Stress tests, liquidity and risk budgets

Asset Allocation

1 A story: the wind tunnel

- Engineers test a bridge in a wind tunnel
- Test the plan against storms, not calm days
- Simulation: thousands of random futures
- Scenarios: a few named storms, chosen by hand

IN THE LESSON

- Before a bridge is built, engineers put a model of it in a wind tunnel.
- They do not ask how the bridge behaves on a calm afternoon.
- They ask what happens in the worst storm they can imagine.
- An asset allocation deserves the same treatment.
- The optimiser gives Priya Nair one answer, built on averages for a single period.
- But Calder Foundation lives through many years in a row, paying out as it goes.
- So she tests the plan in two ways.
- Simulation throws thousands of random futures at it.
- Scenario analysis applies a handful of named storms, chosen by hand, and watches what breaks.

2 Same returns, different order

Chart: Wealth (start 1,000) vs Years (withdraw 60 a year)

- Average return 4.8% either way
- Losses first: ends 620
- Losses last: ends 876

IN THE LESSON

- Here is why a single period average is not enough.
- The horizontal axis is years, nought to ten, and the retiree takes sixty out each year.
- The vertical axis is wealth, starting at one thousand.
- Both lines use exactly the same ten annual returns, which average four point eight per cent a year, including two bad years at the start of one sequence.
- Follow the red line first.
- The losses come early, the withdrawals come out of a smaller pot, and the pot never recovers.
- It ends at six hundred and twenty.
- The green line has the same returns in reverse order, so the bad years come last.
- It ends at eight hundred and seventy six.
- Only the order changed.

3 Example 1: order matters with cash flows

Start 1,000; take out 100 a year. Returns −20%, +10%, +10%

No withdrawals, either order	1,000 × 0.8 × 1.1 × 1.1	**968**
Loss first, with withdrawals	−20%, +10%, +10%	**640**
Loss last, with withdrawals	+10%, +10%, −20%	**703**
Gap caused by order	703 − 640	**63**

IN THE LESSON

- Let us see the same effect with simple numbers.
- A retiree starts with one thousand and takes out one hundred at the start of each year.
- In the first path the returns are minus twenty, plus ten, plus ten per cent.
- In the second path the same three returns come in the opposite order.
- With no withdrawals, both paths end at nine hundred and sixty eight, because multiplying numbers in a different order gives the same product.
- Now add the withdrawals.
- The path with the loss first ends at six hundred and forty.
- The path with the loss last ends at seven hundred and three.
- The gap of sixty three comes entirely from the order, and a single period average cannot see it.

4 How a Monte Carlo simulation works

From assumptions to an answer

Assume returns → Draw many paths → Add cash flows → Count successes

IN THE LESSON

- So how does simulation actually work?
- Think of it as playing out the foundation's future again and again.
- The first box is the assumptions: how each asset class might return, how risky it is, how the classes move together.
- These can be drawn from the normal bell curve, or from shapes with fatter tails if that is more realistic.
- The second box draws a path of returns, year by year.
- The third box applies the real life rules to that path: the spending, the rebalancing, the taxes and the fees.
- Then we repeat thousands of times.
- The last box counts what happened, the spread of ending wealth and the share of paths that met the goal.

5 Example 2: Calder's simulation

Mix X, 60% shares; Calder 200m; spend 4%; goal: 200m at year 20

Base assumptions	8,200 of 10,000 paths	**82%**
Returns one point lower	6,900 of 10,000	**69%**
Fatter-tailed returns	7,600 of 10,000	**76%**
Swing from the inputs	82 − 69	**13 points**

IN THE LESSON

- Now Calder Foundation.
- Priya Nair simulates ten thousand paths for a mix with sixty per cent in shares.
- The foundation holds two hundred million, spends four per cent a year, and the goal is to still have two hundred million at the end of twenty years.
- With her base assumptions, eight thousand two hundred paths succeed, which is eighty two per cent.
- Now she lowers the expected return by one point.
- Only six thousand nine hundred paths succeed, sixty nine per cent.
- Then she keeps the average return but gives the returns fatter tails, and the figure is seventy six per cent.
- Notice the swing.
- A small change in one input moved the answer thirteen points.
- The output is only as sound as the inputs.

6 Scenarios and simulation side by side

Which tool answers which question

Question	Best tool
How likely is the goal met?	Simulation
Does it handle fat tails?	Simulation
What if stagflation hits?	Scenario
What if markets freeze?	Scenario

IN THE LESSON

- The two tools answer different questions, so a careful adviser uses both.
- Simulation answers how likely the goal is.
- It handles fat tails and long paths, and it gives a probability.
- Scenario analysis answers a different question: what happens if this particular storm hits?
- The storm might be a repeat of a past episode or a hypothetical one, like a long stretch of stagflation, or a liquidity crisis in which markets freeze and nothing can be sold at a fair price.
- Scenarios are not forecasts.
- They are chosen to hurt, so the investor can see in advance how the allocation and the spending promise respond.
- Here is a quick guide to which tool fits which question.

7 Example 3: a stagflation scenario

Shares 60%, bonds 30%, real assets 10%; Calder 200m

Shares	60% × −20%	**−12.0%**
Bonds	30% × −8%	**−2.4%**
Real assets	10% × +6%	**+0.6%**
Portfolio return	−12.0 − 2.4 + 0.6	**−13.8%**
Assets after spending 8m	200m × 0.862 − 8m	**164.4m**

IN THE LESSON

- Here is a scenario for Calder.
- Stagflation means weak growth and high inflation together.
- Priya Nair assumes shares fall twenty per cent, bonds fall eight per cent and real assets rise six per cent.
- The mix holds sixty per cent shares, thirty per cent bonds and ten per cent real assets.
- Shares contribute minus twelve points, bonds minus two point four, and real assets add point six.
- The portfolio loses thirteen point eight per cent.
- But the foundation must still pay its spending of eight million, which is four per cent of the starting two hundred million.
- So the two hundred million becomes about one hundred and sixty four point four million.
- The trustees can now ask whether they could live with that.

8 The trap: trusting the output

- Output is only as good as the inputs
- Change an assumption and the odds change
- Report a range, not one probability
- Ask which storms the model left out

IN THE LESSON

- Here is the trap.
- A simulation prints a neat probability, such as eighty two per cent, and it looks like a fact.
- It is not.
- Every number it reports was produced by the assumptions put in: the average returns, the risk, the correlations, the shape of the distribution.
- We saw one point of return move the answer thirteen points.
- A scenario has the opposite weakness.
- It shows one storm, so it can say nothing about the storms nobody thought of.
- The sound habit is to rerun the simulation with different inputs, report a range rather than a single figure, and ask which shocks the model has quietly left out.

9 A story: the house you cannot sell by Friday

- A house is worth what the valuer says
- The true price shows only on sale day
- Private assets are priced by appraisal
- Smooth prices make them look safer than they are

IN THE LESSON

- Suppose you own a house, and a valuer says it is worth three hundred thousand.
- Is that its price?
- Not really.
- It is an opinion.
- You will only find the true price on the day you sell, and that might take months and cost a fee.
- Private equity, private credit, real estate and infrastructure are like that house.
- They are hard to trade, and their returns are built from appraisals instead of daily trades.
- Appraisals move slowly, so the reported ups and downs look gentle.
- Gentle numbers make the asset seem safer than it is, and an optimiser, which believes the data, will fall in love with them.
- That is the danger this section is about.

10 From easy to sell to hard to sell

- Cash
- Listed shares
- Private credit
- Real estate
- Private equity

IN THE LESSON

- Let us place the asset classes on one line.
- On the far left is the easy end.
- Cash can be spent at once, and listed shares can be sold in seconds at a visible price.
- Moving right, private credit is lent out for years and cannot be sold quickly.
- Real estate takes months to sell and is valued by appraisal.
- At the far right sit private equity and infrastructure, where money is locked in funds for many years.
- The more illiquid the class, the fewer investable indices there are, and the more the investor's result depends on which manager it picks, rather than on the asset class itself.
- Priya Nair must keep this line in mind whenever an optimiser praises something on the right.

11 Undoing the smoothing

Unsmoothed vol = reported vol × √((1 + φ) ÷ (1 − φ))
Sharpe = (E(R) − Rf) ÷ volatility

Reported volatility	Taken from appraisals
Smoothing weight	Share of old value carried over
Unsmoothed volatility	Always higher than reported
Sharpe ratio	Return per unit of risk

IN THE LESSON

- Appraisal smoothing can be partly undone.
- Suppose each reported return is a blend, with a share called the smoothing weight coming from last period's appraisal.
- Then the reported volatility is too low, and the true figure is the reported volatility times the square root of one plus the weight, over one minus the weight.
- This assumes the true returns are unrelated from one period to the next.
- The heavier the smoothing, the bigger the correction.
- Then we compare assets using the Sharpe ratio you already know: the expected return less the risk free rate, divided by volatility.
- Correcting the volatility lowers the Sharpe ratio, and the real comparison starts from there.

12 Example 4: private equity looks too good

Private equity reports 11%, vol 8%, φ 0.7; Rf 3%

Unsmoothing factor	√(1.7 ÷ 0.3)	**2.38**
Unsmoothed volatility	8% × 2.38	**19.0%**
Sharpe, as reported	(11 − 3) ÷ 8	**1.00**
Sharpe, unsmoothed	(11 − 3) ÷ 19.0	**0.42**
Sharpe, listed equities	(8 − 3) ÷ 16	**0.31**

IN THE LESSON

- Priya Nair is considering private equity for Calder.
- The fund reports eleven per cent a year with volatility of eight per cent, and its smoothing weight is point seven.
- The risk free rate is three per cent.
- As reported, the Sharpe ratio is eleven minus three, over eight, which is one point zero.
- That beats everything.
- Now unsmooth.
- One point seven over point three, square rooted, is two point three eight.
- Eight times that gives true volatility of about nineteen per cent.
- The corrected Sharpe ratio is eight over nineteen, which is point four two.
- Listed equities offer eight per cent with sixteen per cent volatility, so their Sharpe ratio is five over sixteen, point three one.
- The edge was mostly an illusion.

13 Reported, then corrected

Sharpe ratio

Reported	**1.00**
Unsmoothed	**0.42**
Listed shares	**0.31**

First two bars are private equity

IN THE LESSON

- Here are those three Sharpe ratios as bars.
- The vertical axis is the Sharpe ratio, return per unit of risk, and the first two bars are private equity.
- The tall first bar is what the optimiser sees if it is given the reported data: one point zero, the best on offer, so it would pile in.
- The second bar corrects the volatility for smoothing, and the ratio drops by more than half, to point four two.
- Private equity may still beat listed shares, and a real illiquidity premium could explain some of that.
- The third bar is listed equities at point three one, which Calder could hold instead.
- The fair comparison is between the second and third bars, not the first and third.

14 Ways to handle illiquid classes

Four options an adviser can use

Option	Idea
Leave out of the optimiser	Set weight apart
Use a listed stand-in	Adjust inputs
Unsmooth the returns	Raise volatility
Add a liquidity limit	Cap the weight

IN THE LESSON

- So what can an adviser do about illiquid classes?
- There are four common options.
- First, leave them out of the optimiser altogether and set their weight separately, by judgement, because the data cannot be trusted.
- Second, use a listed stand in, such as listed real estate or small company shares, with inputs adjusted to reflect the true risk.
- Third, unsmooth the returns, as we just did, which raises volatility and correlations before they go into the optimiser.
- Fourth, add an explicit liquidity limit, a cap on how much of the portfolio may sit in assets that cannot be sold quickly.
- Each option accepts that the raw data is unreliable, and each puts a human check on the optimiser.

15 Example 5: can Calder afford it?

Calder 200m; fund minimum 10m; 5 managers, 3 start years

Commitments needed	5 managers × 3 years	**15**
Money needed	15 × 10m	**150m**
Share of Calder	150m ÷ 200m	**75%**
Share of a 5,000m endowment	150m ÷ 5,000m	**3%**

IN THE LESSON

- Last, the investor matters.
- A trustee admires a big university endowment that spreads private equity across many managers.
- Could Calder copy it?
- Each fund asks for a minimum commitment of ten million.
- To spread risk, Priya Nair wants five managers, and commitments in three different start years, because the vintage of a fund matters.
- Five times three is fifteen commitments.
- At ten million each that is one hundred and fifty million.
- Against Calder's two hundred million, that is seventy five per cent of everything it has.
- For a five thousand million endowment, the same programme is only three per cent.
- The trap is letting an optimiser overallocate to private assets on smoothed data.
- For a small investor, a listed stand in or a small capped holding is the realistic route.

16 A story: the household budget

- Budget the worry, not only the money
- Cap the total risk, then share it out
- Capital weights: where the money sits
- Risk weights: where the risk comes from

IN THE LESSON

- A sensible family budgets its money.
- Rent gets this much, food gets that much, and the total is capped by what comes in.
- A risk budget does the same thing for worry.
- The investor decides how much total risk it is willing to carry, then shares that risk out across asset classes or managers, and checks that each slice earns its keep.
- The key idea is that the money and the risk are not shared out the same way.
- A small slice of capital in a wild asset can use up most of the risk.
- So a risk budget looks at where the risk comes from, not just where the money sits.

17 Absolute and relative risk budgets

Same idea, two ways to set the limit

Budget	Measured as
Absolute budget	Total volatility
Relative budget	Tracking error
Calder's absolute budget	12% volatility
Calder's relative budget	2% tracking error

IN THE LESSON

- There are two ways to set the limit.
- An absolute risk budget caps the total risk of the portfolio, usually as volatility, or as a largest acceptable loss.
- It answers how much risk the foundation can carry in all.
- A relative risk budget caps how far the portfolio may stray from a benchmark, measured as tracking error, which is the volatility of the difference between the portfolio and the benchmark.
- It limits the active risk that managers are allowed to take.
- Calder's trustees set both.
- The total volatility must stay within twelve per cent, and the active management must keep tracking error within two per cent of the policy benchmark.

18 Contribution to risk

Marginal contribution = β × σp
Contribution = weight × marginal contribution

Weight	Share of the money
Beta to portfolio	Moves with the whole portfolio
Portfolio volatility	Total risk of the mix
Contribution to risk	These add up to σp

IN THE LESSON

- Now the arithmetic of splitting risk.
- For each asset we need its marginal contribution to risk: how much the portfolio's volatility would change, per unit, if we added a little of that asset.
- It equals the asset's beta to the portfolio, times the portfolio's volatility.
- Beta here measures how closely the asset moves with the whole portfolio.
- Then the absolute contribution is the asset's weight times its marginal contribution.
- Add up the contributions of every asset and you get exactly the portfolio volatility, which is what makes them usable as a budget.
- Dividing each contribution by the total gives the asset's share of the risk.

19 Example 6: where Calder's risk comes from

Shares 60% (vol 16%), bonds 40% (vol 5%), correlation 0.2

Portfolio volatility	√(92.16 + 4 + 7.68)	**10.19%**
Share beta	160 ÷ 103.84	**1.54**
Share contribution	0.6 × 1.54 × 10.19	**9.42**
Bond contribution	0.4 × 0.19 × 10.19	**0.77**
Shares' share of total risk	9.42 ÷ 10.19	**92%**

IN THE LESSON

- Take Calder's policy mix: sixty per cent shares with volatility of sixteen per cent, and forty per cent bonds with volatility of five per cent, and a correlation of point two between them.
- The portfolio variance is ninety two point one six, plus four, plus seven point six eight, a total of one hundred and three point eight four.
- The square root, the portfolio volatility, is ten point one nine per cent.
- The shares' beta to the portfolio is one hundred and sixty over one hundred and three point eight four, which is one point five four.
- So the shares contribute point six times one point five four times ten point one nine, about nine point four two.
- The bonds contribute point seven seven.
- Together that is ten point one nine.
- Shares supply ninety two per cent of the risk.

20 Capital weights against risk weights

Share of total (%)

Shares: money	**60%**
Shares: risk	**92%**
Bonds: money	**40%**
Bonds: risk	**8%**

Same portfolio, two views

IN THE LESSON

- Here is that result as a picture.
- The vertical axis is the share of the total, in per cent.
- Look at the first pair of bars, for shares.
- The first is the share of the money, sixty per cent.
- The second is the share of the risk, ninety two per cent.
- Now the bonds.
- They hold forty per cent of the money, but only eight per cent of the risk.
- So a sixty forty portfolio is, in risk terms, almost entirely a share portfolio.
- This is exactly the trap the exam sets: treating capital weights as if they were risk weights.
- A risk budget works from the second pair of bars, and asks whether the risk being carried is the risk the investor actually wants.

21 The efficient risk budget

(E(Ri) − Rf) ÷ marginal contribution i = same for every asset

Excess return	Return above the risk-free rate
Marginal contribution	Beta × portfolio volatility
Return per unit of risk	Equal at the optimum

IN THE LESSON

- How do we know a risk budget is well spent?
- Use the same idea as a shopper comparing prices.
- For each asset, divide its expected excess return, the return above the risk free rate, by its marginal contribution to risk.
- That is the return you earn per unit of risk the asset adds.
- The allocation is optimal when this ratio is the same for every asset.
- If one asset earns more per unit of risk than another, move some risk budget toward it, and away from the weaker one, until the ratios meet.
- The ratio does not depend on how much money is in each asset, which is why capital weights alone cannot tell you whether the budget is well spent.

22 Example 7: is the budget well spent?

Excess returns: shares 5%, bonds 1.5%

Share marginal contribution	1.54 × 10.19	**15.70**
Bond marginal contribution	0.19 × 10.19	**1.92**
Shares: return per risk	5 ÷ 15.70	**0.32**
Bonds: return per risk	1.5 ÷ 1.92	**0.78**
Verdict	bonds earn more per unit	**Shift to bonds**

IN THE LESSON

- Back to Calder's mix.
- The expected excess return is five per cent for shares and one and a half per cent for bonds.
- First the marginal contributions.
- For shares it is one point five four times ten point one nine, about fifteen point seven.
- For bonds the beta is point one nine, so the marginal contribution is about one point nine two.
- Now the ratios.
- Shares earn five over fifteen point seven, point three two, per unit of risk.
- Bonds earn one point five over one point nine two, point seven eight.
- Bonds earn more than twice as much per unit of risk.
- So this allocation is not optimal for these inputs.
- Priya Nair should shift weight from shares toward bonds, and the two ratios will move closer together.

23 Example 8: a relative budget

Tracking-error budget 2%; two active managers

Manager P tracking error		**1.2%**
Manager Q tracking error		**1.6%**
Combined, unrelated	√(1.44 + 2.56)	**2.00%**
Combined, correlation 0.6	√(4.00 + 2.304)	**2.51%**
Overspent against budget	2.51 − 2.00	**0.51 points**

IN THE LESSON

- Now the relative budget.
- Calder's trustees allow two per cent tracking error.
- They hire two active managers.
- Manager P uses one point two per cent of it, and Manager Q uses one point six per cent.
- If their active bets are unrelated, the risks add in squares, not in a straight line.
- One point two squared is one point four four, one point six squared is two point five six, together four.
- The square root is two point zero, exactly the budget.
- But suppose both managers lean on the same value tilt, so their active returns have a correlation of point six.
- Then add two times point six times one point two times one point six, which is two point three.
- The total is the square root of six point three, two point five one.
- The budget is overspent.

24 Mistakes to avoid so far

- Trusting simulated odds without questioning inputs
- Treating a scenario as a forecast
- Letting smoothed private returns look safe
- Equating capital weights with risk weights
- Ignoring correlation between active managers

IN THE LESSON

- Let us collect the common mistakes from the first three topics.
- One: trusting a simulated probability without asking what assumptions produced it.
- Two: treating a scenario as a forecast, when it is only a chosen storm.
- Three: feeding smoothed private asset returns into an optimiser, which makes them look safe and encourages it to overallocate.
- Four: equating capital weights with risk weights, so that shares at sixty per cent of the money are missed as ninety two per cent of the risk.
- Five: adding the tracking errors of managers as if they were unrelated, when similar styles make them move together and the budget is used up faster.

25 A story: clients speak in worries

- Clients say: I cannot lose more than this
- Or: I must earn enough to keep spending
- Or: no tobacco, and sell my donor's shares
- Each worry becomes an objective or a constraint

IN THE LESSON

- Few clients describe their fears as variance.
- A client says, I could not stand to lose more than this in a year.
- Or, I must earn enough to keep paying my grants.
- Or, my wealth is tied up in one company, and I do not want to own tobacco.
- Priya Nair's job is to turn those sentences into the language of the optimiser.
- A risk aversion number goes into the utility function.
- A minimum return becomes a shortfall or safety first rule.
- A tolerated loss becomes a limit on value at risk.
- A worry about one exposure becomes a hedge, a cap or an exclusion.
- Each of these shapes the answer, and each has a price.

26 Turning worries into numbers

$U = E(R) - 0.005 \times \lambda \times \sigma^2$
Loss, 1 year in 20 = $E(R) - 1.65 \times \sigma$
Safety-first = $(E(R) - R_{min}) \div \sigma$

Risk aversion	Higher means more cautious
Expected return	Average return expected
Volatility	Risk of the mix
Minimum return	The floor the client needs

IN THE LESSON

- Here are three ways to write a worry as a number.
- First, the utility function.
- Expected return minus point zero zero five times lambda times variance, where lambda is the client's risk aversion.
- A higher lambda penalises risk more heavily.
- Second, a loss limit.
- Under a normal approximation, the loss that is exceeded only one year in twenty is the expected return minus one point six five times volatility.
- That is value at risk.
- Third, the safety first ratio.
- It is the expected return minus the client's minimum return, divided by volatility.
- The mix with the highest ratio has the lowest chance of falling short of the minimum.
- Each is a different way to say what matters.

27 Example 9: three candidate mixes

$\lambda = 1.5$; limit 12%, one year in twenty; C (6%, vol 9) is on the next slide

Utility of A (8%, vol 15%)	$8 - 0.0075 \times 225$	**6.31**
Utility of B (7%, vol 11%)	$7 - 0.0075 \times 121$	**6.09**
Bad-year result, A	$8 - 1.65 \times 15$	**−16.75%**
Bad-year result, B	$7 - 1.65 \times 11$	**−11.15%**
Verdict	the limit is 12%	**Choose B**

IN THE LESSON

- Priya Nair has three efficient mixes.
- A offers eight per cent with fifteen per cent volatility.
- B offers seven per cent with eleven.
- C offers six per cent with nine, and we meet it on the next slide.
- Calder's risk aversion is one point five.
- Utility for A is eight minus point zero zero seven five times two hundred and twenty five, which is six point three one.
- For B it is seven minus point zero zero seven five times one hundred and twenty one, six point zero nine.
- The utility function alone picks A.
- But the trustees also said they could not accept losing more than twelve per cent in a year, one year in twenty.
- For A the loss is eight minus one point six five times fifteen, minus sixteen point seven five.
- For B it is minus eleven point one five.
- A breaks the limit and B does not.

28 The worst year against the limit

One-in-twenty loss (%)

A: return 8%	**16.75%**
B: return 7%	**11.15%**
C: return 6%	**8.85%**

Client limit: 12%

IN THE LESSON

- Here is the same choice as a picture.
- The vertical axis is the loss that would be exceeded only one year in twenty, in per cent.
- Each bar is one mix, with its expected return in the label.
- A, with the highest return of eight per cent, has the tallest bar, sixteen point seven five.
- That is well beyond the twelve per cent the trustees said they could bear.
- B, with a return of seven per cent, has a loss of eleven point one five, inside the limit.
- C, with a return of six per cent, has a loss of eight point eight five.
- It is safer still, but it gives up more return than the limit requires.
- So the statement about losses, not the utility score, decides the choice, and B is the answer.

29 Worries and how they are written in

Calder's worries, and how Priya Nair writes them in

Client's worry	Written in as
Cannot lose over 12%	Loss limit
Needs 4% to spend	Safety-first floor
Donor shares, 20%	Hedge or cap
No tobacco	Exclusion

IN THE LESSON

- This table lists the common worries and the usual way to write each into an allocation.
- A stated largest loss becomes a value at risk or drawdown limit.
- A need to earn enough to keep spending becomes a safety first floor or a shortfall constraint.
- A concentrated position, such as shares in the donor's company, is hedged, sold down, or capped.
- An ethical preference, like no tobacco, becomes an exclusion that narrows the opportunity set a little.
- Beyond the table, a client who hates losses more than they love gains may be given goals based buckets, so each pot has a clear purpose.
- Limits can also cap leverage, shorting or currency exposure.

30 Example 10: what the preferences cost

Compare the client's choice with the free optimum

Return given up, A to B	8 − 7	**1.0 point**
Utility given up	6.31 − 6.09	**0.22**
Donor stake in one company	20% × 200m	**40m**
Verdict	a price the client chose	**Keep limit**

IN THE LESSON

- Every preference has a cost, and an honest adviser states it.
- Moving from A to B gives up one point of expected return, eight minus seven.
- In utility terms the loss is six point three one minus six point zero nine, point two two.
- That is the price of the twelve per cent loss limit.
- Calder also holds shares in the donor's company, twenty per cent of its two hundred million, which is forty million tied to one business.
- Hedging or selling that stake removes a risk the optimiser never saw.
- The trap is to ignore a stated preference because the optimiser's answer looks better on paper.
- If the client said it, it is a constraint, and the cost is a price the client chose to pay.

31 Recap: the lesson in five lines

- Simulation and scenarios test paths and shocks
- Illiquid classes: unsmooth, limit, mind your size
- Risk budget: weight × beta × portfolio volatility
- Best mix: equal return per unit of marginal risk
- Client worries become limits, and each costs return

IN THE LESSON

- Let us pull the lesson together.
- First, simulation draws thousands of paths and reports the chance of meeting a goal, while scenarios apply named storms, and both depend on the inputs.
- Second, illiquid assets have smoothed returns, so unsmooth them, limit them, or use listed stand ins, and remember that a small investor cannot diversify across managers.
- Third, a risk budget caps total risk or tracking error, and contribution equals weight times beta times portfolio volatility, with the best mix earning equal return per unit of marginal risk.
- Fourth, client worries become limits and exclusions.
- Fifth, every constraint costs return, and that cost belongs to the client.

32 Try three quick questions

1. Reported vol 6%, smoothing weight 0.6. Unsmoothed vol?
   Answer: 6 × √(1.6 ÷ 0.4) = 6 × 2 = 12%.
2. Weight 50%, beta 1.2, portfolio vol 10%. Risk contribution?
   Answer: 0.5 × 1.2 × 10 = 6.0, which is 60% of the risk.
3. Mix: 9%, vol 14%. Limit: lose 10%, one year in 20. Allowed?
   Answer: 9 − 1.65 × 14 = −14.1%. No, it breaks the limit.

IN THE LESSON

- Now try three quick questions yourself.
- Pause after each one.
- Question one.
- A private fund reports volatility of six per cent, with a smoothing weight of point six.
- What is the unsmoothed volatility?
- One point six over point four is four, the square root is two, so six times two gives twelve per cent.
- Question two.
- An asset has a weight of fifty per cent, a beta of one point two to the portfolio, and the portfolio volatility is ten per cent.
- What does it contribute to risk?
- Point five times one point two times ten is six, which is sixty per cent of the total risk.
- Question three.
- A mix offers nine per cent with fourteen per cent volatility, and the client will not lose more than ten per cent, one year in twenty.
- Nine minus one point six five times fourteen is minus fourteen point one per cent, so the mix breaks the limit and the answer is no.