Mean-variance optimisation and the client's whole picture

Asset Allocation

1 A story: the menu and the scorer

- The frontier is a menu of best mixes
- Each risk level: the highest expected return
- A score matches one mix to the client
- Inputs: returns, volatilities, correlations

IN THE LESSON

- Imagine a restaurant that lists only its best dish at each price.
- Nothing on the list is both cheaper and better than another dish, so every dish is the best you can get for what you pay.
- That list is the efficient frontier: for each level of risk, the mix with the highest expected return.
- Picking from the list is then a matter of taste.
- A diner who hates spice chooses a mild dish, and a daring one chooses a hot dish.
- Mean variance optimisation does two jobs.
- It draws the menu from expected returns, volatilities and correlations.
- Then it scores each dish against how much the client dislikes risk, and the top score wins.
- Meet Ivor and his adviser Nell.

2 The frontier and Ivor's pick

Chart: Expected return (%) vs Volatility (%)

- A (6, 4.8)
- B (9, 6.2)
- C (12, 7.4)
- D (16, 8.4)

IN THE LESSON

- Here is Nell's menu, drawn.
- The horizontal axis is risk, measured as volatility in per cent.
- The vertical axis is expected return in per cent.
- The grey rising curve is the efficient frontier, and it flattens because each extra step of risk buys a little less extra return.
- Nell finds four mixes on it.
- Mix A sits at volatility six and return four point eight.
- Mix B sits at nine and six point two.
- Mix C sits at twelve and seven point four.
- Mix D sits at sixteen and eight point four.
- Now the red curve.
- It joins every point that Ivor finds equally good, and it just touches the frontier at mix B.
- Everything below the red curve is worse for him, so B is his pick.

3 The utility score

$U = E(R) - 0.005 \times \lambda \times \sigma^2$

Expected return	Return the mix should earn
Risk aversion	How much risk is disliked
Variance	Volatility squared, in percent

IN THE LESSON

- Now the scorer.
- Utility, written U, equals expected return, minus zero point zero zero five times lambda times variance.
- Lambda is the risk aversion coefficient: the higher it is, the more the client dislikes risk.
- Variance is volatility squared.
- Returns and volatilities are both in per cent, which is why the odd little number, zero point zero zero five, sits in front.
- Read the formula as a score: the return you hope for, less a penalty for the risk you carry.
- The penalty grows with the square of volatility, so doubling the risk quadruples the penalty.
- A mix that earns more can still score lower if its penalty is bigger.

4 Example 1: scoring four mixes

Ivor: $\lambda = 4$, so $U = E(R) - 0.02 \times \sigma^2$

Mix A	$4.8 - 0.02 \times 36$	**4.08**
Mix B	$6.2 - 0.02 \times 81$	**4.58**
Mix C	$7.4 - 0.02 \times 144$	**4.52**
Mix D	$8.4 - 0.02 \times 256$	**3.28**

IN THE LESSON

- Let us score the four mixes.
- Nell estimates Ivor's risk aversion at four, so the penalty is zero point zero two times the variance.
- Mix A: four point eight, minus zero point zero two times thirty six, gives four point zero eight.
- Mix B: six point two, minus zero point zero two times eighty one, gives four point five eight.
- Mix C: seven point four, minus zero point zero two times one hundred forty four, gives four point five two.
- Mix D: eight point four, minus zero point zero two times two hundred fifty six, gives three point two eight.
- Mix D has the highest return but the lowest score, because its risk penalty swallows the extra return.
- Mix B scores highest, with C only just behind.
- Remember that close call, because we come back to it.

5 Same menu, different clients

Same four mixes, three different clients

Risk aversion (λ)	Best mix, score
λ = 2, bold	C, 5.96
λ = 4, Ivor	B, 4.58
λ = 8, cautious	A, 3.36

IN THE LESSON

- Now see why the score matters.
- Change only the risk aversion and the pick changes, on the very same menu.
- With lambda of two, the penalty is light, and mix C scores five point nine six, the best.
- With lambda of four, which is Ivor, mix B scores four point five eight.
- With lambda of eight, a cautious investor, the penalty is heavy, and mix A scores three point three six, the best.
- The optimiser does not know the right answer.
- It needs the client's risk aversion as an input, and estimating that number is judgement, not arithmetic.
- A different adviser with a different guess of lambda could defend a different mix.

6 Example 2: forecasts that wobble

Markets X, Y: vol 15% and 16%, corr 0.9, λ = 4

Bottom line	$15^2 + 16^2 - 2\times0.9\times15\times16$	**49**
Equal forecasts	(256 − 216) ÷ 49	**81.6%**
X up 0.5	(12.5 + 40) ÷ 49 = 107%	**100%**
Y up 0.5	(−12.5 + 40) ÷ 49	**56.1%**

IN THE LESSON

- Now the weakness.
- Suppose the menu includes two equity markets that are close cousins, market X and market Y.
- Their volatilities are fifteen and sixteen, their correlation is point nine, and risk aversion is four.
- The weight in X is a top line divided by a bottom line.
- The bottom line is fifteen squared plus sixteen squared, minus two times point nine times fifteen times sixteen, which is forty nine.
- With equal forecasts the top line is two hundred fifty six minus two hundred sixteen, forty, so X gets eighty one point six per cent.
- Raise X's forecast by half a point, and the top gains twelve point five.
- The weight passes one hundred per cent and is capped.
- Raise Y's by half a point instead, and X falls to fifty six point one per cent.

7 Half a point swings the weight

Weight in market X (%)

Equal forecasts	**82%**
X up 0.5	**100%**
Y up 0.5	**56%**

Only half a point of forecast moved

IN THE LESSON

- Here are those three answers as bars.
- The vertical axis is the optimal weight in market X, in per cent.
- The first bar is equal forecasts, with X at eighty two per cent.
- The second bar is X's forecast raised by half a point, and the optimiser puts everything in X, one hundred per cent: a corner solution, with nothing at all in Y.
- The third bar is Y's forecast raised by half a point, and X drops to fifty six per cent.
- A half point is well inside the usual error of a long run return forecast.
- Yet it moves the weight across almost the whole range.
- The optimiser is precise about inputs that are anything but precise, and that is the trap.

8 Strengths and weaknesses

- Strength: clear link from inputs to weights
- Weak: small forecast errors give corner weights
- Weak: ignores liabilities, spending, other wealth
- Weak: one period; variance misses skew and tails
- Weak: hides factor overlaps between asset classes

IN THE LESSON

- So what is mean variance optimisation good and bad at?
- Its great strength is transparency.
- You can trace every weight back to a return, a volatility and a correlation, and you can explain it to a client.
- The weaknesses are serious, and exam questions love them.
- Small errors in expected returns produce extreme, concentrated weights, as we just saw.
- It ignores liabilities and the client's spending, so it can pick a mix that fails the very purpose of the money.
- It looks at one period only, and it treats risk as variance, so it misses skewness and fat tails.
- And it can hide factor overlaps, where two asset classes quietly carry the same risk.
- The trap is accepting a corner solution without question.

9 Repairs for each weakness

Each weakness has a standard repair

Problem	Repair
Forecasts are noisy	Reverse optimise
One view swings weights	Black-Litterman
One input set is fragile	Resampling
Corner solutions	Constraints

IN THE LESSON

- Fortunately there are repairs, and each answers one problem.
- Because forecasts are noisy, reverse optimisation starts from the returns implied by market weights, so the starting point is neutral.
- Because a single view can swing the weights, Black-Litterman moves away from that start only as far as the investor's confidence allows.
- Because one set of inputs is fragile, resampling runs the optimiser on many plausible sets of inputs and averages the weights.
- And because corner solutions are not believable, sensible constraints cap what any one asset class can hold.
- Each repair gives up a little theoretical precision for a portfolio you can actually trust.
- We use the first two in the last part of this lesson.

10 Screen, rank, check, defend

Ivor: volatility ≤ 13%, liquid ≥ 10%, floor 4%

Screen by limits → Rank survivors → Check liquidity → Defend choice

IN THE LESSON

- A frontier gives you many portfolios, but the client needs one.
- Think of hiring someone.
- First a screen removes anyone missing a must have.
- Then interviews rank those who are left.
- Then you check references.
- Then you explain your choice to the board.
- Nell does the same for Ivor.
- His rules are three.
- Volatility must stay at or below thirteen per cent.
- At least ten per cent must sit in cash and short bonds, because he plans to buy a house soon.
- And he fears earning less than four per cent in a year.
- So Nell screens by the limits, ranks the survivors, checks liquidity and concentration, and then writes down why.

11 Example 3: the screen

Ivor: volatility ≤ 13%, liquid assets ≥ 10%

Mix A	vol 6.0%, liquid 25%	**passes**
Mix B	vol 9.0%, liquid 18%	**passes**
Mix C	vol 12.0%, liquid 12%	**passes**
Mix D	vol 16.0%, liquid 6%	**fails both**

IN THE LESSON

- First the screen.
- Mix A has volatility six and twenty five per cent liquid, so it passes.
- Mix B has volatility nine and eighteen per cent liquid, and passes.
- Mix C has volatility twelve and twelve per cent liquid.
- That is under the thirteen limit and above the ten per cent liquidity floor, so it passes, with little room to spare.
- Mix D has the highest return, but its volatility of sixteen breaks the limit, and its liquid share of six per cent is too small.
- It fails both tests, so it is out before any ranking starts.
- This is the step many candidates skip.
- They pick the highest return that meets the target and never test the risk constraint.

12 Example 4: rank by utility

Utility, λ = 4, for the three survivors

Mix A	4.8 − 0.02 × 36	**4.08**
Mix B	6.2 − 0.02 × 81	**4.58**
Mix C	7.4 − 0.02 × 144	**4.52**
Highest score	B ahead of C by 0.06	**B**

IN THE LESSON

- Now rank the three survivors with the utility score, using lambda of four.
- Mix A scored four point zero eight, mix B scored four point five eight, and mix C scored four point five two, as we found earlier.
- So utility ranks B first.
- But look at the gap between B and C.
- It is only zero point zero six.
- Forecast errors are far larger than that, so utility alone cannot honestly separate the two.
- This is a good moment to ask what Ivor actually worries about, because the answer tells us which test should lead.

13 Roy's safety-first ratio

SFR = (E(R) − RL) ÷ σ

Expected return	Return the mix should earn
Minimum return	The floor the client must not miss
Volatility	Spread of the mix's returns

IN THE LESSON

- Ivor's worry is not variance in general.
- It is earning less than four per cent.
- For that kind of client we use Roy's safety first ratio.
- Take the expected return, subtract the minimum return the client must not miss, written R L, and divide by the volatility.
- The answer says how many standard deviations the expected return sits above the floor.
- If returns are roughly normal, the mix with the biggest ratio has the smallest chance of falling short.
- Notice what it does not need: no risk aversion coefficient at all, only the floor.
- A mix with a lower return can still win, if its volatility is small enough compared with the cushion above the floor.

14 Example 5: rank by safety first

Floor RL = 4%, survivors A, B and C

Mix A	(4.8 − 4) ÷ 6.0	**0.13**
Mix B	(6.2 − 4) ÷ 9.0	**0.24**
Mix C	(7.4 − 4) ÷ 12.0	**0.28**
Highest ratio	C ahead of B	**C**

IN THE LESSON

- Apply it with a floor of four per cent.
- For mix A, four point eight minus four, divided by six, is zero point one three.
- For mix B, six point two minus four, divided by nine, is zero point two four.
- For mix C, seven point four minus four, divided by twelve, is zero point two eight.
- Now mix C ranks first.
- Its extra return more than pays for its extra volatility when the yardstick is the floor.
- So the same menu and the same client give B under one test and C under the other.
- That is not a contradiction.
- The two tests simply answer different questions.

15 Which test fits which client

Match the test to the client's own words

Client says	Lead with
Happy to trade return for risk	Utility
Must not fall below a floor	Safety-first
Cannot breach a hard limit	Screen first
Needs cash for known bills	Liquidity check

IN THE LESSON

- Which test to lead with comes from what the client says.
- A client happy to trade return for risk, in the language of a risk aversion score, is a utility client.
- A client who says there is a minimum they must not miss is a safety first client.
- A client with a hard limit, such as a maximum loss, needs the screen first, before any ranking at all.
- And a client with known bills to pay needs a liquidity check, whatever the ranking says.
- Ivor is a bit of all of these, but the floor is the worry he stated most firmly, so safety first leads.
- A good answer names the test and says why the client's words call for it.

16 Defending mix C

Why C, and what could go wrong

Volatility limit	12.0 against 13.0	**passes**
Liquid assets	12% against 10%	**passes**
Utility gap to B	4.58 − 4.52	**0.06**
Chance of missing 4%	C 38.8%, B 40.3%	**small edge**
Choice	C, tested and reviewed yearly	**C**

IN THE LESSON

- Now write the defence of mix C, honestly.
- It passes the volatility limit, twelve against thirteen.
- It passes the liquidity rule, twelve per cent against ten.
- Safety first ranks it first, and its gap to B on utility is only zero point zero six.
- The chance of earning under four per cent is about thirty eight point eight per cent for C and about forty point three per cent for B, so the edge is real but small.
- What could go wrong?
- C has the least room under the volatility limit, so a bad year will hurt, and every number rests on forecasts.
- So Nell tests C again with nudged forecasts and reviews it yearly.
- That honest list of reasons and risks is what a good answer contains.

17 Mistakes to avoid so far

- Accepting a corner solution without question
- Trusting precise weights from rough forecasts
- Highest return that meets target, risk untested
- Utility for a floor client, or the reverse
- Skipping liquidity and concentration checks

IN THE LESSON

- Time to stop for the common mistakes so far.
- One: accepting a corner solution without asking whether the forecast behind it is believable.
- Two: trusting precise weights that come from rough forecasts.
- Three: picking the highest expected return that meets the target without testing the risk constraint.
- Four: using the utility score when the client has set a return floor, or the other way round.
- Five: forgetting the liquidity and concentration checks that the optimiser cannot see.
- A slow check of each of these protects an answer from losing marks.

18 A story: the iceberg

- The portfolio is the tip of the iceberg
- Human capital, home, pension lie below
- The whole iceberg sets the real risk
- Does the portfolio offset or add more?

IN THE LESSON

- Think of an iceberg.
- The tip above the water is what you can see, and for an investor that is the financial portfolio.
- But most of the ice sits below the surface: the value of future earnings, the home, and any pension promise.
- A ship is judged by the whole iceberg, not by the tip.
- The same goes for risk.
- The economic balance sheet lists everything the client owns and owes, with human capital, the present value of future earnings, counted as an asset.
- Nell's question for Ivor is simple.
- Does the portfolio lean against the risks hidden below the water, or does it add more of the same?

19 Ivor's economic balance sheet

Assets side, $000. Ivor, 38, risk manager at Calder Bank

Financial portfolio		**600**
Home		**450**
Human capital	PV of salary and bonus	**1,800**
Bond-like pension	PV of promised benefit	**150**
Total assets		**3,000**

IN THE LESSON

- Here is Ivor's economic balance sheet, the assets side, in thousands of dollars.
- His financial portfolio is six hundred.
- His home is four hundred fifty.
- The present value of his future earnings, his human capital, is one thousand eight hundred.
- His pension is a defined benefit promise worth one hundred fifty, and because the promise is fixed it behaves like a bond.
- Together the assets come to three thousand.
- His liabilities, the mortgage and a lifetime of spending, sit on the other side.
- The first thing to notice is size.
- The investments the optimiser looked at are only a fifth of what Ivor has.

20 Where Ivor's wealth sits

Share of total wealth (%)

Human capital	**60%**
Financial	**20%**
Home	**15%**
Pension	**5%**

Share of $3,000k of total assets

IN THE LESSON

- Draw that as bars.
- The vertical axis is the share of total wealth, in per cent.
- Human capital is sixty per cent, the tall bar.
- Financial assets are twenty per cent.
- The home is fifteen.
- The pension is five.
- So Ivor's earning power dominates everything he owns.
- And what drives it?
- He works at Calder Bank, and his bonus and job security rise and fall with the banking sector.
- That makes the biggest bar a bet on banks.
- Now the question is whether his portfolio bets on banks too.

21 Example 6: is the bank bet doubled?

$000. Calder shares 60, global fund 360 (20% financials)

Bank-linked human capital	70% × 1,800	**1,260**
Bank exposure in portfolio	60 + 20% × 360	**132**
Share of total wealth	(1,260 + 132) ÷ 3,000	**46.4%**
After selling, ex-financials fund	1,260 ÷ 3,000	**42.0%**

IN THE LESSON

- Nell judges that seventy per cent of Ivor's human capital is tied to the banking sector.
- That is one thousand two hundred sixty.
- Now his portfolio.
- He holds sixty in Calder Bank's own shares, and his global equity fund of three hundred sixty has twenty per cent in financial companies, which is seventy two.
- So the portfolio holds one hundred thirty two in banks.
- Add them up: one thousand three hundred ninety two out of three thousand, so forty six point four per cent of his total wealth rides on banks.
- Yet the portfolio alone looks diversified.
- Nell sells the Calder shares and moves to a fund without financial companies.
- The exposure falls to the forty two per cent that his career carries by itself.
- The portfolio cannot cure that, but it must not add to it.

22 Time changes the answer

Illustration: Ivor's wealth mix as he ages

Share of Ivor's total wealth	Share
Human capital at 38	60%
Human capital at 58	11%
Portfolio at 38	20%
Portfolio at 58	43%

IN THE LESSON

- Time changes the answer.
- Ivor is thirty eight.
- Now imagine him at fifty eight.
- Each year he works, he uses up a slice of his human capital, so it shrinks, while his savings grow.
- The table is an illustration.
- At thirty eight, human capital is sixty per cent of his wealth and the portfolio twenty.
- At fifty eight, human capital is only eleven per cent and the portfolio forty three.
- The portfolio becomes the main event, and it usually has to become less aggressive, because there is no big pool of future earnings left to absorb a market fall.
- There is one exception: a large bond-like pension.
- It already supplies safe income, so it can make room for more equity risk elsewhere.

23 Judging the fit

- Size each part: earnings, home, pension, portfolio
- Ask what drives the risk of each part
- Offset the shared exposure, do not compound it
- Older client: less human capital, less aggressive
- A large bond-like pension can fund more equity

IN THE LESSON

- Let us make this a routine for any client.
- First, size each part of the balance sheet: earning power, home, pension and portfolio.
- Second, ask what drives the risk of each part: a bank, a sector, inflation, interest rates.
- Third, compare with the portfolio.
- Does it offset those risks or compound them?
- Fourth, adjust the portfolio away from the shared exposure.
- Fifth, remember age.
- As human capital shrinks, the portfolio usually becomes less aggressive, unless a big bond-like pension carries the safe side.
- The trap is approving an allocation that looks diversified by itself while the client's total wealth is concentrated.

24 A story: the crowd's guess

- Market weights are what all investors hold together
- Reverse optimisation reads the returns they imply
- Views and client facts move us off them, with reasons
- Risk tolerance sets how much to hold, not what

IN THE LESSON

- Imagine a village fair where hundreds of people guess the weight of an ox.
- Most single guesses are poor, but the average is surprisingly good.
- The global market portfolio is like that average.
- It holds every investable asset in proportion to its market value, so its prices already blend the views of millions of investors.
- It is not perfect, but it is a neutral place to start, and it needs no forecasts from us.
- From that start Nell may move away only as far as she can justify, one reason for each move.
- Risk tolerance then decides how much of the portfolio to hold, not what is in it.

25 From market weights to Ivor's portfolio

Building

Market weights → Implied returns → Add own views

Fitting

Client facts → Scale with cash

IN THE LESSON

- Here is the recipe, in two rows.
- The first row is building.
- Start from the market weights, the share of each asset class in the world.
- Run the optimiser backwards, which is reverse optimisation, to find the returns that make those weights the best choice.
- Those are the implied returns.
- Then add your own views with Black-Litterman, which blends each view with the implied returns according to your confidence, so you get tilts and not extreme bets.
- The second row is fitting.
- Adjust for the client's own facts: liabilities, currency, liquidity, taxes and restrictions.
- Finally, scale the risky portfolio up or down with cash to match risk tolerance.

26 Reverse optimisation

Π = λ × Σ × w(market)

Implied returns	Return the market must expect
Risk aversion	Risk aversion of the market
Covariances	All volatilities and correlations
Market weights	Share of each asset by value

IN THE LESSON

- The reverse step has one formula.
- The implied return of each asset, above cash, equals lambda times the covariance matrix times the market weights.
- Lambda is the risk aversion of the market as a whole.
- The covariance matrix collects every volatility and correlation in one table.
- Multiply by the weights and you get, for each asset, the return the market must be expecting in order to be happy holding that much of it.
- The more risk an asset adds to the market portfolio, the higher its implied return.
- These are the neutral inputs that Black-Litterman starts from.
- No forecaster's opinion has gone in.

27 Example 7: the implied returns

65% equities (vol 15%), 35% bonds (vol 6%), corr 0.2, λ = 3

Covariance	0.2 × 0.15 × 0.06	**0.0018**
Implied equity return	3 × (0.65×0.0225 + 0.35×0.0018)	**4.58%**
Implied bond return	3 × (0.65×0.0018 + 0.35×0.0036)	**0.73%**
Market volatility	√0.010766	**10.38%**

IN THE LESSON

- Take Nell's market proxy.
- It is sixty five per cent equities with volatility fifteen, and thirty five per cent bonds with volatility six.
- The correlation is point two and the market's risk aversion is three.
- First the covariance: point two times point one five times point zero six, which is zero point zero zero one eight.
- The implied equity return is three times, sixty five per cent of the equity variance, zero point zero two two five, plus thirty five per cent of the covariance.
- That is four point five eight per cent above cash.
- The implied bond return, using the bond variance of zero point zero zero three six, is zero point seven three per cent.
- And the market's own volatility is ten point three eight per cent.
- Equities carry more risk, so the market must expect more from them.

28 Example 8: scale to Ivor's risk

Ivor tolerates 8% volatility; market mix has 10.38%

Share in market mix	8.0 ÷ 10.38	**77.1%**
Equities	65% × 77.1%	**50.1%**
Bonds	35% × 77.1%	**27.0%**
Cash	100% − 77.1%	**22.9%**

IN THE LESSON

- Ivor can tolerate volatility of eight per cent, but the market mix has ten point three eight.
- Nell does not reshuffle the mix.
- She scales it.
- Eight divided by ten point three eight is seventy seven point one per cent, so that share of the portfolio holds the market mix, and the rest goes to cash.
- Equities are sixty five per cent of seventy seven point one, which is fifty point one.
- Bonds are thirty five per cent of that, twenty seven point zero.
- Cash is the remaining twenty two point nine per cent.
- The mix inside the risky part stays at market weights.
- Changing the composition arbitrarily to cut risk would throw away the very reference point we chose.

29 Fitting the portfolio to the client

Each departure needs a reason

Client fact	Adjustment
Spends in dollars	Hedge some
Income tied to banks	Exclude banks
House purchase soon	Hold more cash
Taxable account	Low turnover

IN THE LESSON

- Then the client's own facts adjust the portfolio, and each adjustment needs a reason.
- Ivor spends in dollars, so Nell hedges some of the overseas currency exposure.
- His income is tied to banks, so she excludes bank shares, which we saw in the last part.
- He plans to buy a house soon, so she holds more cash.
- And his account is taxable, so she favours low turnover.
- Every one of these moves has a reason attached, and that is the standard: explain every material deviation from the market weights.

30 How far from the market?

- Scaled market mix
- Small view tilts
- Client fit changes
- 20% home equities

IN THE LESSON

- Here is a line of distance from the market weights.
- On the left is the neutral end.
- On the right is far from the market, which needs the strongest reasons.
- The scaled market mix sits at the far left.
- Small view tilts are a little way along.
- Client fit changes, such as excluding banks and hedging currency, are further along, but each one has a reason.
- At the far right sits Ivor's own idea of twenty per cent in home country equities, which has no reason beyond familiarity.
- The further a point sits from the left, the longer and better the justification must be.

31 Example 9: the home bias test

Home country share of global equities assumed to be 4%

Market weight in portfolio	50.1% × 4%	**2.0%**
Ivor's wish	20% of portfolio	**20.0%**
Wish against market	20.0 ÷ 2.0	**10 times**
Defended tilt	6% ÷ 2.0, with reasons	**3 times**

IN THE LESSON

- Ivor says he knows home country companies best and wants twenty per cent of his portfolio in them.
- Nell checks the market weight.
- Suppose home country companies are four per cent of global equities.
- Ivor's equity share is fifty point one per cent, so the market weight in his portfolio is about two per cent.
- Twenty per cent is ten times that.
- Familiarity is not a reason.
- A smaller tilt can be defended: his spending is in dollars, and holding local assets lowers hedging cost.
- Nell proposes six per cent, three times the market weight, and writes the reasons down.
- Adopting a large home overweight with no reason is the classic trap.

32 Recap: the lesson in four lines

$U = E(R) - 0.005 \times \lambda \times \sigma^2$
$SFR = (E(R) - RL) \div \sigma$
$\Pi = \lambda \times \Sigma \times w$

Utility score	Pick highest, then test inputs
Safety-first ratio	Lead with it for a floor
Share of total wealth	Judge risk on the whole sheet
Implied returns	Neutral start, then justify

IN THE LESSON

- Let us pull the lesson together.
- The optimiser scores each mix on the frontier with utility, the return less a risk penalty, and it is only as good as its forecasts, so we repair it.
- For a client with a floor, safety first picks the mix with the biggest cushion above the floor, per unit of volatility.
- Judge every allocation against the share of total wealth the client already carries in the same risk, because human capital is the biggest asset.
- And start from market weights, using implied returns, then justify each step away from them, and scale with cash for risk tolerance.

33 Try three quick questions

1. λ = 5; E(R) 6.5%, σ 10%. Utility?
   Answer: 6.5 − 0.025 × 100 = 4.0.
2. Floor 3%; E(R) 7%, σ 8%. Safety-first ratio?
   Answer: (7 − 3) ÷ 8 = 0.5.
3. Home is 5% of global equities; equity is 40% of portfolio. Home market weight?
   Answer: 40% × 5% = 2.0%.

IN THE LESSON

- Now try three quick questions yourself.
- Pause after each one.
- Question one.
- Risk aversion is five.
- A mix has an expected return of six point five per cent and volatility of ten per cent.
- What is its utility?
- Six point five minus zero point zero two five times one hundred is four.
- Question two.
- A client's floor is three per cent.
- A mix has an expected return of seven and volatility of eight.
- What is the safety first ratio?
- Seven minus three, divided by eight, is point five.
- Question three.
- Home country shares are five per cent of global equities, and the client holds forty per cent of the portfolio in global equities.
- What is the market weight in home country shares?
- Two per cent.
- Well done, and keep practising.