Allocation approaches, risk and risk factors

Asset Allocation

1 A story: one family, three pots

- Dara Whitlock advises the Aldrich Group, a family business
- Pot one: a reserve fund that owes nobody anything
- Pot two: a staff pension plan that owes benefits
- Pot three: family wealth for three personal goals

IN THE LESSON

- Dara Whitlock is an adviser at Pennine Advisers.
- Her client, the Aldrich Group, is a family business with three quite different pots of money.
- A pot is just a pile of money kept for one purpose.
- The first pot is a reserve fund that owes nobody anything.
- The second is a staff pension plan, which has promised benefits to workers and must pay them.
- The third is the family's own wealth, kept for a few personal goals.
- Here is the big idea of this lesson.
- You cannot judge a portfolio until you know what it is for.
- The same shares and bonds can be a good answer for one pot and a poor answer for another.

2 Three objectives on one line

- Asset-only: best trade-off
- Liability-relative: fund obligations
- Goals-based: fund each goal

IN THE LESSON

- Here are the three approaches on a single line.
- On the left, the money is judged by itself.
- On the right, the money is judged against the goals of a person.
- Start on the left with the asset only approach.
- It looks at the assets alone and tries to find the best trade off between expected return and risk.
- Liabilities, if there are any, only matter indirectly.
- In the middle is the liability relative approach.
- Its job is to fund a defined set of obligations, such as pension benefits, so it looks at assets minus liabilities, which we call the surplus.
- On the right is the goals based approach.
- It splits wealth into sub portfolios, one for each goal, each with a required chance of success.

3 Which approach fits each pot?

Dara sorts the Aldrich Group's pots by purpose

Aldrich pot	Approach
Reserve fund, 800m	Asset-only
Staff pension plan, 110m	Liability-relative
Family wealth, 10m	Goals-based

IN THE LESSON

- Now match the approach to the pot.
- The reserve fund holds eight hundred million and has no payout commitments.
- Its board simply wants the best long run return for the risk taken, so the asset only approach fits.
- The staff pension plan holds one hundred and ten million of assets against promised benefits, so the liability relative approach fits.
- The family has ten million of personal wealth and three things it wants to do with it, so the goals based approach fits.
- Notice what decides the answer.
- It is not the size of the pot or the assets it holds.
- It is what the owner needs the money to do, and whether anyone is owed anything.

4 The asset-only objective: a utility score

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

Utility score	Higher score is better
Expected return	Return in per cent
Risk aversion	Bigger means more risk shy
Variance	Volatility squared, per cent

IN THE LESSON

- The asset only investor needs one number that scores a portfolio for return and risk together.
- A common choice is a utility score.
- You start from the expected return and subtract a penalty for risk.
- The penalty is the risk aversion coefficient, written lambda, times the variance, and a constant of zero point zero zero five makes the units work when return and volatility are in per cent.
- A more nervous investor has a bigger lambda and so pays a bigger penalty.
- The portfolio with the highest score wins.
- The other common choice you met earlier is the Sharpe ratio, which is the excess return per unit of volatility.
- Either way, the question is only about the assets themselves.

5 Example 1: choosing for the reserve fund

Risk aversion 4, so the penalty is 0.02 × variance

Growth: 7.0 return, risk 14	7.0 − 0.02 × 196	**3.08**
Balanced: 6.0 return, risk 10	6.0 − 0.02 × 100	**4.00**
Cautious: 5.0 return, risk 6	5.0 − 0.02 × 36	**4.28**

IN THE LESSON

- Example one.
- The reserve fund board has a risk aversion coefficient of four, so the penalty is zero point zero two times the variance.
- Dara tests three portfolios.
- The cautious one expects five per cent with a volatility of six.
- Its variance is thirty six, so the score is five minus point zero two times thirty six, which is four point two eight.
- The balanced one expects six with a volatility of ten and a variance of one hundred, so its score is four.
- The growth one expects seven with a volatility of fourteen and a variance of one hundred and ninety six, so its score is three point zero eight.
- The highest return is not the winner.
- For this board, the cautious portfolio gives the best trade off.

6 Example 2: the same test for the pension plan

Assets 110, liabilities 100, duration 10. Rates −1 point, shares −10%

Asset-only view	X 6.2% against Y 5.0%	**X wins**
Liabilities after shock	100 × 1.10	**110.0**
X: 60% shares, bonds duration 4	59.4 + 45.8 = 105.2	**−4.8**
Y: 30% shares, bonds duration 12	29.7 + 86.2 = 115.9	**+5.9**

IN THE LESSON

- Example two moves to the staff pension plan.
- Its assets are one hundred and ten, its liabilities are one hundred, and the liabilities have a duration of ten.
- Portfolio X is sixty per cent shares and forty per cent bonds with duration four, and expects six point two.
- Portfolio Y is thirty per cent shares and seventy per cent bonds with duration twelve, and expects five.
- On an asset only view, X wins.
- Now test one shock: rates fall one point and shares fall ten per cent.
- The liabilities rise ten per cent to one hundred and ten.
- X ends at one hundred and five point two, a surplus of minus four point eight.
- Y ends at one hundred and fifteen point nine, a surplus of plus five point nine.

7 Goals-based: one pot split into three

Family wealth of 10m, split by goal

Goal bucket	Needed success
Lifestyle spending, 4.5m	95%
House in 8 years, 2.0m	90%
Legacy for grandchildren, 3.5m	60%

IN THE LESSON

- Now the family.
- The goals based approach divides the ten million into sub portfolios, and each is given its own goal and its own required chance of success.
- Lifestyle spending gets four and a half million, and because running out would be serious, it needs a ninety five per cent chance of success.
- The house, wanted in eight years, gets two million and needs ninety per cent.
- The legacy for the grandchildren gets three and a half million and needs only sixty per cent, because missing it would be disappointing rather than damaging.
- See the difference from the other two approaches.
- There is no single portfolio to score.
- There are three, each with a different promise to keep.

8 Example 3: what the house bucket must earn

House bucket: 2.0m today, goal of 2.6m in 8 years

Bucket today		**2.0m**
Goal in 8 years		**2.6m**
Growth needed	2.6 ÷ 2.0	**1.30**
Yearly return needed	1.30^(1/8) − 1	**3.3%**

IN THE LESSON

- Example three shows how a goal turns into a return target.
- The house bucket holds two million today.
- The family wants two point six million in eight years.
- The growth needed is two point six divided by two, which is one point three.
- To turn that into a yearly return, take the eighth root of one point three and subtract one.
- The bucket needs about three point three per cent a year.
- Notice the order of thinking.
- We started from the goal and worked back to a return, rather than starting from a portfolio and asking how good it is.
- The chance of actually getting there is a risk question, and that comes in the next part.

9 A story: one thermometer, three scales

- Objective decides what risk means
- Asset-only: how much the return wobbles
- Liability-relative: surplus wobbles, or goes below zero
- Goals-based: failing a goal within its horizon

IN THE LESSON

- Imagine three people checking the same weather.
- A sailor cares about wind, a farmer about rain, and a runner about heat.
- The sky is the same, but each person reads a different number, because each has a different thing to protect.
- Investors are like that.
- Once you know what an investor is trying to achieve, you also know what risk means to them.
- For the asset only investor, risk is how much the portfolio's return wobbles.
- For the pension plan, risk is the surplus wobbling, or falling below zero.
- For the family, risk is simply failing to reach a goal in time.
- The same holding can look safe on one scale and dangerous on another.

10 Three definitions of risk

What risk means under each approach

Approach	Risk means
Asset-only	Volatility, tracking error
Downside add-on	CVaR
Liability-relative	Surplus risk, shortfall
Goals-based	Chance of missing goal

IN THE LESSON

- Here are the three definitions side by side.
- For the asset only investor, risk is volatility, the standard deviation of portfolio return.
- If there is a policy benchmark, it is tracking error, the volatility of the gap to the benchmark.
- Some investors add downside measures such as conditional value at risk, which averages the worst losses beyond a chosen cut off.
- For the liability relative investor, risk is the volatility of the surplus, or of the funded ratio, plus shortfall risk, which is the chance that assets fall below liabilities.
- For the goals based investor, risk is the probability of failing to meet each goal within its horizon.

11 Surplus volatility for a pension plan

$$\sigma S^2 = \sigma A^2 + (L/A)^2 \sigma L^2 - 2\,(L/A)\,\rho\,\sigma A\,\sigma L$$

Asset volatility	Volatility of asset returns
Liability volatility	Swing in liability value
Correlation	How assets move with liabilities
Liabilities over assets	Scales the liability side

IN THE LESSON

- To measure surplus risk we need a formula.
- The surplus return is the change in assets minus the change in liabilities, divided by assets.
- Its variance looks like the two asset portfolio formula, with the liabilities treated as a short position.
- That is why the correlation term is subtracted.
- Start with asset volatility squared.
- Add the liability volatility squared, scaled by liabilities over assets, all squared.
- Then subtract two times liabilities over assets, times the correlation, times both volatilities.
- The higher the correlation between assets and liabilities, the bigger the subtraction, and the smaller the surplus risk.
- Here liability volatility means how much the value of the liabilities swings, in per cent of the liabilities.

12 Example 4: cash against long bonds

Assets 110, liabilities 100, liability volatility 12%

Liabilities over assets	100 ÷ 110	**0.909**
Scaled risk, squared	$(0.909 \times 12)^2 = 10.91^2$	**119.0**
Cash, volatility 1	√(1 + 119.0 − 0)	**10.95%**
Long bonds, risk 11, corr 0.95	√(121 + 119.0 − 228.0)	**3.47%**

IN THE LESSON

- Example four tests the trustee who wants to move the pension plan into cash.
- Liabilities over assets are one hundred divided by one hundred and ten, which is zero point nine zero nine.
- The liabilities swing twelve per cent, so the scaled liability risk is ten point nine one, and its square is about one hundred and nineteen.
- Cash has a volatility of one and no correlation with the liabilities.
- The surplus variance is one plus one hundred and nineteen, and the square root gives a surplus volatility of ten point nine five.
- Long bonds have a volatility of eleven and a correlation of point nine five.
- The variance is one hundred and twenty one plus one hundred and nineteen minus two hundred and twenty eight, which gives a surplus volatility of three point four seven.

13 Safe alone, risky against the liabilities

Volatility (% a year)

Cash: assets	**1.0**
Cash: surplus	**10.95**
Bonds: assets	**11.0**
Bonds: surplus	**3.47**

Same plan, same two portfolios, two ways to measure

IN THE LESSON

- Here are those results as bars.
- The vertical axis is volatility in per cent a year.
- The first bar is cash on its own.
- Its asset volatility is one, almost nothing.
- The second bar is the same cash measured as surplus volatility, and it jumps to ten point nine five.
- Cash does not move with the liabilities, so the liability swing is left completely unhedged.
- The third bar is long bonds on their own, with an asset volatility of eleven.
- An asset only investor would call them dangerous.
- The fourth bar is the long bonds' surplus volatility, only three point four seven, because they move with the liabilities and most of the swing cancels out.
- For this plan the so called risk free asset is the risky one.

14 Example 5: the chance of a shortfall

Cushion today; normal approximation, expected return ignored

Cushion in per cent of assets	(110 − 100) ÷ 110	**9.09%**
Cash: z score	9.09 ÷ 10.95	**0.83**
Cash: chance of shortfall	1 − Φ(0.83)	**20%**
Long bonds: z score	9.09 ÷ 3.47	**2.62**
Long bonds: chance of shortfall	1 − Φ(2.62)	**0.4%**

IN THE LESSON

- Shortfall risk turns surplus volatility into a probability.
- Today the plan has a cushion: assets exceed liabilities by ten, which is nine point zero nine per cent of assets.
- A shortfall means a loss bigger than the cushion.
- As a rough illustration we treat the surplus as normally distributed and ignore expected return.
- For cash, divide the cushion by the surplus volatility, nine point zero nine over ten point nine five, to get a z score of point eight three.
- The chance of falling below zero is about twenty per cent.
- For long bonds the z score is nine point zero nine over three point four seven, which is two point six two, and the chance is only about point four per cent.

15 Example 6: risk for the house bucket

House 2.0m to 2.6m; need 3.3%; offer 4.5%, risk 6%

Log growth needed	ln(2.6 ÷ 2.0)	**0.262**
Expected log growth	8 × (4.5% − 0.18%)	**0.346**
Spread over 8 years	6% × √8	**0.170**
z score	(0.262 − 0.346) ÷ 0.170	**−0.49**
Chance of success	Φ(0.49); needed 90%	**69%**

IN THE LESSON

- For the goals based family, risk is the chance of missing a goal.
- Take the house bucket again, with two million to grow into two point six in eight years.
- Dara proposes a portfolio expecting four point five per cent a year with a volatility of six.
- That beats the three point three per cent we needed, but is it enough?
- As a simple illustration, use log growth.
- The growth needed is the natural log of one point three, which is point two six two.
- The expected log growth is eight times four point five per cent minus half the variance, which is point three four six.
- The spread is six per cent times the square root of eight, point one seven.
- The z score is minus point four nine, so the chance of success is only about sixty nine per cent.

16 Example 7: fixing a goal that fails

Move 0.3m from legacy; chance was 69%

House bucket raised to	from 2.0m	**2.3m**
Log growth needed	ln(2.6 ÷ 2.3)	**0.123**
z score	(0.123 − 0.346) ÷ 0.170	**−1.31**
Chance of success	Φ(1.31); needed 90%	**91%**
Legacy bucket now	3.5m − 0.3m	**3.2m**

IN THE LESSON

- The house bucket needs ninety per cent but offers only sixty nine, so it fails its own test.
- Dara has two levers.
- She can lower the goal, or she can put more money in the bucket.
- Suppose the family moves three tenths of a million from the legacy bucket, raising the house bucket to two point three million.
- The growth needed falls to the natural log of two point six over two point three, which is point one two three.
- The expected log growth and the spread stay at point three four six and point one seven.
- The z score becomes minus one point three one, and the chance of success rises to ninety one per cent.
- The legacy bucket falls to three point two million, a trade the family must agree to.

17 Mistakes to avoid so far

- Asset-only objective for a plan that owes benefits
- Calling cash risk-free against long-dated liabilities
- Return above the need does not make success likely
- Judging a goal by volatility, not the chance of failing

IN THE LESSON

- Four traps before we move on.
- First, applying an asset only objective to an investor who has liabilities.
- In example two, portfolio X looked better on return, yet one shock left the plan under water.
- Second, calling cash risk free for a pension plan.
- Cash has low asset volatility, but against long liabilities it leaves a large surplus volatility.
- Third, in the goals based approach, assuming that an expected return above the return needed is enough.
- The house bucket expected more than the return it needed, and still fell short of its required chance.
- Fourth, judging a goal by volatility instead of by the chance of failing it.

18 A story: building blocks

- Asset class: assets sharing systematic risk exposures
- Allocation becomes a choice of exposures
- A bad list of classes gives the optimiser bad blocks
- Tests keep the list sensible

IN THE LESSON

- Think of a child with a box of building blocks.
- If the box holds ten blocks that are really the same block, the child can build very little.
- If it holds a few different shapes, the child can build almost anything.
- Asset classes are the blocks of an allocation.
- An asset class is a group of assets that share systematic risk exposures, meaning exposures that come from the whole market rather than from one company.
- So the allocation decision becomes a choice of exposures.
- If the list of classes is badly drawn, the optimiser works on a bad box of blocks.
- It may treat two copies of the same exposure as different things, and the result is noise.

19 The tests for an asset class

Within a class

Similar members → No overlap

Across the set

Diversifies the rest → Covers most wealth → Absorbs real money

IN THE LESSON

- There are five tests, and they come in two groups.
- The top row is about each class on its own.
- Its members should be similar in risk and return, so the class behaves like one exposure.
- And the classes should not overlap, so each asset belongs to only one class.
- The bottom row is about the set of classes as a whole.
- Each class should diversify the others, which means that including it expands the efficient frontier.
- Together the classes should cover most of the investable wealth that exists.
- And each should be able to absorb a meaningful share of the portfolio without liquidity problems.
- The test people apply too loosely is the diversification test, and we meet it next.

20 Residual risk: how much is truly new?

Residual risk = σ × √(1 − R²)

Class volatility	Volatility of the candidate
R squared	Share explained by old classes
Residual risk	What is left unexplained

IN THE LESSON

- How do we test diversification properly?
- Not by looking at correlations in pairs.
- A new class can look different from each existing class one at a time, and still be explained by all of them together.
- So we regress the new class on all the existing classes at once.
- The R squared tells us the share of its variance that the existing classes explain.
- What is left over is the residual risk.
- It equals the new class's volatility times the square root of one minus R squared.
- A big residual means the class brings something new to the portfolio.
- A small residual means it is a noisy copy of what you already own.

21 Example 8: three candidate classes

Existing classes: shares, bonds, real estate. Volatility in %

Domestic large-cap, R^2 0.97	17 × √(1 − 0.97)	**2.9%**
Infrastructure, R^2 0.88	15 × √(1 − 0.88)	**5.2%**
Commodity futures, R^2 0.25	20 × √(1 − 0.25)	**17.3%**
Verdict	only commodities pass	**add**

IN THE LESSON

- Example eight.
- The Aldrich reserve fund holds three classes now: shares, bonds and real estate.
- Dara regresses each candidate on those three together.
- Domestic large cap shares have a volatility of seventeen and an R squared of point nine seven.
- The residual risk is seventeen times the square root of point zero three, which is two point nine.
- That is almost nothing, so this is not a class, only a slice of shares.
- Listed infrastructure has a volatility of fifteen and an R squared of point eight eight, giving five point two.
- Commodity futures have a volatility of twenty and an R squared of point two five, giving seventeen point three.
- Only commodity futures pass the diversification test.

22 A new class pushes the frontier out

Chart: Expected return vs Risk (volatility)

- Existing classes
- With commodity futures

IN THE LESSON

- Here is why diversification matters, as a picture.
- The horizontal axis is risk, measured as volatility.
- The vertical axis is expected return.
- The grey curve is the efficient frontier built from the existing classes, the best return available at each level of risk.
- The green curve is the frontier after commodity futures are added.
- It lies above and a little to the left of the grey one.
- At the same risk the investor can expect more return, or the same return with less risk.
- That is what the diversification test is checking.
- Had we added domestic large cap shares instead, the two curves would sit on top of each other, and the only change would be extra estimation error.

23 Example 9: can the class absorb real money?

Reserve fund 800m; a niche market worth 40m in all

Fund size		**800m**
Proposed 5% of the fund	800 × 0.05	**40m**
Whole market		**40m**
Share of the market held	40 ÷ 40	**100%**
At 1% instead	8 ÷ 40	**20%**

IN THE LESSON

- The last test is capacity.
- Suppose a trustee wants the reserve fund to hold five per cent in a niche market, say vintage wine, where the whole investable market is worth forty million.
- The fund is eight hundred million, so five per cent is forty million.
- That is one hundred per cent of the entire market.
- The fund could not buy that without moving prices, and it could not sell it back easily either.
- Even one per cent, which is eight million, would be twenty per cent of the market.
- A class like this may look attractive on its own numbers, but it fails the test of absorbing a meaningful share of the portfolio.
- At best it is a small satellite holding, not a building block.

24 A story: the label and the ingredients

- Asset classes are bundles of risk factors
- Factors: growth, inflation, real rates, credit, duration
- Different labels can hide the same factor
- Look through the label to the ingredient

IN THE LESSON

- Look at a fruit smoothie.
- The cup says strawberry banana, but the real question is what is inside: how much sugar, how much fibre, how much water.
- An asset class label is the name on the cup.
- A risk factor is an ingredient.
- A corporate bond carries interest rate duration and credit spread exposure.
- Shares carry growth, inflation and liquidity exposure.
- Factors include inflation, real rates, economic growth, credit, duration, size, value, momentum, liquidity and volatility.
- Two classes with different labels may share the same ingredient.
- A portfolio can look nicely spread across four labels and still be one big bet.
- Looking through the label to the ingredient is how you find that out.

25 Bundles of factors

Illustrative main exposures of four asset classes

Asset class	Main factors
Government bond	Duration
Corporate bond	Duration, credit
Shares	Growth, inflation, liquidity
Inflation-linked bond	Real rate, inflation

IN THE LESSON

- Here are four classes and their main ingredients.
- A government bond is mostly duration, its sensitivity to interest rates.
- A corporate bond carries duration and credit spread, so it holds one more ingredient.
- Shares carry growth, and also inflation and liquidity exposure.
- An inflation linked bond carries the real rate and inflation.
- Now notice what this means.
- Corporate bonds and shares both lean on economic growth, because credit spreads widen when growth disappoints.
- So a portfolio holding both is less diversified than the two labels suggest.
- The class view says two exposures.
- The factor view says there is a shared one hiding between them.

26 Exposure to one factor

Exposure = Σ weight × factor exposure

Weight	Share of the portfolio
Factor exposure	Holding's sensitivity to factor
Portfolio exposure	Sum across all holdings

IN THE LESSON

- To see through the labels we add up exposures.
- For any one factor, the portfolio's exposure is the sum, over every holding, of the weight in that holding times the holding's exposure to the factor.
- Take the equity market factor, where the exposure is the beta.
- A listed share fund has a beta near one.
- A government bond has a beta near zero.
- Private equity and hedge funds have betas that need care, because smoothed prices hide their true sensitivity, so analysts often unsmooth the returns first.
- Once you have the exposure to each factor, you can compare the portfolio against what the investor actually wants, instead of against a list of labels.

27 Example 10: the hidden equity bet

Government bonds 10%, equity beta 0. Betas are illustrative

Shares 30%, beta 1.0	0.30 × 1.0	**0.30**
High yield 20%, beta 0.5	0.20 × 0.5	**0.10**
Hedge funds 20%, beta 0.4	0.20 × 0.4	**0.08**
Private equity 20%, beta 1.1	0.20 × 1.1	**0.22**
Portfolio equity beta	0.30 + 0.10 + 0.08 + 0.22	**0.70**

IN THE LESSON

- Example ten.
- The Aldrich growth portfolio is thirty per cent in listed shares, twenty in high yield bonds, twenty in hedge funds, twenty in private equity and ten in government bonds.
- Its trustees think of it as thirty per cent in equities.
- Dara looks through to the equity market factor.
- Shares have a beta of one, so they add point three zero.
- High yield bonds have a beta of point five, adding point one zero.
- Hedge funds have point four, adding point zero eight.
- Private equity, after unsmoothing, has one point one, adding point two two.
- Government bonds add nothing.
- The total is point seven zero.
- The trustees believe in thirty, but the real equity exposure is seventy.

28 Where the equity risk comes from

Equity beta added (hundredths)

Shares 30%	**30**
High yield 20%	**10**
Hedge funds 20%	**8**
Private eq. 20%	**22**

Non-equity labels add 40 of 70, which is 57%

IN THE LESSON

- Here is that result as bars.
- The vertical axis is the equity beta each holding adds, in hundredths.
- The first bar is listed shares at thirty per cent of the portfolio, adding thirty.
- That is the only part the trustees were counting.
- The second bar is high yield bonds, at twenty per cent, adding ten.
- The third is hedge funds, at twenty per cent, adding eight.
- The fourth is private equity, at twenty per cent, adding twenty two.
- Together the three holdings that do not carry the equity label add forty out of seventy, which is fifty seven per cent of the total.
- Four labels, one dominant exposure.
- The board thought it owned a balanced portfolio, and it owned mostly a growth bet.

29 Example 11: a zero-investment factor position

Value factor: long cheap shares, short expensive shares

Long cheap shares	50m × 9%	**+4.5m**
Short expensive shares	50m × 5%	**−2.5m**
Net money invested	50 − 50	**0**
Factor profit	4.5 − 2.5	**2.0m**
Gain on the whole fund	2.0 ÷ 800	**0.25%**

Gross exposure 100m, 12.5% of the fund

IN THE LESSON

- Allocating by factor means choosing exposures directly, usually with long short portfolios that need no net investment.
- Suppose Dara wants more value exposure for the eight hundred million reserve fund.
- She buys fifty million of cheap shares, which return nine per cent, and sells short fifty million of expensive shares, which return five.
- The long side gains four point five million.
- The short side costs two point five.
- Net money invested is nothing, because the short sale funds the purchase.
- The factor profit is two million, which is point two five per cent of the fund.
- The price is complexity: the position needs shorting and borrowing, and its gross exposure is one hundred million, twelve and a half per cent of the fund.

30 Asset classes against factors

Factor-based allocation, judged against classes

Question	Factor lens
Reveals shared exposures	Yes
Needs shorting, leverage	Often
Easy to explain to a board	Less so
Can reach a similar mix	Yes

IN THE LESSON

- So should Aldrich allocate by class or by factor?
- Here is the comparison.
- The factor lens reveals shared exposures that distinct looking classes hide, which is its great advantage.
- The drawbacks are practical.
- Factor portfolios are usually long short, so they need shorting and often leverage, and the implementation is more complex.
- They are also harder to explain to a board that thinks in terms of shares and bonds.
- The reassuring part is that the two routes can reach similar efficient portfolios.
- Because of that, many investors keep a class based allocation and use the factor lens to analyse it, as we did with the growth portfolio.

31 Mistakes with classes and factors

- Overlapping classes add only estimation error
- Testing diversification pair by pair, not by regression
- Ignoring capacity: a class must absorb real money
- Reading labels, not exposures: 30% equity is not 30% risk
- Treating classes and factors as rivals

IN THE LESSON

- Five traps to close.
- First, overlapping classes.
- Domestic large cap shares beside global shares add nothing except estimation error.
- Second, testing diversification pair by pair.
- Regress the candidate on all the existing classes together, as in example eight.
- Third, forgetting capacity.
- A class has to absorb a meaningful share of the portfolio.
- Fourth, reading labels instead of exposures.
- Thirty per cent in equities did not mean thirty per cent of equity risk.
- Fifth, treating classes and factors as rivals.
- They are two views of the same portfolio, and the factor view is often best used to check the class view.

32 Recap: the lesson in five lines

Utility, then surplus risk
Residual risk, then factor exposure

Asset-only	Score return against risk
Liability	Surplus volatility, shortfall
Goals	Chance of missing each goal
New class	Needs large residual risk
Factors	Look through labels to exposure

IN THE LESSON

- Let us pull it together.
- The asset only objective scores the assets for return against risk, using utility.
- The liability relative objective grows the surplus, and its risk is the volatility of the surplus and the chance of a shortfall, so cash can be the risky choice.
- The goals based objective funds each goal, and its risk is the chance of missing it.
- Asset classes are building blocks that must be similar inside, not overlap, diversify one another, cover most wealth and absorb real money, and diversification is tested by residual risk.
- Finally, factors are the ingredients inside the labels, and exposure is the sum of weight times exposure.

33 Try three quick questions

1. Pension plan, long liabilities: riskier against them, cash or long bonds?
   Answer: Cash: it leaves the liability swing unhedged.
2. Volatility 14, R^2 0.75 against existing classes. Residual risk?
   Answer: 14 × √0.25 = 7.0%.
3. 40% beta 1.0, 30% beta 0.6, 30% beta 0. Equity beta?
   Answer: 0.40 + 0.18 + 0 = 0.58.

IN THE LESSON

- Now try three quick questions yourself.
- Pause after each one.
- Question one.
- A pension plan has long dated liabilities.
- Which is riskier measured against the liabilities, cash or long bonds?
- Cash, because it leaves the whole liability swing unhedged, while long bonds move with the liabilities.
- Question two.
- A candidate class has a volatility of fourteen and an R squared of point seven five against the existing classes.
- What is its residual risk?
- Fourteen times the square root of point two five, which is seven.
- Question three.
- A portfolio holds forty per cent in shares with a beta of one, thirty per cent in a fund with a beta of point six, and the rest in bonds with a beta of zero.
- What is its equity beta?
- Point five eight.