CHARTA LEARNING · LESSON NOTES

Value at risk: three methods and what comes after

Portfolio Management

1 A story: the bad day line

- VaR = a loss level, a probability and a time horizon
- $200k VaR at 5%: beaten 1 day in 20
- A threshold on the left tail, not the worst case
- Three ways: parametric, historical, Monte Carlo

IN THE LESSON

- Imagine you run a small fund and your boss asks one question.
- On a really bad day, how much do we lose?
- You cannot say never, and you cannot say everything.
- So you pick a line.
- You say, on nineteen days out of twenty we lose less than this amount.
- That amount is the value at risk at five per cent.
- It has three ingredients: a probability, a time horizon such as one day, and a sum of money.
- So a one day five per cent VaR of two hundred thousand dollars reads like this.
- On one day in twenty, we expect to lose at least that much.
- It is a threshold, not a worst case.

2 VaR on the return curve

Chart: Likelihood vs Daily return

- VaR cut-off
- Most days are here
- Worst 5% lies left of the line

IN THE LESSON

- Here is the picture.
- Along the bottom are daily returns, from losses on the left to gains on the right.
- Up the side is how likely each return is.
- The bell shaped curve is the normal distribution, with most days near the middle.
- Now find the vertical line on the left.
- Everything to the left of it is the worst five per cent of days.
- The line itself is the VaR cut-off.
- The tail to its left holds the days that are worse than the VaR, and notice the picture says nothing about how far that tail stretches.
- Remember that, because it comes back when we discuss weaknesses.
- For a normal curve the line sits one point six five standard deviations below the mean.

3 Where each answer comes from

Parametric

Estimate means, volatilities, correlations → Assume a normal curve → z times volatility

Historical and Monte Carlo

Scenarios: past days or model draws → Revalue today's portfolio → Read the percentile

IN THE LESSON

- Here is how the three methods differ.
- Read the top row first.
- The parametric method boils the portfolio down to expected returns, volatilities and correlations.
- It assumes a normal curve, and then simply multiplies a z value by the volatility.
- It never looks at the individual positions again.
- Now the bottom row, shared by the other two methods.
- Both create a long list of scenarios, then revalue today's portfolio in every scenario, then read off the loss at the chosen percentile.
- The only difference is where the scenarios come from.
- Historical simulation takes them from a real past window.
- Monte Carlo generates them from a model that you build.

4 Method one: parametric

- Inputs: expected returns, volatilities, correlations
- Assumes normal returns and linear positions
- Strength: fast, simple, easy to explain
- Weakness: fat tails and options are misstated

IN THE LESSON

- The parametric method, also called variance covariance, is the quick one.
- Think of a recipe card.
- You need the expected return of each holding, its volatility, and the correlations between them.
- From those you work out the portfolio's mean and volatility, assume returns follow the normal curve, and read the loss from a table of z values.
- It is fast and needs no stored history.
- But it has two serious weaknesses.
- Real returns have fat tails, so big losses are more common than the normal curve says.
- And options do not move in a straight line with the underlying price, so a method that treats every position as linear misstates their risk.

5 Method two: historical simulation

- Replay actual past returns on today's portfolio
- No distribution assumed; full revaluation handles options
- Rank the outcomes, read the percentile
- Weakness: only as severe as the window

IN THE LESSON

- Historical simulation is the replay method.
- Take the last five hundred trading days.
- For each day, ask what would have happened if that day's market moves hit today's portfolio, and revalue every position.
- You now have five hundred possible outcomes.
- Sort them from worst to best, and read off the one at your percentile.
- It makes no assumption about the shape of returns, and because it revalues each position it handles options properly.
- The weakness is the window.
- If those five hundred days were calm, the answer will be calm too.
- If a crash is missing from the window, the method cannot imagine it.
- That is the exam trap: history cannot show a risk it never contained.

6 Method three: Monte Carlo

- Draw thousands of scenarios from an assumed model
- Fully revalue the portfolio in each scenario
- Strength: flexible, suits complex instruments
- Weakness: model dependent and heavy to compute

IN THE LESSON

- Monte Carlo simulation is the build your own history method.
- You write a model for how the risk factors behave, including, if you wish, fat tails and changing volatility.
- A computer then draws thousands of random scenarios from that model.
- Each scenario is used to fully revalue the portfolio, and the outcomes are ranked just as in historical simulation.
- It is the most flexible method, and the best fit for complex instruments.
- The price you pay is effort.
- It takes a lot of computing time, and the answer is only as good as the model.
- A badly chosen model produces thousands of precise looking numbers that are all wrong in the same way.

7 Example 1: three methods, one book

$10m book, one day, 5%; daily volatility 1.0%, mean zero

Parametric	1.65 × 1.0% × $10m	**$165,000**
Historical	25th worst of 500: 2.1%	**$210,000**
Monte Carlo	500th worst of 10,000: 1.9%	**$190,000**
Gap	$210,000 − $165,000	**$45,000**

Parametric is lowest: it assumes away tails and curvature

IN THE LESSON

- Let us compare all three on a single book.
- An analyst runs a ten million dollar book of shares and index options, over one day at five per cent.
- Daily volatility is one per cent with zero mean.
- The parametric method gives one point six five times one per cent times ten million, which is one hundred and sixty five thousand dollars.
- Historical simulation replays five hundred days, and the twenty fifth worst loses two point one per cent, so two hundred and ten thousand dollars.
- Monte Carlo takes the five hundredth worst of ten thousand draws from a fat tailed model, one point nine per cent, so one hundred and ninety thousand.
- The parametric figure is lowest because it ignores fat tails and option curvature.

8 Choosing between the methods

The main weak spot of each method

Method	Weak spot
Parametric	fat tails, options
Historical	window may be calm
Monte Carlo	model and compute cost
Best practice	run more than one

IN THE LESSON

- Here is the summary to keep.
- For each method we show the weak spot, which is what an exam will test.
- Parametric: it misstates risk for fat tails and for options.
- Historical simulation: it depends entirely on the window, so a calm window gives a calm answer.
- Monte Carlo: it depends on the model, and it is heavy to compute.
- In practice good teams run more than one method and ask why the answers differ.
- If the parametric figure is far below the other two, that is a warning that the portfolio has tail or option risk that the normal curve cannot see.
- A gap is information, not a nuisance.

9 Portfolio volatility from the pieces

$$\sigma_p = \sqrt{(w_1^2\sigma_1^2 + w_2^2\sigma_2^2 + 2w_1w_2\rho\sigma_1\sigma_2)}$$

Weights (w)	How much of each asset
Volatilities (σ)	How much each one wobbles
Correlation (ρ)	How much the wobbles cancel
Portfolio volatility (σp)	The answer: portfolio risk

IN THE LESSON

- Before the VaR itself we need the portfolio's volatility.
- For two assets it is the square root of the first weight squared times the first volatility squared, plus the second weight squared times the second volatility squared, plus two times both weights times the correlation times both volatilities.
- Look at the pieces.
- The weights say how much you hold of each.
- The volatilities say how much each wobbles.
- The correlation tells you how much the wobbles cancel.
- The lower the correlation, the smaller the last term, and the smaller the portfolio volatility.
- That shrinkage is the benefit of diversification, and it feeds straight into VaR.

10 The parametric VaR formula

$$VaR = (z \times \sigma - \mu) \times V$$

z value (z)	Tail size: 1.65 at 5%
Daily volatility (σ)	Standard deviation for the horizon
Daily mean return (μ)	Expected return for the horizon
Portfolio value (V)	Money value of the portfolio

IN THE LESSON

- Here is the parametric VaR.
- As a return, it is the z value times the daily standard deviation, minus the daily mean return.
- Multiply by the portfolio value and you have the loss in money.
- The z value depends on the probability.
- It is one point six five at five per cent, one point nine six at two and a half per cent, and two point three three at one per cent.
- Take the pieces.
- The mean return pushes the loss figure down a little, because on average you earn something.
- The standard deviation does the heavy lifting.
- Always convert to the same horizon first: a daily VaR needs a daily mean and a daily standard deviation.

11 Example 2: the inputs (set-up)

$20m: 60% equities (10%, 20%), 40% bonds (4%, 6%), correlation 0.2

Expected return	0.6 × 10 + 0.4 × 4	**7.6%**
Equity term	(0.6 × 20)²	**144**
Bond term	(0.4 × 6)²	**5.76**
Correlation term	2 × 0.6 × 0.4 × 0.2 × 20 × 6	**11.52**
Volatility	√(144 + 5.76 + 11.52)	**12.70%**

IN THE LESSON

- Example two.
- A twenty million dollar portfolio holds sixty per cent in equities with an expected return of ten per cent and volatility of twenty per cent.
- The other forty per cent is bonds, at four per cent return and six per cent volatility.
- The correlation is zero point two, and there are two hundred and fifty trading days a year.
- Step one, the expected return.
- Zero point six times ten plus zero point four times four gives seven point six per cent a year.
- Step two, the volatility.
- The three terms under the root are one hundred and forty four, five point seven six, and eleven point five two.
- They add to one hundred and sixty one point two eight, and the root is twelve point seven per cent.

12 Example 2: daily figures and the VaR

250 trading days a year; 5% VaR, z = 1.65

Daily mean	7.6 ÷ 250	**0.030%**
Daily volatility	12.70 ÷ √250	**0.803%**
VaR as a return	1.65 × 0.803 − 0.030	**1.295%**
VaR in money	1.295% × $20m	**$259,000**

IN THE LESSON

- Step three, move to daily figures.
- The mean is divided by the number of trading days, so seven point six over two hundred and fifty is zero point zero three zero per cent.
- The volatility is divided by the square root of the number of days.
- The root of two hundred and fifty is about fifteen point eight, so twelve point seven over that is zero point eight zero three per cent.
- Step four, the five per cent VaR.
- One point six five times zero point eight zero three, minus zero point zero three zero, is one point two nine five per cent.
- Times twenty million dollars, that is about two hundred and fifty nine thousand dollars.
- On one day in twenty, we expect to lose at least that much.

13 Rarer means larger

One day VaR ($000)

5%, z 1.65	**259**
2.5%, z 1.96	**309**
1%, z 2.33	**368**

Same portfolio; only the probability changes

IN THE LESSON

- Here is the same calculation at three probabilities.
- The bars show the one day parametric VaR for our twenty million dollar portfolio, in thousands of dollars.
- Only the z value changes.
- The first bar is five per cent, with a z of one point six five, giving two hundred and fifty nine.
- The second is two and a half per cent, with a z of one point nine six, giving three hundred and nine.
- The third is one per cent, with a z of two point three three, giving three hundred and sixty eight.
- Asking about a rarer day gives a bigger number, about forty two per cent bigger from five to one.
- That is why a VaR quoted without its probability and horizon cannot be compared with anything.

14 Example 3: historical and Monte Carlo

Same $20m portfolio; 500 past days replayed

Observations in the 5% tail	5% × 500	**25**
Cut-off outcome	25th worst of 500	**−1.48%**
Historical VaR	1.48% × $20m	**$296,000**
Monte Carlo cut-off	5% × 10,000	**500th worst**
Versus parametric	$296,000 − $259,000	**$37,000**

Other conventions interpolate; the answer moves only slightly

IN THE LESSON

- Now the other two methods on the same portfolio.
- For historical simulation, you revalue today's portfolio on each of the last five hundred days and rank the outcomes.
- Five per cent of five hundred is twenty five, so a common convention takes the twenty fifth worst outcome.
- Say it is a loss of one point four eight per cent.
- Times twenty million dollars, that is two hundred and ninety six thousand.
- It is higher than the parametric two hundred and fifty nine thousand, which fits a fat left tail in the actual record.
- Monte Carlo works the same way, but the ranking is of simulated outcomes.
- With ten thousand draws, the five hundredth worst is the five per cent point.

15 Example 4: the scaling trap

Annual volatility 12.70%; which divisor for a daily figure?

Wrong: divide by 250	12.70 ÷ 250	**0.051%**
VaR from the wrong figure	about	**$10,700**
Right: divide by √250	12.70 ÷ 15.81	**0.803%**
Right VaR	from the last example	**$259,000**
Understatement	259,000 ÷ 10,700	**about 24 times**

IN THE LESSON

- The most common mistake is in the units.
- Suppose a candidate divides the annual volatility by two hundred and fifty instead of by its square root.
- Twelve point seven divided by two hundred and fifty is zero point zero five one per cent, not zero point eight zero three.
- The VaR then comes out at about ten thousand seven hundred dollars, against two hundred and fifty nine thousand.
- That understates the risk by a factor of about twenty four.
- Why the square root?
- Variances add across independent days, so the variance grows with the number of days, and the standard deviation grows with its square root.
- The mean, by contrast, grows in direct proportion to time.

16 What the result actually means

- 5% one day VaR: a loss at least this big on 1 day in 20
- About 12 or 13 breaches a year are expected, not a failure
- Not the maximum loss: it says nothing about the tail
- Assumes the portfolio is unchanged over the horizon

IN THE LESSON

- Let us read a result aloud, because the exam asks for it.
- A one day five per cent VaR of two hundred and fifty nine thousand dollars means this.
- If the model is right and tomorrow is a typical day, there is a five per cent chance of losing at least two hundred and fifty nine thousand.
- It also means that on about one trading day in twenty, or roughly twelve days a year, we should expect to breach it.
- It does not mean the most we can lose.
- A breach could cost three hundred thousand or three million, and the number itself is silent on which.
- It also assumes the portfolio is held unchanged for the horizon.

17 Mid-lesson mistakes: estimating VaR

- Volatility scales with √time, not with time
- Keep the mean, the volatility and the horizon consistent
- Subtract the mean; it matters over long horizons
- Use z = 1.65 at 5%, not 1.96
- History cannot show a risk the window never contained

IN THE LESSON

- Let us stop and collect the traps so far.
- First, scaling volatility by the number of days rather than its square root.
- Second, mixing horizons, for example using an annual mean with a daily standard deviation.
- Third, forgetting that the mean is subtracted, and that over longer horizons it matters more.
- Fourth, using the wrong z value, such as one point nine six when the question says five per cent.
- One point nine six belongs to a two and a half per cent tail, or to a two sided ninety five per cent interval.
- Fifth, claiming historical simulation captures risks that are absent from its window.
- Check these five before every calculation and you will win marks others lose.

18 Strengths of VaR

- One money figure that non-specialists can grasp
- Comparable across desks and asset classes
- Supports risk limits and capital allocation
- Widely accepted by regulators

IN THE LESSON

- Why did value at risk become the standard anyway?
- Think of a dashboard in a car.
- You do not need to understand the engine to see the fuel gauge.
- VaR is a fuel gauge for risk.
- It condenses a whole portfolio into one money figure that a board member can grasp.
- It can be compared across desks and asset classes, because the unit is always money at a probability.
- It supports risk limits and the allocation of capital, since you can give each desk a budget.
- And it is widely accepted by regulators, so it is a common language between a firm and its supervisors.
- Those are real benefits, and they explain its popularity.

19 Weaknesses of VaR

- Silent on how big losses beyond the VaR can be
- Depends on method, window, normality
- Understates risk when correlations rise in a crisis
- Not always subadditive; ignores the upside

IN THE LESSON

- Now the other side.
- The same compression that makes VaR easy to read also hides things.
- First, it says nothing about how large losses beyond the threshold can be.
- Second, the estimate depends heavily on the method, the look back period, and assumptions such as normality, so two teams can report different numbers for the same book.
- Third, it tends to understate risk when correlations rise in a crisis, just when you need it most.
- Fourth, it is not always subadditive, which we will see in a moment.
- And it looks only at the left tail, ignoring upside.
- The trap is treating a low VaR as proof of low risk when the tail beyond it is fat.

20 Same VaR, different tails

Loss ($ million)

Normal: VaR	**2.0**
Normal: avg beyond	**2.5**
Fat tail: VaR	**2.0**
Fat tail: avg beyond	**4.1**

Illustrative desks, same one day 5% VaR

IN THE LESSON

- Look at this picture.
- Two desks both report the same one day five per cent VaR of two million dollars.
- On a limit sheet they look identical.
- The second bar in each pair shows the average loss on the days when the threshold is crossed.
- If returns are normal, that average is about one and a quarter times the VaR, so two and a half million.
- Now picture a desk that sells out of the money options, or holds assets that gap down on bad news.
- Its tail can be much deeper, so the average loss beyond the same threshold might be four million or more.
- The four point one here is an illustration, not a rule.
- The threshold cannot tell these two desks apart.

21 Example 5: VaR is not always subadditive

Two independent bonds, each loses 100 with probability 4%

Bond A alone, 5% VaR	4% is under 5%	**0**
Bond B alone, 5% VaR	the same	**0**
At least one default	$1 - 0.96^2$	**7.84%**
Combined 5% VaR	7.84% is over 5%	**100**

Whole (100) exceeds the sum of the parts (0)

IN THE LESSON

- Here is a famous failure.
- Two bonds each lose one hundred if they default, which each does with probability four per cent, independently.
- Otherwise they lose nothing.
- Take bond A alone.
- Four per cent is less than five per cent, so the five per cent cut-off lands on the no loss outcome, and its VaR is zero.
- Bond B is the same, so the sum of the parts is zero.
- Now hold both.
- The chance that at least one defaults is one minus zero point nine six squared, which is seven point eight four per cent.
- That is more than five per cent, so the combined VaR is one hundred.
- The whole shows more risk than the parts added together, which breaks subadditivity.

22 The extensions of VaR

- Conditional VaR: average loss beyond VaR
- Incremental VaR: VaR after − VaR before
- Marginal VaR: change in VaR per small change in size
- Relative VaR: shortfall against the benchmark

IN THE LESSON

- Plain VaR answers one question, so risk managers have built extensions for the next questions.
- Conditional VaR, also called expected shortfall or expected tail loss, is the average loss in outcomes worse than the VaR.
- It is always at least as large as the VaR at the same probability.
- Incremental VaR is the change in the portfolio's VaR from adding, removing or resizing a position: VaR after, minus VaR before.
- Marginal VaR is the change in VaR per small change in a position's size, used to see which holdings drive the risk.
- Relative VaR measures the potential shortfall against a benchmark rather than in absolute terms, and is closely related to ex ante tracking error.

23 Incremental and conditional VaR

Incremental VaR = VaR after − VaR before
Conditional VaR = average loss beyond VaR

Incremental VaR	What the trade really adds
VaR after the trade	Rerun with the new position
VaR before the trade	The portfolio as it stands
Conditional VaR	Mean of the tail days

IN THE LESSON

- Two extensions have formulas.
- Incremental VaR is simply the VaR with the new position, minus the VaR without it.
- It is the number to check against a risk limit when you decide on a trade, because it already includes the offsetting effect of everything else you hold.
- Conditional VaR is the average of the losses beyond the VaR threshold.
- Because it averages only the tail days, it must be at least as big as the VaR.
- The trap is confusing a position's standalone VaR, which pretends it is the only holding, with its incremental VaR, which reflects diversification with the rest of the portfolio.

24 Example 6: incremental VaR

$200m fund, one day 5% VaR $5.0m; new position standalone $1.2m

VaR with the position	rerun the model	**$5.4m**
Incremental VaR	$5.4m − $5.0m	**$0.4m**
Standalone VaR	ignores the rest of the book	**$1.2m**
Diversification	$1.2m − $0.4m	**$0.8m**

IN THE LESSON

- Example six.
- A two hundred million dollar fund has a one day five per cent VaR of five million dollars.
- A proposed position has a standalone VaR of one point two million.
- Rerunning the model with the position included gives five point four million.
- The incremental VaR is five point four minus five point zero, which is zero point four million.
- That is what the trade really adds, not one point two.
- The difference, zero point eight million, is diversification: the existing holdings offset it because the new position does not lose on exactly the same days.
- A limit check should use the smaller, incremental number.
- Marginal VaR would then rank holdings by how much each drives the total.

25 Example 6 continued: conditional VaR

100 simulated days; 5th worst = VaR = $5.0m

Four worse days ($m)	10, 8, 6.5, 5.5	**sum 30**
Conditional VaR	30 ÷ 4	**$7.5m**
Against the VaR	$7.5m ÷ $5.0m	**1.5 times**
Worst five incl. VaR day	35 ÷ 5	**$7.0m**

IN THE LESSON

- Now conditional VaR for the same fund.
- It has one hundred simulated days.
- Five per cent of one hundred is five, so the fifth worst day is the VaR, a loss of five million dollars.
- The four worse days lose ten, eight, six point five and five point five million.
- The conditional VaR is the average of those four, which is thirty million divided by four, or seven point five million.
- That is half as much again as the VaR.
- On the days the VaR is breached, the fund loses seven and a half million on average.
- If you instead average the five worst days, including the VaR day, you get seven million.
- Conventions vary slightly, but conditional VaR always exceeds VaR.

26 Example 7: relative VaR

$200m fund, tracking error 2% a year, 250 days, 5%

Daily tracking error	2 ÷ √250	**0.126%**
Relative VaR as a return	1.65 × 0.126	**0.209%**
Relative VaR in money	0.209% × $200m	**$417,000**

A shortfall against the benchmark, not an absolute loss

IN THE LESSON

- Finally relative VaR.
- A manager is judged against a benchmark, so absolute losses are not the point.
- Suppose the fund holds two hundred million dollars and its ex ante tracking error, the expected volatility of its return minus the benchmark return, is two per cent a year.
- Convert to a daily figure by dividing by the square root of two hundred and fifty, which gives zero point one two six per cent.
- Times one point six five, that is zero point two zero nine per cent.
- Times two hundred million, it is about four hundred and seventeen thousand dollars.
- On one day in twenty, we expect to trail the benchmark by at least that much.
- The fund could be up in absolute terms and still breach it.

27 Sensitivity measures

- Value change for a small move in one risk factor
- Beta; duration and convexity; delta, gamma and vega
- Precise about size and direction of exposure
- No probability; local; hard to add across asset types

IN THE LESSON

- Now the other family of risk measures.
- Think of a car again.
- Sensitivity is how much the speed changes when you press the pedal a little.
- A sensitivity measure shows how value changes for a small move in one risk factor.
- Beta does it for the equity market.
- Duration and convexity do it for interest rates.
- Delta, gamma and vega do it for options, against the price, its curvature and volatility.
- They are precise about the direction and size of an exposure.
- But they carry no probability, they are only reliable for small moves, and they come in different units, so a beta cannot easily be added to a duration across asset types.

28 Scenario measures

- Loss under a specified set of moves in many factors at once
- Drawn from history or built as hypothetical events
- Concrete; captures combined moves and extreme events
- Covers only listed events; gives no probability

IN THE LESSON

- Scenario measures ask a different question.
- Instead of nudging one factor, they say, what would this portfolio lose if a defined set of moves happened together?
- The scenario might come from history, such as a repeat of a past crash, or be hypothetical, such as equities falling and yields rising at the same time.
- Stress tests are scenarios chosen to be extreme.
- They are concrete, and they show combined damage across many factors, including events a model may rate as impossible.
- But a scenario is only as good as the imagination behind it, it covers only the events someone listed, and it says nothing about how likely they are.

29 Example 8: three measures, one portfolio

$50m equities (beta 1.2), $30m bonds (duration 6), VaR $1.1m

Equity sensitivity	1.2 × 1% × $50m	**$0.6m per 1%**
Bond sensitivity	6 × 0.01% × $30m	**$18,000 per bp**
Scenario: equities	1.2 × 20% × $50m	**$12.0m**
Scenario: bonds	6 × 1% × $30m	**$1.8m**
Scenario total	$12.0m + $1.8m	**$13.8m**

Ignores convexity; a 20% fall is a large move

IN THE LESSON

- Example eight.
- A portfolio holds fifty million dollars of equities with a beta of one point two, and thirty million of bonds with modified duration six.
- Its one day five per cent VaR is one point one million.
- The equity sensitivity is one point two times one per cent times fifty million, which is six hundred thousand dollars per one per cent market fall.
- The bond sensitivity is six times one basis point, which is zero point zero one per cent, times thirty million, or eighteen thousand per basis point.
- Those two cannot be added.
- Now the committee's scenario, equities down twenty per cent and yields up one hundred basis points together.
- Equities lose twelve million and bonds one point eight, a total of thirteen point eight million.

30 Weighing the three against each other

What each measure gives, and what it misses

Measure	Gives / misses
Sensitivity	exposure size / probability
Scenario	event loss / probability
VaR	probability / any named event
Best practice	use all three

IN THE LESSON

- Here is the comparison.
- The scenario loss of thirteen point eight million dollars is more than twelve times the VaR of one point one million, yet neither number is wrong.
- They answer different questions.
- Sensitivities tell desks what to hedge, and are precise but local, with no probability.
- Scenarios tell the committee what a named event would cost, and are concrete but cover only what someone listed, again with no probability.
- VaR is the only one of the three that attaches a probability, by passing all exposures through one distribution, but it names no particular event.
- Comparing a scenario loss with a VaR compares a severity with a frequency.
- Use all three together.

31 Recap

- Parametric z × σ; historical replay; Monte Carlo
- Volatility × √time; z = 1.65 at 5%
- VaR is silent on the tail; CVaR averages it
- Incremental = after − before; only VaR gives a probability

IN THE LESSON

- Let us recap.
- The three methods answer the same question.
- Parametric multiplies z by volatility and assumes normal, linear positions.
- Historical simulation replays a past window and is only as severe as that window.
- Monte Carlo draws scenarios from a model and is only as good as the model.
- For the numbers, volatility scales with the square root of time and the mean with time, and z is one point six five at five per cent.
- VaR is easy to communicate but silent about the tail, and can fail subadditivity.
- Conditional VaR averages the tail, incremental VaR is after minus before, and sensitivities and scenarios add detail but give no probability.

32 Quick quiz

1. Annual to daily volatility: what is the divisor?
   Answer: The square root of the trading days, not the days
2. VaR rises from $5.0m to $5.4m; standalone is $1.2m. Incremental?
   Answer: $0.4m: the change in portfolio VaR
3. Which gives a probability: sensitivity, scenario or VaR?
   Answer: Only VaR

IN THE LESSON

- Question one.
- An annual volatility is converted to a daily one.
- What is the right divisor?
- The answer: the square root of the number of trading days, not the number itself.
- Question two.
- A fund's VaR goes from five million to five point four million after adding a position whose standalone VaR is one point two million.
- What is the incremental VaR?
- The answer: zero point four million dollars, the change in the whole portfolio's VaR.
- Question three.
- Which of the three measures, sensitivity, scenario or VaR, gives a probability of loss?
- The answer: only value at risk does, which is why we keep it beside the other two measures.