Key terms: risk and return

Key terminology

1 How to use these terms

- Each term: a story, a plain meaning, a tiny example
- Spread and shape, then moving together
- Then the price of risk, then portfolios and indexes
- Every number is a made-up round example

IN THE LESSON

- Each term comes with a short story, a plain meaning and a tiny example that we compute together.
- Nothing is real data.
- The terms fall into families.
- First, spread and shape: how bumpy returns are.
- Second, moving together: how two investments relate, and what that does to risk.
- Third, the price of risk: what return is fair for the risk you carry.
- Fourth, portfolios, indexes and honest testing.
- Each term uses the ones before it, so go in order the first time.
- Later, when a topic uses a word you forget, come back and jump straight to its slide.

2 Variance and standard deviation

Variance = Σ p × (R − mean)²
Standard deviation = √Variance
Past data: ÷ (n − 1)

The mean	The average outcome
Variance	Average squared gap from mean
Standard deviation	Square root: back in return units

IN THE LESSON

- Picture two bus routes with the same average journey time.
- On route one the bus takes the same time every day.
- On route two it is quick one day and painfully slow the next.
- Same average, very different ride.
- Variance captures that difference.
- Find the mean, the average outcome.
- Measure how far each outcome sits from it.
- Square each gap, so gaps above and below do not cancel.
- Then average the squares, weighting each by its probability.
- The units are squared, which is awkward, so take the square root.
- That is the standard deviation, back in the units of the returns, and the usual measure of risk.
- For past data, divide by the number of observations minus one instead.
- You will meet both in Quantitative Methods, Equity, Alternative Investments and Portfolio Management.

3 Example: the Corvin Fund's year

Equal chance of −10% or +30% next year

Mean return	½ × (−10) + ½ × 30	**10%**
Gaps from the mean	−10 − 10 and 30 − 10	**−20, +20**
Squared gaps	20^2 and 20^2	**400, 400**
Variance	average of 400 and 400	**400**
Standard deviation	√400	**20%**

Bell-shaped: about 2 years in 3 within one SD

IN THE LESSON

- A tiny example.
- The Corvin Fund, a made-up fund, has an equal chance of losing ten per cent or gaining thirty next year.
- The mean is half of minus ten plus half of thirty, ten per cent.
- The gaps from the mean are minus twenty and plus twenty.
- Square each and you get four hundred.
- The average of those is four hundred.
- That is the variance, in per cent squared, which nobody can picture.
- Take the square root and you get twenty per cent: the standard deviation.
- In a bell shaped world, about two years in three land within one standard deviation of the mean.
- Remember: the standard deviation counts a pleasant surprise as risk just as much as a nasty one.

4 Value at risk

Portfolio $10m, 1,000 days of history

Worst 5% of 1,000 days	5% × 1,000	**50 days**
The 50th-worst day	count in from the worst	**−$1.0m**
Value at risk, 5%, one day	that 50th-worst loss	**$1.0m**
Meaning	1 day in 20 loses at least this	**$1.0m+**
Worst day of all	beyond the line	**−$4.0m**

5% means 1 day in 20

IN THE LESSON

- A shop owner does not ask what the worst day could ever be.
- They ask: how much could I lose on a bad day, one that comes only now and then?
- Value at risk answers that.
- It is a loss figure for a set period and probability.
- Take a ten million dollar portfolio with a thousand days of history.
- Five per cent of a thousand days is fifty days, one day in twenty.
- Line the days up from the worst.
- The fiftieth worst loss is one million dollars.
- So the five per cent, one day value at risk is one million dollars: a loss at least that large on about one day in twenty.
- The warning: the worst day lost four million, and the figure says nothing about it.
- You will meet it in Quantitative Methods, Alternative Investments and Portfolio Management.

5 Skewness and negative skew

Teasel Strategy: ten months of returns

Nine months at +2%	9 × 2	**+18%**
One month at −18%	the rare flood	**−18%**
Mean of ten months	(18 − 18) ÷ 10	**0%**
Median month	middle of the sorted list	**+2%**
Mean below median	long tail on the left	**Negative skew**

IN THE LESSON

- Think of a shop that earns a small profit nearly every day, and once in a while a flood wipes out a month of takings.
- That lopsidedness is what skewness measures.
- A symmetrical shape has no skew.
- Negative skew means a long tail on the left: many small gains and the occasional large loss, with the mean pulled below the median, the middle value.
- Take the Teasel Strategy, a made-up one, over ten months.
- Nine months it gains two per cent.
- One month it loses eighteen.
- The mean is zero, yet the median month is plus two, so the mean sits below the median: negative skew.
- Standard deviation treats gains and losses alike, so it understates the risk that matters.
- Skewness turns up in Quantitative Methods and Portfolio Management, negative skew in Alternative Investments too.

6 Fat tails

Months in 1,000

Bell curve predicts	**1**
Made-up record	**9**

Months 12 points or more below the mean (3 SD, SD = 4%)

IN THE LESSON

- If rainfall followed a neat bell curve, a once in a century flood would come once in a century.
- Real rivers flood more often, and returns behave the same way.
- A distribution with fat tails puts more probability on extreme outcomes than the normal bell curve does.
- Its measure is positive excess kurtosis.
- Look at the bars.
- A thousand monthly returns have a standard deviation of four per cent.
- A fall of three standard deviations is twelve points.
- The bell curve says only about one month in a thousand should be that bad.
- In this made-up record, nine were.
- So any risk measure built on the bell curve, value at risk included, understates how often large losses occur.
- You will meet fat tails in Quantitative Methods, Alternative Investments and Portfolio Management.

7 Covariance and correlation

Cov(A, B) = Σ (A − mean A)(B − mean B) ÷ (n − 1)
ρ = Cov ÷ (σA × σB)

Covariance	Sign tells; size has units
Correlation (rho)	Always from −1 to +1
Risk of A	Standard deviation of A
Risk of B	Standard deviation of B

IN THE LESSON

- Two dancers can move in step, in opposition, or ignore each other.
- Covariance puts a sign on that.
- For each period, multiply how far the first investment sits from its own average by how far the second sits from its own.
- Add these up and divide by the number of observations minus one.
- Positive means they rise and fall in step.
- Negative means they offset.
- But the size carries the units of both investments, so you cannot compare one pair with another.
- Correlation fixes that.
- Divide the covariance by both standard deviations and the answer runs from minus one to plus one.
- It is written with the Greek letter rho.
- Covariance appears in Quantitative Methods and Portfolio Management.
- Correlation appears in many topics, from Equity to Derivatives.

8 Example: Quince Foods and Nettle Power

Quince −4, 0, 4, 4, 6 | Nettle 0, 4, 2, 4, 5

Means	Quince 2, Nettle 3	**2 and 3**
Gap products added	18 − 2 − 2 + 2 + 8	**24**
Covariance	24 ÷ (5 − 1)	**6**
Standard deviations	Quince 4, Nettle 2	**4 and 2**
Correlation	6 ÷ (4 × 2)	**0.75**

IN THE LESSON

- Here are five months for two made-up companies.
- Quince Foods returned minus four, zero, four, four and six per cent.
- Nettle Power returned zero, four, two, four and five.
- The means are two and three.
- Multiply each pair of gaps and you get eighteen, minus two, minus two, two and eight, which add to twenty-four.
- Divide by four, one less than the five months, and the covariance is six.
- Positive, so they move together.
- The standard deviations are four and two.
- Divide six by four times two and the correlation is zero point seven five: a clear, not perfect, tendency to move in step.

9 Rho: one letter, two jobs

Same Greek letter, two jobs

Two meanings of rho	What it is
In portfolios: ρ	Correlation
In options: rho	Rate sensitivity
Call, rates up 1 point	+0.50
Put, rates up 1 point	−0.40

IN THE LESSON

- One letter does two jobs.
- In portfolio work, rho is simply the correlation we just met.
- You will see it in Quantitative Methods and Portfolio Management, so a question that gives you rho is giving you a correlation.
- The second job is in options.
- A call is the right to buy at a fixed strike price, a put the right to sell at it.
- The Greek rho there measures how an option's value changes when the risk free interest rate rises by one point.
- A call gains, say fifty cents, because a higher rate shrinks the present value of the strike you will pay.
- A put loses, say forty cents.
- That second meaning is a preview of later study.

10 Correlation and the risk of a mix

Two funds, each risk 20%, half and half

Correlation	Risk of the mix
+1.0, move in step	20.0%
+0.5	17.3%
0.0, unrelated	14.1%
−1.0, opposite	0.0%

IN THE LESSON

- Correlation decides what combining investments does to risk.
- Take two funds, each with a risk of twenty per cent, and put half your money in each.
- The weighted average of their risks is twenty.
- If they move perfectly in step, correlation plus one, the mix has a risk of exactly twenty: no help at all.
- At plus zero point five the risk of the mix falls to seventeen point three.
- With no relationship, correlation zero, it falls to fourteen point one.
- In perfect opposition, correlation minus one, the risks cancel and the mix has none.
- The rule to keep: any correlation below plus one pulls portfolio risk beneath the weighted average of the parts.
- You can check these with the portfolio variance formula in Quantitative Methods.

11 Diversification

Two funds, each: expected 8%, risk 20%, correlation 0

Each fund alone	expected, risk	**8%, 20%**
Half and half: expected	0.5 × 8 + 0.5 × 8	**8%**
Half and half: risk	20 × √0.5	**14.1%**
Risk removed	20.0 − 14.1	**5.9 points**

IN THE LESSON

- An umbrella seller does well when it rains.
- A sunscreen seller does well when it is sunny.
- A shopkeeper who owns both stalls has a steadier income than either owner alone, and has not given up any average income.
- That is diversification: holding things that do not move together, so their independent disappointments cancel.
- Two funds each expect eight per cent a year, with a risk of twenty and a correlation of zero.
- Put half in each.
- The expected return of the mix is still eight.
- Its risk is twenty times the square root of one half, fourteen point one.
- Risk fell by five point nine points and expected return did not move.
- Almost nothing else offers that.
- You will meet diversification in many topics, from Quantitative Methods to Alternative Investments.

12 Systematic risk

Portfolio risk (%)

1 share	**30.0**
4 shares	**20.7**
10 shares	**18.2**
40 shares	**16.9**
Very many	**16.4**

Each share risk 30%, correlation 0.3: the floor is systematic risk

IN THE LESSON

- Think of bakeries in a city.
- A baker who burns the bread has a problem of their own.
- A power cut across the city hurts every baker at once.
- The first is specific risk, and holding many bakeries makes it fade.
- The second is systematic risk: risk that moves the whole market and cannot be diversified away.
- The bars show it.
- Each share has a risk of thirty per cent, and any two have a correlation of zero point three.
- One share carries thirty.
- Four shares bring it to twenty point seven.
- Ten, eighteen point two.
- Forty, sixteen point nine.
- With very many it levels at sixteen point four and never goes lower.
- That floor is systematic risk, and it is the only kind the market pays you for.
- You will meet it in Equity Investments and Portfolio Management.

13 Beta

Beta = Cov(share, market) ÷ Var(market) = ρ × σshare ÷ σmarket

Beta	1 = market; >1 more; <1 less
Covariance with market	How share and market move together
Market variance	Market's swings, squared
Correlation (rho)	How tightly it follows the market

IN THE LESSON

- Imagine a dog on a lead, and the market is the person walking it.
- Beta says how far the dog moves when the walker moves one step.
- A beta of one means the dog follows step for step.
- Above one, the dog overshoots, amplifying the market.
- Below one, it trails behind, damping it.
- Beta measures systematic risk, the only risk the market pays for.
- It equals the covariance of the share with the market, divided by the variance of the market.
- Rewritten, it is the correlation, times the share's standard deviation, over the market's.
- A wild share with a weak link to the market can have a small beta.
- You will meet it in Quantitative Methods, Corporate Issuers, Equity Investments and Portfolio Management.

14 Example: same risk, different betas

Market 15%. Each share 24%. Correlation: Quince 0.75, Nettle 0.5

Market variance	15^2	**225**
Quince covariance	0.75 × 24 × 15	**270**
Quince beta	270 ÷ 225	**1.2**
Nettle beta	0.5 × 24 ÷ 15	**0.8**
Market falls 10%: Quince, Nettle	1.2 × −10, 0.8 × −10	**−12%, −8%**

IN THE LESSON

- The market has a standard deviation of fifteen per cent, so its variance is two hundred and twenty-five.
- Quince Foods and Nettle Power each have a standard deviation of twenty-four, so they look equally risky.
- Quince has a correlation with the market of zero point seven five.
- Its covariance is zero point seven five times twenty-four times fifteen, two hundred and seventy.
- Divide by the market variance and the beta is one point two.
- Nettle follows the market less tightly, a correlation of zero point five.
- Its beta is zero point five times twenty-four over fifteen, zero point eight.
- If the market falls ten per cent, Quince tends to fall about twelve and Nettle about eight.
- Equal risk, different beta.

15 Dividend yield

Share at $50 pays $2 a year; a year later it is $52

Dividend yield	2 ÷ 50	**4%**
Price return	(52 − 50) ÷ 50	**4%**
Total return	4% + 4%	**8%**
Same $2 at a price of $100	2 ÷ 100	**2%**

IN THE LESSON

- Think of an apple tree.
- Each year it drops apples, and over the years it also grows taller.
- The apples are the dividends.
- The growth in height is the rise in the share price.
- The dividend yield is the apples alone: the annual dividend per share divided by the share price.
- A share priced at fifty dollars pays two dollars a year.
- The yield is four per cent.
- If the price rises to fifty-two, the price return is four per cent too, and the total return is the two added together, eight.
- The yield moves with the price: the same two dollars on a share priced at one hundred is only two per cent.
- In Derivatives it is often a continuous yield, written delta, which lowers the forward price and the value of a call.
- You will meet yield in Quantitative Methods, Corporate Issuers and Equity Investments.

16 The risk-free rate

A government bill pays 3% for one year

Lend $100 safely	100 × 1.03	**$103**
Risky asset expects	market view	**8%**
Excess over risk-free	8% − 3%	**5%**
PV of a strike of 100	100 ÷ 1.03	**$97.09**

IN THE LESSON

- Suppose you lend money to a government for a year.
- You are paid only for waiting.
- That payment is the risk free rate: the return on an investment with no default risk over your horizon, usually proxied by short term government bills.
- No asset is perfectly safe, but this is the yardstick against which every risky return is measured.
- With a rate of three per cent, one hundred dollars grows to one hundred and three.
- A risky asset expecting eight per cent promises an excess of five.
- Option models also use it, discounting the exercise price at that rate, so the present value of a strike of one hundred is ninety-seven point zero nine.
- You will meet it in most topics, from Quantitative Methods to Portfolio Management.

17 Risk aversion

Sure 5%, or a gamble: expected 8%, risk 10% (variance 0.01)

Moderate, A = 4: utility	0.08 − ½ × 4 × 0.01	**6.0%**
Cautious, A = 10: utility	0.08 − ½ × 10 × 0.01	**3.0%**
Sure thing's utility	no variance, no penalty	**5.0%**
Who takes the gamble?	cautious keeps the sure 5%	**Moderate**

IN THE LESSON

- Offer a friend a sure prize, or a coin toss for nothing or double the prize.
- The average is the same.
- Most people take the sure prize.
- They are risk averse.
- Risk aversion is preferring less uncertainty for the same expected return, and accepting a lower expected return to get it.
- That does not mean refusing all risk, only wanting to be paid for it.
- A common score is utility: expected return minus one half times A times the variance.
- A bigger A means more dislike of risk.
- Offer a sure five per cent, or a gamble expecting eight with a risk of ten.
- A moderate investor, A equal to four, scores the gamble six per cent and takes it.
- A cautious one, A equal to ten, scores it three and keeps the sure five.
- You will meet it in Derivatives, Alternative Investments and Portfolio Management.

18 The equity risk premium

- Premium = expected market return − Rf
- Example: 8% − 3% = 5% a year
- It cannot be observed; it is estimated
- From history, models of dividends, or surveys

IN THE LESSON

- Think of hazard pay.
- A firm sending a worker up a tall mast pays more than for a desk job, because the job is riskier.
- The equity risk premium is hazard pay for owning shares: the return investors expect from equities above the risk free rate.
- If the market is expected to return eight per cent and the safe rate is three, the premium is five per cent a year.
- The awkward fact is that nobody can observe it.
- So it is estimated, from long historical averages, from models that use dividends or earnings, or from surveys of investors.
- You will meet it in Quantitative Methods, Corporate Issuers, Equity Investments and Portfolio Management.

19 The capital asset pricing model (CAPM)

Required return = Rf + β × ERP, where ERP = E(Rm) − Rf

Risk-free rate	Reward for waiting
Beta	Market part of the risk
Equity risk premium	Expected market return less Rf

IN THE LESSON

- Now we can put the pieces together.
- The capital asset pricing model, said C A P M, says the return an investor should require from an asset is the risk free rate, plus beta times the equity risk premium.
- The first part pays you for waiting.
- The second pays you for worry, and the worry is measured only by beta, the market part of the risk.
- The rest of the risk can be diversified away, so the model says it earns nothing.
- A share with no market risk earns only the risk free rate.
- A share that moves with the market earns the market's return.
- You will meet the model in Corporate Issuers, where it sets the cost of equity, and in Equity Investments and Portfolio Management.

20 Example: required returns from the CAPM

Rf = 3%, market expected 8%, so ERP = 5%

Quince, beta 1.2	3% + 1.2 × 5%	**9.0%**
Nettle, beta 0.8	3% + 0.8 × 5%	**7.0%**
The market, beta 1.0	3% + 1.0 × 5%	**8.0%**
Beta 0	3% + 0 × 5%	**3.0%**

IN THE LESSON

- The risk free rate is three per cent and the market is expected to return eight, so the premium is five.
- Quince Foods has a beta of one point two.
- Its required return is three plus one point two times five, which is nine per cent.
- Nettle Power, beta zero point eight, needs three plus four, seven.
- The market itself, beta one, needs eight.
- An asset with a beta of zero needs just three.
- A higher beta demands a higher return.
- Two slips to avoid: forgetting to add the risk free rate back, and using the market return in place of the premium, which gives an answer far too high.

21 The Sharpe ratio

Sharpe ratio = (Rp − Rf) ÷ σp

Portfolio return	What the portfolio earned
Risk-free rate	What safe cash earns
Portfolio risk	Standard deviation of returns

IN THE LESSON

- Two students sit an exam.
- One scores a little higher after months of study, the other after hardly any.
- Who did better?
- Judged by the score alone, the first.
- Judged by the score per hour of effort, the second.
- The Sharpe ratio does this for investing.
- Take the return in excess of the risk free rate, the reward, and divide it by the portfolio's standard deviation, the effort, meaning total volatility.
- On a chart of return against risk it is the slope of the line from the risk free asset to the portfolio.
- A higher ratio is better.
- But because the standard deviation counts upside and downside alike, the ratio flatters negatively skewed strategies, whose risk shows up rarely.
- You will meet it in Alternative Investments and Portfolio Management.

22 Example: comparing two funds

Rf = 3%. Vetch: return 11%, risk 16%. Thistle: return 9%, risk 8%

Vetch: excess return	11% − 3%	**8%**
Vetch: Sharpe ratio	8 ÷ 16	**0.50**
Thistle: excess return	9% − 3%	**6%**
Thistle: Sharpe ratio	6 ÷ 8	**0.75**
More reward per unit of risk		**Thistle**

IN THE LESSON

- The risk free rate is three per cent.
- The Vetch Fund returned eleven per cent with a standard deviation of sixteen.
- Its excess return is eight, and eight over sixteen gives a Sharpe ratio of zero point five.
- The Thistle Fund returned only nine, but with a standard deviation of eight.
- Its excess return is six, and six over eight gives zero point seven five.
- The Vetch Fund earned more, but the Thistle Fund earned more reward for each unit of risk.
- The ratio is a ranking number, not a percentage return.

23 Mistakes and look-alikes

- 1. Variance is squared units: quote the standard deviation
- 2. Covariance has units; correlation runs from −1 to +1
- 3. A wild share is not a high-beta share
- 4. Sharpe uses total risk: negative skew can flatter it
- 5. Value at risk is a line, not the worst loss

IN THE LESSON

- Five look-alikes that cost marks.
- One.
- Variance and standard deviation.
- Variance is in squared units, so quote the standard deviation.
- Two.
- Covariance and correlation.
- Covariance carries the units of both assets, correlation runs only from minus one to plus one.
- Three.
- A wild share and a high beta share.
- A wild share with a weak link to the market can have a small beta, because beta counts only the market part.
- Four.
- The Sharpe ratio uses total risk, so a negatively skewed strategy can look better than it is.
- Five.
- Value at risk is a line that bad days seldom cross, not the worst loss.

24 The minimum-variance portfolio

Fennel: 6%, risk 20. Tansy: 10%, risk 30. Correlation 0.25

Covariance	0.25 × 20 × 30	**150**
Weight in Fennel	(900 − 150) ÷ (400 + 900 − 300)	**75%**
Expected return of the mix	0.75 × 6 + 0.25 × 10	**7.0%**
Risk of the mix	√337.5	**18.4%**
Fennel alone	no Tansy	**20.0%**

IN THE LESSON

- Imagine sliding your money from one fund to another and tracking the risk at every step.
- The risk traces a bowl, and the bottom is the minimum variance portfolio: the combination of risky assets with the lowest possible variance.
- Take the Fennel Fund, expecting six per cent with a risk of twenty, and the Tansy Fund, expecting ten with a risk of thirty.
- Their correlation is zero point two five, so the covariance is a hundred and fifty.
- The weight in Fennel at the bottom is nine hundred minus a hundred and fifty, over four hundred plus nine hundred minus three hundred, which is zero point seven five.
- The mix has a risk of eighteen point four, lower than Fennel alone at twenty.
- A bit of the riskier fund made the whole safer.
- You will meet it in Quantitative Methods and Portfolio Management.

25 The efficient frontier

Chart: Expected return (%) vs Risk (standard deviation, %)

- Minimum variance: 18.4%, 7.0%
- Efficient frontier
- Beaten: same risk, less return

IN THE LESSON

- Now draw the menu.
- The bottom axis is risk, standard deviation in per cent.
- The side axis is expected return in per cent.
- Every mix of the Fennel and Tansy funds is a point, and together they trace a curve.
- The leftmost point is the minimum variance portfolio, at a risk of eighteen point four and a return of seven.
- From there the green curve climbs up and to the right.
- That upper part is the efficient frontier: the portfolios with the highest expected return for each level of risk.
- Now the red branch curling back down.
- Straight above any point on it sits a green point with the same risk and a higher return, so anything beneath the frontier is beaten by something on it.
- Which point is best depends on the investor's risk aversion.
- You will meet it in Quantitative Methods and Portfolio Management.

26 Three ways to weight an index

- Price-weighted: add the share prices, divide by a divisor
- Market-capitalisation-weighted: price × number of shares
- Equal-weighted: every member gets the same weight
- Same market, three rules, three answers

IN THE LESSON

- A village hall committee must pick a representative for the street.
- One rule says the person with the priciest hat speaks loudest.
- A second says the biggest landowner.
- A third gives everyone an equal vote.
- Indexes face the same choice.
- A price-weighted index adds up the share prices and divides by a divisor, so a high priced company sways it whatever its size, and a share split changes the answer with no economics behind it.
- A market-capitalisation-weighted index weights each share by the market value of its shares, price times number of shares.
- An equal-weighted index gives every member the same weight.
- You will meet all three in Quantitative Methods and Equity Investments.

27 Weights under three rules

Prices $10, $30, $100; shares 60m, 10m, 1m

Mallow, price-weighted	100 ÷ 140	**71.4%**
Mallow, cap-weighted	100 ÷ 1,000	**10.0%**
Mallow, equal-weighted	1 ÷ 3	**33.3%**
Sorrel, cap-weighted	600 ÷ 1,000	**60.0%**
Mallow splits 2 for 1	50 ÷ 90	**55.6%**

IN THE LESSON

- Three made-up shares.
- Sorrel costs ten dollars with sixty million shares.
- Bramble costs thirty with ten million.
- Mallow costs a hundred with one million.
- Sorrel is worth six hundred million and Mallow a hundred million, out of a total of one thousand million.
- The prices add to a hundred and forty.
- Look at Mallow, the smallest company.
- Price weighting gives it a hundred over a hundred and forty, seventy-one point four per cent.
- Market capitalisation weighting gives it only ten per cent.
- Equal weighting gives it a third.
- Sorrel, the biggest, carries sixty per cent by value.
- Now let Mallow split two for one.
- Its price halves to fifty, the price total drops to ninety, and its price weight falls to fifty-five point six per cent, though nothing real has changed.

28 One year, three answers

Index return (%)

Price-weighted	**5.7**
Cap-weighted	**4.0**
Equal-weighted	**3.3**

Sorrel +10%, Bramble −10%, Mallow +10%

IN THE LESSON

- Let a year pass.
- Sorrel and Mallow each rise ten per cent and Bramble falls ten.
- The bars show the index returns, in per cent.
- The tallest is the price-weighted index, up five point seven: Mallow has the biggest price, so the biggest say, and it rose.
- The middle bar is the market-capitalisation-weighted index, up four.
- Sorrel is the biggest by value.
- The third is the equal-weighted index, up three point three, the plain average of plus ten, minus ten and plus ten.
- Not one share behaved differently: the whole gap comes from the weighting rule.

29 Where each rule leans

What each weighting rule favours

Rule	Leans towards
Price-weighted	High-priced shares
Cap-weighted	Largest companies
Equal-weighted	Smaller companies
Lowest turnover	Cap-weighted

IN THE LESSON

- None of the rules is neutral.
- A price-weighted index leans towards high priced shares, which says nothing about size.
- A market-capitalisation-weighted index leans towards the largest companies.
- Its weights move with prices on their own, so a fund copying it trades only when shares outstanding change, giving the lowest turnover.
- An equal-weighted index tilts towards smaller companies, because a small firm counts as much as a giant.
- But prices keep pulling the weights away from equal, so the fund must trade back after every move.
- The index pays no trading costs, so its published return overstates what a fund copying it nets.

30 Scenario analysis

Fund $100m: $60m in shares, $40m in bonds

Shares fall 20% alone	20% × $60m	**−$12m**
Bonds fall 10% alone	10% × $40m	**−$4m**
Worst single test	the larger of the two	**−$12m**
Stagflation: both together	12 + 4	**−$16m**
Loss one-at-a-time missed	16 − 12	**$4m**

IN THE LESSON

- Think of a fire drill.
- The marshal does not test one door at a time.
- They imagine a whole fire: smoke, a blocked stairwell and a power cut at once.
- Scenario analysis changes a coherent set of inputs together to describe a recognisable state of the world, such as a recession.
- It captures risks that arrive together, which one at a time testing misses.
- The made-up Hawthorn Fund holds a hundred million dollars: sixty in shares, forty in bonds.
- Shares falling twenty per cent alone cost twelve million.
- Bonds falling ten alone cost four.
- But in a stagflation scenario both happen together: sixteen million, four more than the worst single test.
- You will meet it in Economics, Financial Statement Analysis and Equity Investments.

31 Survivorship bias

100 funds started; 60 survive at 8%, 40 closed at 2%

Survivors only report	database view	**8%**
Weighted sum, all funds	60 × 8 + 40 × 2	**560**
True average, all funds	560 ÷ 100	**5.6%**
Closed funds average	dropped out	**2%**
Record overstates by	8.0 − 5.6	**2.4 points**

IN THE LESSON

- Imagine judging restaurants by walking down a street of the ones still open.
- They all look good, because the bad ones have closed.
- That is survivorship bias: the upward tilt from studying only the funds or firms that still exist, because the failures have dropped out of the database.
- Let us size it.
- A hundred funds started.
- Sixty survive, averaging eight per cent a year.
- Forty closed, averaging two.
- A database holding only survivors reports eight.
- The true average over all hundred is sixty times eight plus forty times two, over a hundred, which is five point six per cent.
- The record overstates by two point four points, and it understates risk, because the busts are missing.
- You will meet it in Ethical and Professional Standards, Quantitative Methods, Alternative Investments and Portfolio Management.

32 Recap: the relationships to remember

SD = √Variance, Beta = ρ × σs ÷ σm
Required = Rf + β × ERP
Sharpe = (Rp − Rf) ÷ σp

Standard deviation	Risk in return units
Diversification	Cuts risk, keeps return
Beta	Market risk only
CAPM	Pay for waiting and worry
Sharpe ratio	Reward per unit of total risk

IN THE LESSON

- Let us pull it together.
- Risk is measured by the standard deviation, the square root of the variance, though negative skew and fat tails mean the bell curve understates large losses, value at risk included.
- Covariance, rescaled, becomes correlation, and any less than perfect correlation lets diversification lower risk without lowering expected return.
- What remains is systematic risk, which beta measures and which alone is paid for.
- The capital asset pricing model turns beta into a required return: the risk free rate plus beta times the equity risk premium.
- The Sharpe ratio scores reward per unit of total risk.
- Along the efficient frontier, risk aversion picks the spot.
- And whatever the index, ask how it is weighted and what was left out.

33 Try three quick questions

1. 5% one-day value at risk of $1m: what does it mean?
   Answer: A loss of $1m or more on about 1 day in 20.
2. Rf 2%, ERP 4%, beta 1.5. Required return?
   Answer: 2% + 1.5 × 4% = 8.0%.
3. Closed funds missing, or many small gains and rare big loss?
   Answer: First: survivorship bias. Second: negative skew.

IN THE LESSON

- Now try three quick questions yourself.
- Pause after each one.
- Question one.
- What does a five per cent, one day value at risk of one million dollars mean?
- A loss of at least one million dollars on about one day in twenty, and nothing about how bad the worst day gets.
- Question two.
- The risk free rate is two per cent, the equity risk premium is four per cent and beta is one point five.
- What is the required return?
- Two plus one point five times four is eight per cent.
- Question three.
- One record leaves out the funds that closed.
- Another has many small gains and the occasional big loss.
- Which is survivorship bias and which is negative skew?
- The first is survivorship bias, the second negative skew.
- Well done.