CHARTA LEARNING · LESSON NOTES

Validating VaR Models

FRM Part II

1 A story: the model that passed

- Kestrel reports a 1-day, 99% VaR every day
- Losses break through about as often as predicted
- Head of trading: 'so the model is proven'
- Validator asks: what does it assume and leave out?
- A good record can hide a model that is wrong

IN THE LESSON

- Kestrel Bancorp has used the same V a R model for years.
- Each day it reports a one day, ninety nine per cent V a R for its trading book, and over the last year the actual losses broke through that figure about as often as the model predicts.
- The head of trading says the model is therefore proven.
- A new validator, Priya, is not convinced.
- She asks what the model actually assumes, which risks it leaves out, how fragile the number is, and how far it can be trusted.
- She knows a good track record can hide a model that is wrong, one that passes by luck or because its errors offset.
- Validation asks those questions on purpose.

2 The validator's toolkit

- 1. Conceptual soundness: is the design right?
- 2. Sensitivity analysis: how much does VaR move?
- 3. Confidence intervals: how precise is VaR?
- 4. Benchmarking: compare with an alternative
- Plus backtesting against actual profit and loss

IN THE LESSON

- A validator has several tools, and this lesson takes four of them.
- First, conceptual soundness, which asks whether the design itself makes sense for this portfolio.
- Second, sensitivity analysis, which asks how much the number moves when one choice is changed.
- Third, confidence intervals, which ask how precise the number is, given that it is estimated from limited data.
- Fourth, benchmarking, which compares the model with an independent alternative.
- Alongside all of them sits backtesting, which compares each day's V a R with the actual profit or loss.
- Backtesting looks at outcomes.
- The other four look inside the model, and none of them replaces the others.

3 Conceptual soundness

- Judged before outputs are compared with outcomes
- Sound for one book can be unsound for another
- Review the documentation and the theory
- Check the code does what the documents say
- Key idea: fit to the positions actually held

IN THE LESSON

- Start with conceptual soundness.
- It asks whether the model, as designed, is a defensible way to measure the risk of this particular portfolio, before any output is compared with outcomes.
- A model can be perfectly sound for a book of shares and unsound for a book of options.
- So the validator begins with the documentation and the theory.
- She reads what the model is supposed to do, what it assumes, and why each choice was made.
- Then she checks that the code actually does what the documents say.
- The key idea is fit.
- Every assumption is judged against the positions the bank really holds.

4 The choice of method

Each method has a portfolio where it breaks

Method	Where it tends to break
Variance-covariance	Fat tails; options and nonlinear books
Historical simulation	Sees only its data window
Monte Carlo	Scenario model risk; heavy computing

IN THE LESSON

- Each method has a portfolio where it breaks.
- The first choice to review is the methodology, and why it was chosen, and the table pairs each common method with its weak spot.
- Variance covariance, also called the parametric method, assumes normal returns and linear positions, so it struggles with fat tails and with options.
- Historical simulation replays past market moves on today's positions, so it can only see what is in its data window.
- A calm window means a low V a R.
- Monte Carlo simulates many random scenarios and can revalue options fully, but it is only as good as the model that generates the scenarios, and it is costly to run.
- The validator asks why this method suits this book.

5 Assumptions, estimators and data

- Distribution: the shape given to returns, such as normal
- Independence: today says nothing about tomorrow
- Volatility estimator: equal weights or EWMA
- Correlation estimator: and is it stable?
- Data window: long enough, clean, few proxies

IN THE LESSON

- Next come the assumptions behind the method.
- The distributional assumption is the shape the model gives to returns, for example normal.
- The independence assumption is that today's return tells you nothing about tomorrow's.
- Then the estimators.
- How is volatility measured: a simple average over a window, or an exponentially weighted moving average, written E W M A, which gives recent days more weight?
- How are correlations estimated, and are they stable?
- Last, the data window.
- Is it long enough to include stressed periods, and is the data clean, with no gaps, stale prices or proxies standing in for missing series?

6 Tie each assumption to where it breaks

The exam asks which book a given assumption fails for

Assumption	Breaks for
Normal returns	Fat tails: short options, credit
Independence	Trends, clustering, illiquid books
Stable correlations	Crisis: hedges fail together
Calm data window	A return of stress
Linear approximation	Options: value curves

IN THE LESSON

- For the exam, tie each assumption to the portfolio where it breaks.
- Normal returns break for a book with fat tails, such as short options or credit, where big moves come more often than the normal predicts.
- Independence breaks when returns trend or cluster, which matters for longer horizons and for illiquid positions.
- Stable correlations break in a crisis, when correlations usually rise and hedges and diversification fail together.
- A short or calm data window breaks when stress returns, because the model has never seen a bad day.
- And a linear approximation breaks for options, whose value curves as the price moves.

7 Example 1: a risk factor that is missing

Corporate bonds hedged with Treasury futures. Factors must cover spreads, basis, volatility

Model: Treasury factor only	net rate DV01 near zero	**VaR ≈ $0**
Missing: credit spread risk	spread DV01	**$40,000/bp**
Daily spread volatility	spread moves a day	**4 bp**
99% spread move	2.326 × 4	**9.304 bp**
Hidden spread VaR	40,000 × 9.304	**$372,160**

IN THE LESSON

- Kestrel's bond desk holds corporate bonds and hedges the interest rate risk by selling Treasury futures.
- Risk factors must cover every material exposure, including spreads, basis and volatility.
- Yet the model maps both legs onto the Treasury curve only.
- Net interest rate sensitivity is close to zero, so the model reports a V a R close to zero.
- But the corporate bonds also carry credit spread risk.
- Their spread D V zero one, the loss when the spread widens by one basis point, is forty thousand dollars.
- The spread moves four basis points a day.
- At ninety nine per cent the move is two point three two six times four, nine point three zero four basis points, and the loss is about three hundred and seventy two thousand dollars.

8 Valuing options: full or approximate?

ΔP ≈ δ × ΔS + ½ × Γ × ΔS^2
Delta only: drop the second term

Full revaluation	Reprice exactly: accurate, slow
Delta, δ	Value change per $1 move
Gamma, Γ	How fast delta changes
Price move, ΔS	Move in the underlying price

IN THE LESSON

- The validator also checks how nonlinear positions, such as options, are valued.
- A fast way approximates the change in value.
- The delta gamma approximation says the change is about delta times the price move, plus one half times gamma times the price move squared.
- The delta approximation drops the second term.
- The best way is full revaluation, which reprices every option exactly in every scenario.
- It is accurate but slow.
- Delta, written with the small Greek letter delta, is the change in value for a one dollar move in the underlying price.
- Gamma, the capital Greek letter, is how fast delta itself changes.
- And the price move is the move in the underlying price.
- For a large move or a complex option, only full revaluation is reliable.

9 Example 2: what a delta-only model misses

Delta +10,000 shares, gamma −2,000 per $. Price $100, daily σ 2%

99% price move	2.326 × $2	**−$4.652**
Delta term	10,000 × 4.652	**−$46,520**
Gamma term	½ × 2,000 × 4.652^2	**−$21,641**
Delta-gamma loss	46,520 + 21,641	**$68,161**
Missed by delta only	21,641 ÷ 68,161	**31.7%**

IN THE LESSON

- Kestrel's options desk has a delta of ten thousand shares, but it has sold options, so its gamma is minus two thousand per dollar.
- The share price is one hundred dollars with a daily volatility of two per cent.
- A ninety nine per cent move is two point three two six times two dollars, a fall of four point six five two dollars.
- The delta term loses ten thousand times that, forty six thousand five hundred and twenty dollars.
- The gamma term adds a further loss of one half times two thousand times the move squared, about twenty one thousand six hundred and forty one dollars.
- The delta gamma loss is about sixty eight thousand one hundred and sixty one dollars.
- A delta only model misses almost thirty two per cent of it.

10 Holding period, confidence and scaling

VaR(T days) = VaR(1 day) × √T
Holds only with independent returns, constant positions

Horizon T, in days	Horizon, such as 10 days
Independent returns	No trends or clustering
Constant positions	Book not changed or rehedged
Confidence level	Higher: deeper tail, less data

IN THE LESSON

- The validator then checks the holding period and the confidence level, and any scaling between horizons.
- Many banks compute a one day V a R and scale it to ten days by multiplying by the square root of time.
- The rule says V a R over T days equals one day V a R times the square root of T.
- It holds only under conditions.
- Returns must be independent from day to day, with no trends or clustering.
- And the positions must be held constant over the horizon, which an options book that is rehedged every day does not do.
- The confidence level matters too: the higher it is, the further into the tail the model must reach, where it has least data.

11 Example 3: when scaling understates

1-day 99% VaR $2.61m. Daily returns: autocorrelation 0.2

True 2-day factor	√(2 × (1 + 0.2)) = √2.4	**1.549**
Square-root 2-day factor	√2	**1.414**
2-day VaR: rule vs true	2.61 × 1.414; × 1.549	**$3.69m; $4.04m**
Rule understates by	1 − 1.414 ÷ 1.549	**8.7%**
10-day VaR: rule vs true	2.61 × √10; × √14.375	**$8.25m; $9.9m**

IN THE LESSON

- Kestrel's one day ninety nine per cent V a R is two point six one million dollars, and suppose daily returns have an autocorrelation of point two, a correlation between one day's return and the next.
- Returns that trend break the square root of time rule.
- Over two days the variance is not two times the daily variance, but two times one point two.
- So the true two day factor is the square root of two point four, about one point five four nine, against the square root of two, one point four one four.
- The scaled V a R of three point six nine million understates the true four point zero four million by about eight point seven per cent.
- Over ten days, the rule gives two point six one times the square root of ten, about eight point two five million, but the true variance factor is fourteen point three seven five, so the true figure is about nine point nine million, and the rule understates it by about seventeen per cent.

12 Example 4: a clean backtest that hides two flaws

Simplified: desk VaRs add. Model versus true one-day 99% VaR, $m

Rates desk: conservative	model 2.0, true 1.5	**+0.5**
Options desk: ignores gamma	model 1.0, true 1.5	**−0.5**
Firm total	2.0 + 1.0 = 1.5 + 1.5	**3.0 both**
Expected exceptions	250 × 1%	**2.5**
Verdict	firm-level backtest misses them	**looks clean**

IN THE LESSON

- To keep it simple, suppose the desk figures add up, and compare each model with the true figure.
- This shows why a clean backtest is not proof of soundness.
- Kestrel's rates desk model is too conservative: it reports a V a R of two million dollars when the true figure is one and a half.
- The options desk model ignores gamma: it reports one million against a true one and a half.
- The firm total is three million either way, so the errors offset.
- Over two hundred and fifty days at ninety nine per cent, the expected number of exceptions, days when the loss beats V a R, is two and a half.
- The firm sees about that many, and the backtest looks clean.
- A firm level backtest cannot see the two flaws; a review of the design, or testing desk by desk, finds them.

13 Sensitivity analysis: one input at a time

- Rerun VaR, changing one input at a time
- Record how far VaR moves
- Inputs: confidence level, holding period, data window
- EWMA decay factor, correlations, volatilities
- Large moves show which assumptions drive VaR

IN THE LESSON

- The second tool is sensitivity analysis.
- The analyst reruns the V a R model while changing one input or assumption at a time, and records how far V a R moves.
- Typical inputs are the confidence level, the holding period, the length of the data window, the E W M A decay factor, which sets how fast old data loses weight, the correlation matrix and the volatility estimates.
- The result is a list of movements, and the large ones show which assumptions drive the number.
- To see how it works, we first need Kestrel's base V a R, computed with the simplest method, the parametric one, so that every change can be worked by hand.

14 Kestrel's VaR: two risk factors

σ$ = √(d1^2 + d2^2 + 2ρ × d1 × d2)
VaR = z × σ$

Dollar volatilities d1, d2	Position × daily volatility
Correlation, ρ	Correlation of the two returns
Multiplier, z	2.326 at 99%, one-tailed

IN THE LESSON

- Kestrel's trading book has two risk factors, equities and credit.
- For each one, the dollar volatility, written d, is the position times its daily volatility.
- The portfolio's dollar volatility, sigma dollar, is the square root of d one squared, plus d two squared, plus two times rho times d one times d two.
- Rho is the correlation between the two returns.
- V a R is then z times sigma dollar.
- Here z is the number of standard deviations for the confidence level.
- At ninety nine per cent, one tailed, it is two point three two six, rounded.
- This assumes normal returns with a mean of zero, which is the very kind of assumption sensitivity analysis will test.

15 Example 5: Kestrel's base VaR

Equities $60m, daily σ 1.5%. Credit $40m, daily σ 1.0%. ρ = 0.4

Equity dollar volatility	60 × 1.5%	**$0.90m**
Credit dollar volatility	40 × 1.0%	**$0.40m**
Variance	0.81 + 0.16 + 2×0.4×0.9×0.4	**1.258**
Portfolio σ$	√1.258	**$1.122m**
1-day 99% VaR	2.326 × 1.122	**$2.61m**

IN THE LESSON

- Here are the numbers.
- Kestrel holds sixty million dollars of equities with a daily volatility of one and a half per cent, and forty million of credit with a daily volatility of one per cent.
- The correlation is point four.
- The equity dollar volatility is sixty times one and a half per cent, nine hundred thousand dollars.
- The credit dollar volatility is forty times one per cent, four hundred thousand.
- The variance is point eight one plus point one six plus two times point four times point nine times point four, which is one point two five eight.
- Its square root is about one point one two two million dollars.
- Multiply by two point three two six, and the one day, ninety nine per cent V a R is about two point six one million dollars.

16 Example 6: moving one input at a time

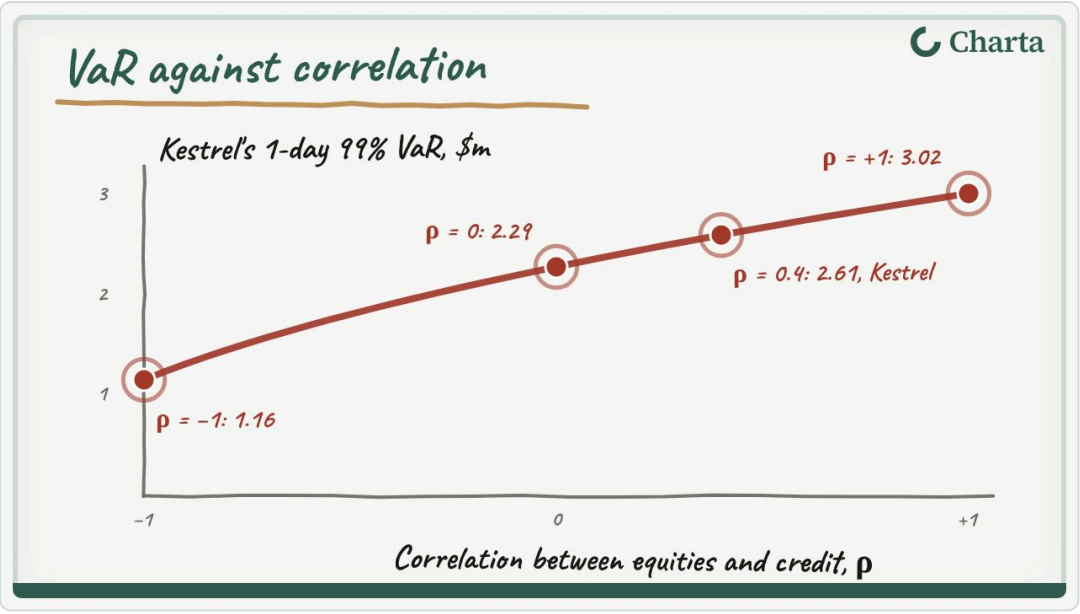
Base: 1-day 99% VaR $2.61m. Change one input, keep the rest

Confidence 97.5% (z = 1.96)	1.96 × 1.122	**$2.20m, −16%**
Holding period 10 days	2.61 × √10	**$8.25m, ×3.16**
Correlation 0.4 → 0.8	σ$ = √1.546 = 1.243	**$2.89m, +11%**
Equity volatility 1.5% → 2%	σ$ = √1.984 = 1.409	**$3.28m, +26%**

IN THE LESSON

- Priya now moves one input at a time against the base of two point six one million.
- Lower the confidence level to ninety seven and a half per cent, where z is one point nine six, and V a R falls to two point two zero million, down about sixteen per cent.
- Lengthen the holding period to ten days, and the square root of time rule takes it to eight point two five, three point one six times larger.
- Raise the correlation from point four to point eight; sigma dollar becomes the square root of one point five four six, one point two four three, and V a R rises to two point eight nine, up about eleven per cent.
- Raise the equity volatility from one and a half to two per cent, as a shorter window might; sigma dollar becomes the square root of one point nine eight four, one point four zero nine, and V a R rises to three point two eight, up about twenty six per cent.

17 VaR against correlation



IN THE LESSON

- A sensitivity can also be drawn as a curve.
- Across the bottom is the correlation between equities and credit, from minus one to plus one.
- Up the side is Kestrel's V a R, with everything else held at the base case.
- At a correlation of plus one, on the right, V a R is three point zero two million, simply the two stand alone figures added.
- At Kestrel's estimate of point four it is two point six one.
- At zero it is two point two nine.
- At minus one it falls to one point one six million.
- The curve is not a straight line, so the effect of a change in correlation depends on where you start, and the correlation itself is an estimate with no single correct value.

18 Example 7: two inputs move together

Base $2.61m. Equity σ 1.5% → 2% and ρ 0.4 → 0.8, alone and together

Equity volatility alone	3.28 − 2.61	**+$0.67m**
Correlation alone	2.89 − 2.61	**+$0.28m**
Sum of separate effects	2.61 + 0.67 + 0.28	**$3.56m**
Both together	2.326 × √2.368 = 2.326 × 1.539	**$3.58m**
Interaction missed	3.58 − 3.56	**$0.02m**

VaR is not additive across inputs

IN THE LESSON

- In a crisis, volatility and correlation tend to rise together, so Priya also moves both at once.
- Raising equity volatility to two per cent alone added three point two eight minus two point six one, point six seven million.
- Raising correlation to point eight alone added point two eight million.
- Adding the two separate effects to the base predicts two point six one plus point six seven plus point two eight, three point five six million.
- But with both changes together, sigma dollar is the square root of two point three six eight, about one point five three nine, and V a R is three point five eight million.
- The interaction, about twenty thousand dollars here, is what one at a time testing misses.
- In a book with options and hedges it can be far larger.
- So V a R is not additive across inputs.

19 The benefits of sensitivity analysis

- Shows which assumptions drive the number
- Shows where model risk is concentrated
- Shows fragility to choices with no single right value
- Helps users understand what moves VaR
- Flags inputs that need better data or monitoring

IN THE LESSON

- What does sensitivity analysis buy?
- First, it shows which assumptions drive the number.
- At Kestrel, the volatility estimate and the holding period matter far more than a modest change in correlation.
- Second, it shows where model risk is concentrated, so validation effort and model reserves can go there.
- Third, it shows how fragile the result is to choices that have no single correct value, such as the window length or the decay factor.
- Fourth, it helps people who use the number understand what moves it, and it flags inputs that deserve better data or closer monitoring.

20 The difficulties

- Many parameters: which ones to move?
- Plausible ranges must be judged
- No objective threshold for an acceptable change
- Effects interact; VaR is not additive across inputs
- Shows dependence, not correctness; not a backtest

IN THE LESSON

- Its difficulties are practical.
- A real model has many parameters, often hundreds, so the analyst must choose which to move.
- For each one, a plausible range must be judged, and that judgement is itself a choice.
- There is no objective threshold for an acceptable change: is a sixteen per cent move a problem or not?
- Effects interact, so moving inputs one at a time can miss combined shocks, as Example seven showed, and V a R is not additive across inputs.
- Above all, sensitivity analysis shows dependence, not correctness.
- It tells you how much the number leans on an assumption, not whether that assumption is right, so it does not replace backtesting.

21 VaR is an estimate, not a fact

- VaR is estimated from a finite sample
- So it carries estimation error
- Ideally reported with a confidence interval
- But VaR sits in the tail, where data is scarce
- Exam: VaR is a point estimate with sampling error

IN THE LESSON

- The third tool asks how precise V a R is.
- A V a R figure is a statistic estimated from a finite sample of data, so it carries estimation error, just like an average computed from a poll.
- Draw a different year of history and you get a different number.
- Ideally, a confidence interval should accompany it: a range that, under the model, would contain the true V a R with a stated probability, such as ninety five per cent.
- The difficulty is that V a R is a quantile far out in the tail, and the tail is exactly where data is scarce.
- For the exam, remember that V a R is a point estimate with sampling error.

22 The standard error of a quantile

SE ≈ √(p × (1 − p) ÷ n) ÷ f(q)
95% interval: VaR ± 1.96 × SE

Tail probability, p	1% for a 99% VaR
Sample size, n	Number of observations
Density at the quantile, f(q)	Small far in the tail
95% interval	VaR ± 1.96 standard errors

IN THE LESSON

- For a V a R read straight from the data, as in historical simulation, there is an approximate formula for the standard error.
- Let p be the tail probability, one per cent at ninety nine per cent confidence.
- Let n be the sample size, the number of observations.
- Let f of q be the density of the profit and loss distribution at the V a R quantile, meaning how thickly observations are packed there.
- The standard error is the square root of p times one minus p, over n, all divided by f of q.
- Far in the tail, f of q is small, so the error is large.
- And because n sits under a square root, the error narrows only slowly.
- A ninety five per cent interval is then V a R plus or minus one point nine six standard errors.

23 Example 8: an interval around Kestrel's VaR

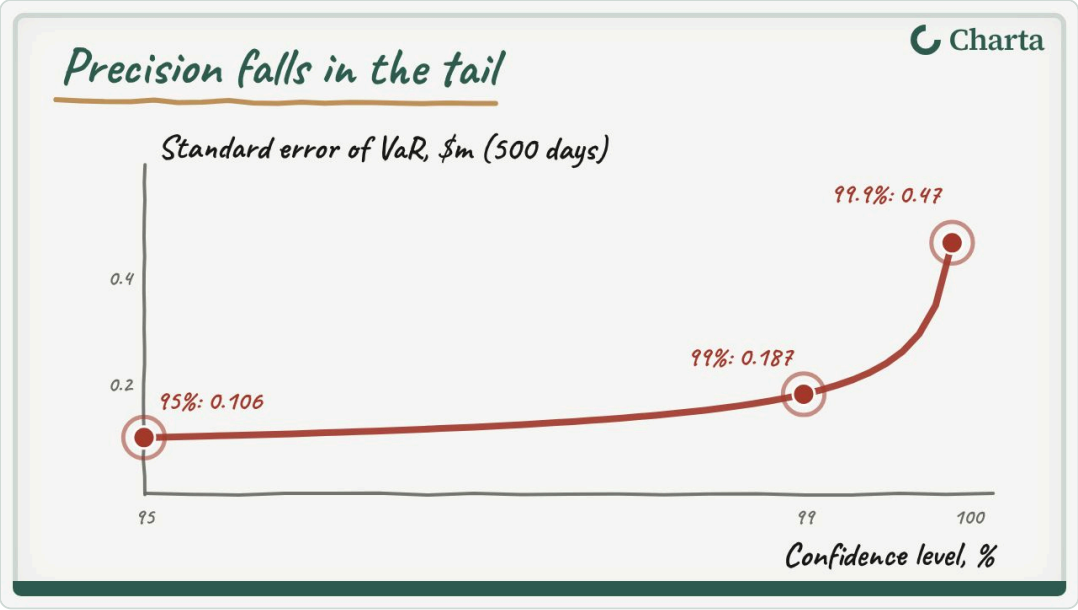
Historical simulation: 99% VaR from 500 days. Assume normal P&L, σ = $1.122m

Days beyond VaR	1% × 500	**5**
√(p(1 − p) ÷ n)	√(0.01 × 0.99 ÷ 500)	**0.00445**
Density at the quantile	0.0267 ÷ 1.122	**0.0238 per $m**
Standard error	0.00445 ÷ 0.0238	**$0.187m**
95% interval	2.61 ± 1.96 × 0.187	**$2.24m to $2.98m**

IN THE LESSON

- Suppose instead that Kestrel uses historical simulation, reading its ninety nine per cent V a R straight from five hundred days of data.
- Assume that its daily profit and loss is normal with a standard deviation of one point one two two million dollars, so the quantile sits near two point six one million.
- Only one per cent of five hundred, just five observations, lie beyond the V a R.
- The square root of point zero one times point nine nine over five hundred is about point zero zero four four five.
- The normal density at two point three two six standard deviations is point zero two six seven.
- Divided by one point one two two million, that gives f of q of about point zero two three eight per million.
- The standard error is about one hundred and eighty seven thousand dollars, and the ninety five per cent interval runs from about two point two four to two point nine eight million.

24 Precision falls in the tail



IN THE LESSON

- This chart shows how the standard error grows as the confidence level rises, for Kestrel with five hundred days of data.
- Across the bottom is the confidence level, from ninety five to one hundred per cent.
- Up the side is the standard error, in millions of dollars.
- At ninety five per cent the standard error is about one hundred and six thousand dollars.
- At ninety nine per cent it is one hundred and eighty seven thousand.
- At ninety nine point nine per cent it is about four hundred and seventy thousand, more than four times the figure at ninety five.
- The curve climbs steeply as we move into the tail, because fewer and fewer observations support the estimate.
- A higher confidence level gives a wider interval.

25 Example 9: what more precision costs

Kestrel: 99% VaR $2.61m, SE $0.187m from 500 days

250 days of data	0.187 × √2	**SE $0.26m**
2,000 days: four times the data	0.187 ÷ 2	**SE $0.094m**
99.9% VaR from 500 days	3.09 × 1.122	**$3.47m**
Its standard error	0.00141 ÷ 0.00300	**$0.47m, ≈ 14%**
10-day periods that do not overlap	500 ÷ 10	**about 50**

IN THE LESSON

- Kestrel's standard error is one hundred and eighty seven thousand dollars from five hundred days.
- How could it be narrowed?
- With only two hundred and fifty days, one year, the standard error grows by the square root of two, to about point two six million dollars.
- To halve the five hundred day figure, to about ninety four thousand, Kestrel would need four times the data, two thousand days, roughly eight years, over which the market and the book will have changed.
- Going further into the tail makes it worse: with z of three point zero nine, the ninety nine point nine per cent V a R is about three point four seven million.
- The root term is point zero zero one four one, the density there is only point zero zero three zero zero per million, and the standard error is about point four seven million, close to fourteen per cent of the estimate.
- Overlapping ten day returns add little: five hundred days hold only about fifty ten day periods that do not overlap.

26 Why intervals are hard in practice

- Analytical intervals assume a distribution
- Clustering and overlapping returns shrink the sample
- Parameters change: more data adds old regimes
- Bootstrap and Monte Carlo: heavy computing
- And they still carry the model's own assumptions

IN THE LESSON

- In practice there are deeper difficulties.
- Analytical intervals, like our formula, rely on distributional assumptions, here a normal density at the quantile, which may be wrong exactly in the tail.
- Volatility clustering and overlapping returns mean the observations are not independent, which shrinks the effective sample.
- Parameters change over time, so a longer window adds data but also adds old regimes.
- Bootstrap intervals, which resample the data many times, and Monte Carlo intervals avoid the formula but are computationally heavy, and they still reflect the model's own assumptions, so they cannot measure model risk.
- So an interval is useful, but it is never exact.

27 Benchmarking

- Compare with an independent alternative
- A different method: historical simulation vs parametric
- A vendor model, bought from an outside firm
- Or a simpler model with few risk factors
- Compare final VaR, or inputs and components

IN THE LESSON

- The fourth tool is benchmarking.
- It compares the V a R model, and its inputs and components, with an independent alternative, to see whether the results are broadly consistent.
- The alternative can be a different methodology, for example historical simulation against a parametric model.
- It can be a vendor model, bought from an outside firm.
- Or it can be a simpler model, such as one built on a handful of risk factors.
- Benchmarking can be done on the final V a R, or piece by piece, on the volatilities, the correlations or the valuation of a single instrument.
- Large, unexplained gaps send the validator back to the design.

28 Difficulties and remedies

Agreement shows consistency, not accuracy

Difficulty	Remedy
No true benchmark	Decompose: factor, vol, correlation
Different level, horizon, data	Align level, horizon, portfolio first
Benchmark shares flaws, unvalidated	Use more than one benchmark
Hard to compare in a big book	Test simple hypothetical portfolios

IN THE LESSON

- Agreement between two models shows consistency, not accuracy.
- With that in mind, benchmarking has well known difficulties, and each has a remedy.
- First, no true benchmark exists, so a gap does not say which model is wrong.
- The remedy is to decompose the difference by risk factor, volatility and correlation, to find its source.
- Second, the models may differ in confidence level, horizon, data, risk factors and valuation.
- So align the confidence level, the horizon and the portfolio first.
- Third, the benchmark may share the same flaws, or be unvalidated itself.
- Use more than one benchmark.
- Fourth, components are hard to compare inside a large book.
- So test both models on simple hypothetical portfolios.

29 Example 10: align, then decompose

Vendor: 10-day 99% VaR $9.328m; equity σ 1.7%, ρ 0.5. Kestrel: 1-day $2.609m

Align the horizon	9.328 ÷ √10	**$2.950m**
True gap, both 1-day	2.950 − 2.609	**$0.341m**
Swap in vendor volatility	2.874 − 2.609	**$0.265m**
Swap in vendor correlation	2.950 − 2.874	**$0.076m**

The gap is mostly the equity volatility estimate

IN THE LESSON

- Kestrel compares its model with a vendor model.
- The vendor reports a ten day, ninety nine per cent V a R of nine point three two eight million dollars for the same book, against Kestrel's one day figure of two point six zero nine.
- The gap looks huge, but first align the horizon.
- Dividing by the square root of ten gives a one day figure of two point nine five zero million.
- The true gap is only point three four one million.
- Now decompose it.
- The vendor estimates equity volatility at one point seven per cent, and correlation at point five.
- Putting the vendor's volatility into Kestrel's model gives two point eight seven four, which explains point two six five million.
- Then the vendor's correlation explains the remaining point zero seven six.
- So the gap is mostly the equity volatility estimate.
- The split depends a little on the order of the swaps, but the conclusion does not.

30 Example 11: a simple hypothetical portfolio

Hypothetical book: $10m of equities only, so correlation plays no part

Kestrel's model	2.326 × 1.5% × 10m	**$348,900**
Vendor model	2.326 × 1.7% × 10m	**$395,420**
Gap: volatility only	395,420 − 348,900	**$46,520**

Agreement or a gap shows consistency, not which model is right

IN THE LESSON

- A simple hypothetical portfolio isolates one component.
- Priya feeds both models a book of ten million dollars of equities and nothing else.
- With one position, correlation plays no part, so any gap must come from volatility.
- Kestrel's model gives two point three two six times one and a half per cent times ten million, three hundred and forty eight thousand nine hundred dollars.
- The vendor's gives two point three two six times one point seven per cent times ten million, three hundred and ninety five thousand four hundred and twenty dollars.
- The gap of forty six thousand five hundred and twenty dollars is entirely the volatility estimate.
- Neither result says which volatility is right, so Priya also checks a second benchmark, and the backtest.

31 Mistakes to avoid

- A clean backtest is not proof of soundness
- Tie each assumption to the book where it breaks
- Sensitivity shows dependence, not correctness
- VaR is a point estimate with sampling error
- Align before benchmarking; agreement is not accuracy

IN THE LESSON

- Here are the mistakes to avoid.
- Do not treat a clean backtest as proof of soundness, because offsetting errors or a conservative design can hide flaws.
- Do not judge an assumption in general: tie it to the portfolio where it breaks.
- Do not read a sensitivity result as a verdict on correctness, since it shows only dependence.
- Do not report V a R as an exact figure, because it is a point estimate with sampling error that widens in the tail.
- And do not compare two models before aligning confidence level, horizon and portfolio, or treat agreement as accuracy.
- Benchmarking and sensitivity analysis supplement backtesting, and none of them replaces it.

32 Try three quick questions

1. Delta 5,000, gamma −1,000 per $. Price falls $3. Delta-gamma loss?
   Answer: 15,000 + ½ × 1,000 × 9 = $19,500.
2. 95% VaR from 475 days; f(q) = 0.1 per $m. Standard error?
   Answer: √(0.05 × 0.95 ÷ 475) = 0.01; SE = $0.1m.
3. Vendor 10-day 99% VaR $12.65m; yours, 1-day 99%, $4m. Disagree?
   Answer: No: 12.65 ÷ √10 = $4.0m once aligned.

IN THE LESSON

- Pause, think, then listen to the answers.
- Question one.
- A book has a delta of five thousand shares and a gamma of minus one thousand per dollar.
- The price falls three dollars.
- What is the delta gamma loss?
- Fifteen thousand from delta, plus one half times one thousand times nine, four thousand five hundred, from gamma: nineteen thousand five hundred dollars.
- Question two.
- A ninety five per cent V a R comes from four hundred and seventy five days, with a density at the quantile of point one per million.
- What is the standard error?
- The root of point zero five times point nine five over four hundred and seventy five is point zero one, so the standard error is point one million dollars.
- Question three.
- A vendor's ten day, ninety nine per cent V a R is twelve point six five million and your one day, ninety nine per cent V a R is four million.
- Do they disagree?
- No.
- Twelve point six five divided by the square root of ten is four million.