − 4 × 1.00	**−1.0%**
Momentum break-even	3.0 ÷ 4	**0.75%**
Value, cost 0.50% / 1.00%	2.0 − 0.5 × cost	**1.75% / 1.50%**

IN THE LESSON

- Example six.
- Heron Momentum shows three per cent a year of gross alpha, with turnover of four hundred per cent, written as four, and an assumed cost of half a per cent.
- Net alpha is three minus four times zero point five, which is one per cent.
- Now double the cost to one per cent.
- Net alpha is three minus four, which is minus one per cent.
- The edge is gone.
- The break even cost is three over four, or zero point seven five per cent, only a quarter of a point of headroom.
- The Heron Value tilt has two per cent gross and turnover of half.
- At the base cost it nets one point seven five, and at double cost one point five.
- It barely notices.

27 Net alpha as the cost assumption rises

Chart: Net alpha (%) vs Cost per unit traded (%)

- Momentum
- Value tilt
- Zero

IN THE LESSON

- Here are those results drawn.
- The horizontal axis is the cost assumed per unit traded, from nothing up to one per cent.
- The vertical axis is net alpha.
- The flat grey line is zero.
- The steep red line is momentum.
- It starts at three and falls four points for every one point of cost, and it crosses zero at zero point seven five.
- The gentle green line is the value tilt.
- It starts at two and falls only half a point for each point of cost, so it stays comfortably above zero.
- The slope of each line is its turnover.
- A picture like this shows at a glance how wrong the cost assumption can be before the conclusion changes.

28 Look-back window: spreading the dials

Trend rule, Sharpe ratio by look-back window

Look-back window	Sharpe ratio
6 months	0.30
12 months	0.62
18 months	0.31
24 months	0.28

IN THE LESSON

- Costs are not the only dial.
- Suppose Heron tests a trend rule with different look back windows, and records the Sharpe ratio for each.
- Six months gives zero point three zero.
- Twelve months gives zero point six two.
- Eighteen months gives zero point three one.
- Twenty four months gives zero point two eight.
- The twelve month result stands alone like a spike, and its neighbours are about half as good.
- The average over the four windows is only about zero point three eight.
- If a report showed just the twelve month column, it would be accurate and still misleading.
- A result that depends on one lucky setting is a warning of overfitting.

29 Robust or overfitted? The final trap

- Robust: similar results across reasonable alternatives
- Fragile: results collapse after a small change
- Show the whole grid, not the best cell
- Prefer a broad plateau to a single peak

IN THE LESSON

- Let us finish with how to judge the results.
- A robust strategy performs similarly across reasonable alternatives: nearby windows, higher costs, a fat tailed distribution.
- A strategy whose results collapse with small changes is probably overfitted, meaning it has learned the noise of one sample.
- The trap is to report only the most favourable parameter set.
- A careful analyst shows the whole grid, says how many versions were tried, and prefers a broad plateau of good results to a single peak.
- In simulation the same test applies: swap the normal for a fat tailed distribution and see how much of the tail risk was just an effect of the model.

30 Recap: what to remember

- Scenarios: regime by regime, few episodes, be modest
- Historical resamples the past; Monte Carlo uses a model
- Read the tail, not the simulated mean
- Sensitivity: change one thing at a time, show every result

IN THE LESSON

- Let us pull it together.
- Historical scenario analysis splits the past into regimes, and shows whether a good average hides a bad regime, and whether diversification fails in stress.
- Only a few episodes sit behind each regime, so be modest.
- Historical simulation resamples real dates, keeping real tails and correlations but never producing anything new.
- Monte Carlo draws from a chosen model, which can create new outcomes but is only as good as the model.
- A simulation needs choices of what, which distribution, which parameters, how many trials and what horizon, and the output to read is the tail.
- Sensitivity analysis varies one input at a time, and robust results survive it.

31 Try three quick questions

1. 200 months at +1%, 50 at −2%. Full-sample mean?
   Answer: (200 × 1 − 50 × 2) ÷ 250 = 0.4% a month.
2. Which approach cannot produce a month worse than its data's worst?
   Answer: Historical simulation (the bootstrap).
3. Gross alpha 2%, turnover 5. Break-even cost?
   Answer: 2 ÷ 5 = 0.4% per unit traded.

IN THE LESSON

- Now try three quick questions.
- Pause after each one.
- Question one.
- A strategy has two hundred months in expansion averaging one per cent and fifty in recession averaging minus two per cent.
- What is the full sample mean?
- Two hundred times one, minus fifty times two, over two hundred and fifty, which is zero point four per cent.
- Question two.
- Which simulation approach can never produce a month worse than the worst in its data?
- Historical simulation, the bootstrap.
- Question three.
- A strategy has gross alpha of two per cent and turnover of five.
- What cost makes its net alpha zero?
- Two divided by five, which is zero point four per cent.
- Well done, you are ready for the practice questions.