CHARTA LEARNING · LESSON NOTES

Estimating Market Risk Measures

FRM Part II

1 A story: how much could we lose tomorrow?

- Ashby Bank's desk: $15m shares + $5m bonds = $20m
- P&L: the change in value; gain +, loss −
- VaR: a loss exceeded only (1 − c) of the time
- 1-day 95% VaR: beaten on about 1 day in 20
- Expected shortfall: the average loss beyond VaR

IN THE LESSON

- Ashby Bank's trading desk holds fifteen million dollars of shares and five million dollars of bonds, twenty million in all.
- Profit and loss, written P and L, is the change in the desk's value over a horizon such as one day: a gain is positive, a loss negative.
- Value at risk, V a R, is a loss level, set at a confidence level c, that is exceeded only a fraction one minus c of the time.
- So a one day, ninety five per cent V a R is a loss we expect to be beaten on about one day in twenty.
- Expected shortfall is the average loss beyond V a R: how bad the tail gets.

2 Historical simulation: let the past speak

- No distribution assumed: nonparametric
- Past P&L is a sample of future P&L
- Apply each past day's returns to today's holdings
- Sort the results from worst to best
- Read off the loss only (1 − c) of outcomes exceed

IN THE LESSON

- The simplest method is historical simulation.
- It assumes no distribution at all, which is why it is called nonparametric.
- It treats past P and L as a sample of future P and L.
- Take each day in a past window, say the last five hundred trading days, and apply that day's asset returns to the holdings the desk has today.
- That gives five hundred possible outcomes for tomorrow.
- Sort them from the worst to the best.
- Then read off the loss that only a fraction one minus c of the outcomes exceed.
- That loss is the historical simulation V a R.

3 How many observations sit in the tail

Tail = n × (1 − c)
VaR = the Tail-th worst loss (a common convention)

Observations, n	Days in the window, e.g. 500
Confidence level, c	e.g. 95%; tail 1 − c = 5%
Tail count	500 × 0.05 = 25: the 25th worst

IN THE LESSON

- The tail holds n times one minus c observations.
- V a R is a single order statistic, not the whole tail: under a common convention, which we use throughout, it is the tail count's worst loss.
- Other conventions take the next one, the first outside the tail, or interpolate.
- Always state yours.
- Here n is the number of days in the window, say five hundred, and c the confidence level, say ninety five per cent, so the tail fraction, one minus c, is five per cent.
- The tail count is then five hundred times point zero five, twenty five, so V a R is the twenty fifth worst loss.
- And keep signs consistent: sort losses or P and L, never a mix of the two.

4 Example 1: building tomorrow's possible P&L

Today: $15m shares, $5m bonds. Apply each past day's returns to today's holdings

Day 312: shares −2.40%	−2.40% × $15m	**−$360,000**
Day 312: bonds +0.30%	+0.30% × $5m	**+$15,000**
Day 312 scenario P&L	−360,000 + 15,000	**−$345,000**
Day 88: +1.10% and −0.50%	+165,000 − 25,000	**+$140,000**
Repeat for all 500 days	not the desk's actual P&L	**then sort**

IN THE LESSON

- We apply each past day's returns to Ashby's holdings today.
- On day three hundred and twelve, shares fell two point four per cent: a loss of three hundred and sixty thousand dollars on today's fifteen million.
- Bonds rose point three per cent, a gain of fifteen thousand on five million.
- So the scenario P and L for that day is minus three hundred and forty five thousand dollars.
- On day eighty eight, shares rose one point one per cent and bonds fell half a per cent: plus one hundred and sixty five thousand, minus twenty five thousand, a gain of one hundred and forty thousand.
- Repeat for all five hundred days, then sort.
- We do not use the desk's actual P and L, which was earned on old positions.

5 Example 2: reading off the 95% VaR

500 losses sorted. 25th worst $398,000; 26th worst $385,000; 95th worst $171,000

Tail count at 95%	500 × (1 − 0.95)	**25**
VaR: the 25th-worst loss	our convention	**$398,000**
Other convention: 26th worst	first one outside the tail	**$385,000**
Or interpolate midway	(398,000 + 385,000) ÷ 2	**$391,500**
Wrong: the 95th-worst loss	confuses c with a rank	**$171,000**

IN THE LESSON

- Ashby's five hundred scenario losses are now sorted.
- The twenty fifth worst is three hundred and ninety eight thousand dollars, the twenty sixth worst three hundred and eighty five thousand, and the ninety fifth worst one hundred and seventy one thousand.
- At ninety five per cent the tail count is five hundred times one minus point nine five: twenty five.
- Under our convention, V a R is the twenty fifth worst loss, three hundred and ninety eight thousand dollars.
- Under the other convention, the twenty sixth worst, it is three hundred and eighty five thousand.
- Interpolating midway gives three hundred and ninety one thousand five hundred.
- A common error is the ninety fifth worst loss, one hundred and seventy one thousand, which confuses the confidence level with a rank.

6 Example 3: the 99% VaR, and a ceiling

Same 500 losses. Five worst: $1,240k, $980k, $870k, $760k, $650k

Tail count at 99%	500 × (1 − 0.99)	**5**
VaR: the 5th-worst loss	our convention	**$650,000**
Worst loss in the window	no estimate can go beyond it	**$1,240,000**
A 100-day window at 99%	100 × (1 − 0.99)	**1 observation**

IN THE LESSON

- Among the same five hundred losses, the five worst are one million two hundred and forty thousand, nine hundred and eighty thousand, eight hundred and seventy thousand, seven hundred and sixty thousand and six hundred and fifty thousand dollars.
- At ninety nine per cent the tail count is five hundred times point zero one, just five, so V a R is the fifth worst loss, six hundred and fifty thousand dollars.
- Notice the ceiling.
- The worst loss in the window is one million two hundred and forty thousand, and historical simulation can never estimate beyond it.
- With only one hundred days at ninety nine per cent, the tail would hold a single observation.
- Historical simulation depends entirely on how many tail observations the window holds.

7 Parametric VaR under a normal

VaR = −μ + z × σ
Over T days: μ × T and σ × √T

Normal quantile, z	1.645 at 95%, 2.326 at 99%
Mean P&L, μ	A positive mean lowers VaR
Standard deviation, σ	Of P&L, in dollars
Horizon, T days	Needs independent daily returns

IN THE LESSON

- The second approach is parametric.
- We assume a distribution for P and L and plug in estimated parameters.
- If P and L is normal with mean mu and standard deviation sigma, V a R equals minus mu plus z times sigma.
- Over a horizon of T days, use the mean times T and sigma times the square root of T.
- Here z is the standard normal quantile: one point six four five at ninety five per cent and two point three two six at ninety nine per cent.
- Because mu enters with a minus sign, a positive expected gain reduces V a R.
- Sigma is the standard deviation of P and L in dollars.
- And the square root rule assumes that daily returns are independent.

8 Example 4: Ashby's one-day normal VaR

Desk $20m; daily mean return 0.04%, daily volatility 1.25%; P&L assumed normal

Mean P&L, μ	0.04% × $20m	**$8,000**
Standard deviation, σ	1.25% × $20m	**$250,000**
95% VaR	1.645 × 250,000 − 8,000	**$403,250**
99% VaR	2.326 × 250,000 − 8,000	**$573,500**
Wrong: adding the mean, 95%	1.645 × 250,000 + 8,000	**$419,250**

IN THE LESSON

- Ashby's desk has a daily mean return of point zero four per cent and daily volatility of one point two five per cent, and we assume P and L is normal.
- The mean P and L is point zero four per cent of twenty million, eight thousand dollars.
- The standard deviation is one point two five per cent of twenty million, two hundred and fifty thousand.
- The ninety five per cent V a R is one point six four five times two hundred and fifty thousand, minus eight thousand: four hundred and three thousand two hundred and fifty dollars.
- At ninety nine per cent it is two point three two six times two hundred and fifty thousand, minus eight thousand: five hundred and seventy three thousand five hundred.
- Adding the mean instead, one point six four five times two hundred and fifty thousand plus eight thousand, wrongly gives four hundred and nineteen thousand two hundred and fifty.

9 Example 5: a ten-day horizon

Same desk; 10-day 99% VaR; daily returns independent

10-day mean	8,000 × 10	**$80,000**
10-day standard deviation	250,000 × √10	**$790,569**
10-day 99% VaR	2.326 × 790,569 − 80,000	**$1,758,864**
Wrong: σ scaled by 10	2.326 × 2,500,000 − 80,000	**$5,735,000**

IN THE LESSON

- Now find a ten day, ninety nine per cent V a R, assuming independent daily returns.
- The mean scales with time: eight thousand times ten, eighty thousand dollars.
- The standard deviation scales with the square root of time: two hundred and fifty thousand times root ten, about seven hundred and ninety thousand five hundred and sixty nine.
- V a R is two point three two six times that, minus eighty thousand: about one million seven hundred and fifty eight thousand eight hundred and sixty four dollars.
- Scaling volatility by ten instead gives two point three two six times two and a half million, minus eighty thousand: five million seven hundred and thirty five thousand, over three times too big.

10 Parametric VaR under a lognormal

VaR = V × (1 − e^{μ − z × σ})
Never above V; below V × (z × σ − μ) for the same μ, σ

Position value, V	Value of the position today
Mean log return, μ	Mean log return over the horizon
SD of log return, σ	Log return SD over the horizon

IN THE LESSON

- A normal P and L allows a loss larger than the whole position, since a normal has no lower limit.
- If instead the price is lognormal, the log return over the horizon is normal, and V a R equals the value V times one minus e to the power mu minus z sigma.
- Because e to any power is positive, this V a R can never exceed the position value.
- With the same mu and sigma and a positive move, it is below the normal approximation, V times z sigma minus mu, because the exponential falls short of the straight line.
- Here V is the value of the position today, and mu and sigma are the mean and standard deviation of the log return over the horizon.

11 Example 6: Ashby's one-day lognormal VaR

Desk $20m; daily log return: mean 0.04%, SD 1.25%; 99% confidence

Exponent: μ − z × σ	0.0004 − 2.326 × 0.0125	**−0.028675**
Lognormal VaR	$20m × (1 − e^{−0.028675})	**$565,355**
Normal VaR, same μ and σ	$20m × 0.028675	**$573,500**
Lognormal is lower by	573,500 − 565,355	**$8,145**

Over 1 day they barely differ; over long horizons the gap is large

IN THE LESSON

- Now assume Ashby's prices are lognormal, with the daily log return having the same mean, point zero four per cent, and standard deviation, one point two five per cent.
- At ninety nine per cent the exponent is point zero zero zero four minus two point three two six times point zero one two five: minus point zero two eight six seven five.
- Lognormal V a R is twenty million times one minus e to that power, about five hundred and sixty five thousand three hundred and fifty five dollars.
- The normal V a R with the same mu and sigma is twenty million times point zero two eight six seven five, five hundred and seventy three thousand five hundred.
- The lognormal is lower by eight thousand one hundred and forty five.
- Over one day they barely differ; over long horizons the gap is large.

12 Example 7: a volatile stake held for a year

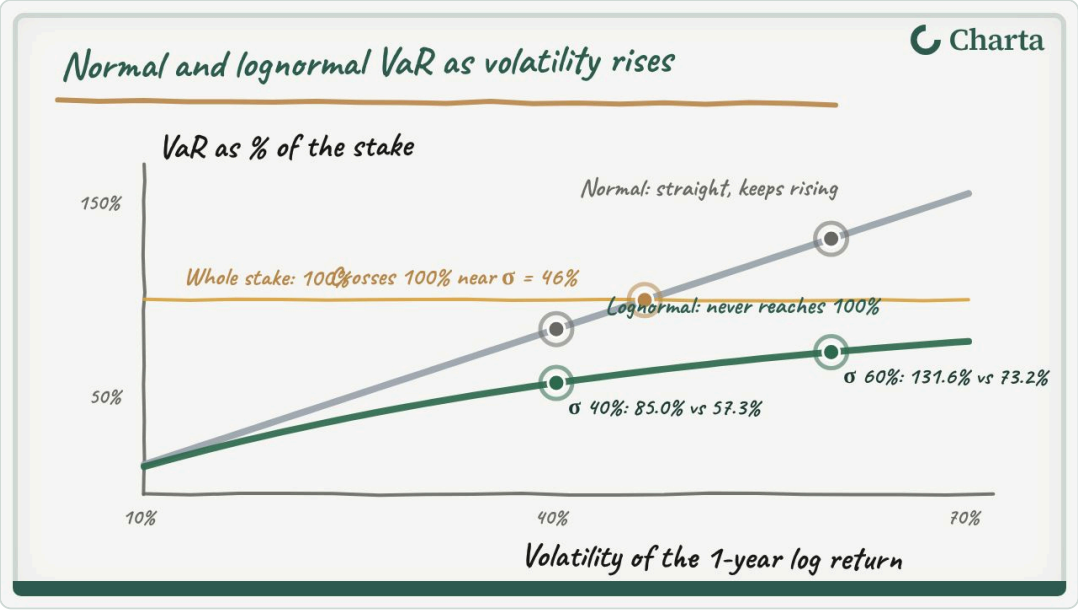
$5m stake in one share; 1-year log return: mean 8%, SD 40%; 99% confidence

Move: z × σ − μ	2.326 × 0.40 − 0.08	**0.8504**
Normal VaR	$5m × 0.8504	**$4,252,000**
Lognormal VaR	$5m × (1 − e^{−0.8504})	**$2,863,780**
If σ were 60%: normal	$5m × (2.326 × 0.6 − 0.08)	**$6,578,000**
If σ were 60%: lognormal	$5m × (1 − e^{−1.3156})	**$3,658,434**

IN THE LESSON

- Take a five million dollar stake in one volatile share, over one year, at ninety nine per cent.
- Its log return has a mean of eight per cent and a standard deviation of forty per cent.
- The move, z sigma minus mu, is two point three two six times point four, minus point zero eight: point eight five zero four.
- Normal V a R is five million times that, four million two hundred and fifty two thousand dollars.
- Lognormal V a R is five million times one minus e to the minus point eight five zero four, about two million eight hundred and sixty three thousand seven hundred and eighty.
- At sixty per cent volatility, the normal gives six million five hundred and seventy eight thousand, more than the stake is worth.
- The lognormal gives three million six hundred and fifty eight thousand four hundred and thirty four.

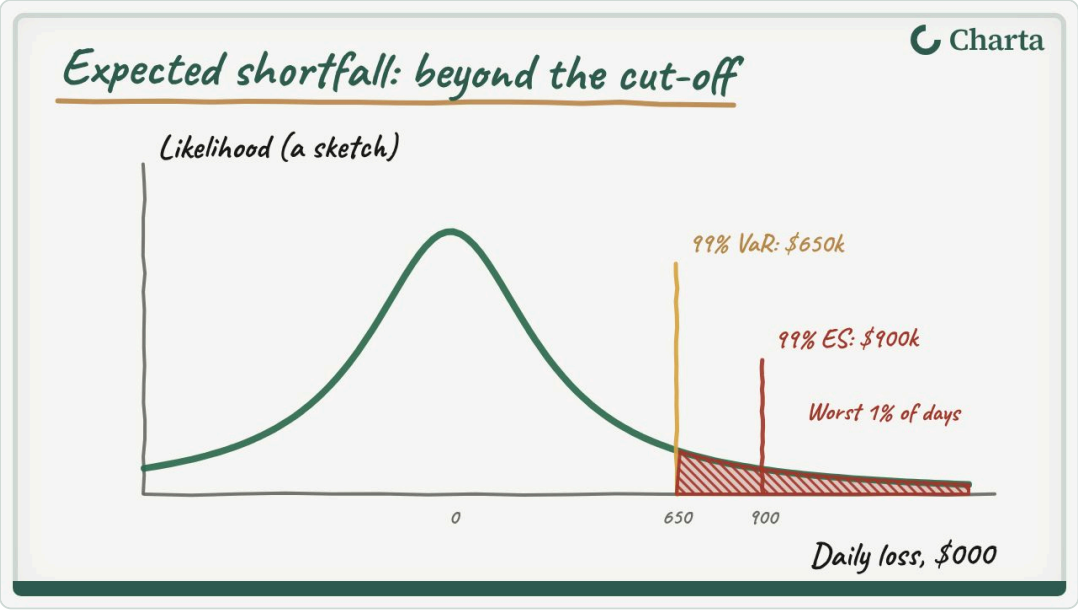
13 Normal and lognormal VaR as volatility rises



IN THE LESSON

- This chart shows V a R as a share of the stake, with the mean fixed at eight per cent, as volatility rises from ten to seventy per cent.
- The grey line is the normal formula: straight, and it keeps rising.
- The gold line marks one hundred per cent, the whole stake, and the normal line crosses it near a volatility of forty six per cent, impossible for a long position.
- The green curve is the lognormal formula.
- It flattens and never reaches one hundred per cent.
- At forty per cent volatility the two give eighty five and fifty seven point three per cent.
- At sixty per cent, one hundred and thirty one point six and seventy three point two per cent.
- Match the formula to the distribution.

14 Expected shortfall: beyond the cut-off



IN THE LESSON

- V a R tells us where the bad tail starts.
- Expected shortfall, E S, tells us how bad it is inside.
- Here is a sketch of Ashby's daily loss distribution, with losses to the right.
- The gold line is the ninety nine per cent V a R from historical simulation, six hundred and fifty thousand dollars.
- The shaded area beyond it holds the worst one per cent of days.
- Expected shortfall is the average loss in that area: the expected loss, given that the loss exceeds V a R.
- For Ashby it is nine hundred thousand dollars, the red line further out.
- Because it averages losses beyond the cut-off, expected shortfall is always at least as large as V a R.

15 Expected shortfall from a sample

Tail count k = n × (1 − c)
ES = (sum of the k worst losses) ÷ k
k = 2.5: ES = (loss 1 + loss 2 + 0.5 × loss 3) ÷ 2.5

Tail count, k	How many losses to average
Sorted losses, worst first	Positive losses, worst first
Expected shortfall	At least as large as VaR

IN THE LESSON

- From a sample, expected shortfall takes three steps.
- Count the tail: k equals n times one minus c.
- Sort the losses from worst to best and average the k worst.
- If k is not a whole number, include the boundary observation fractionally: with k of two and a half, take the two worst losses plus half of the third, and divide by two and a half.
- That is the convention we use; others round or interpolate.
- Here k is how many losses to average, the sorted losses, loss one, loss two and so on, are positive numbers, worst first, and the result is at least as large as V a R.
- For example, two hundred observations at ninety five per cent give k of ten, so E S is the mean of the ten largest losses.

16 Example 8: Ashby's 99% expected shortfall

500 days at 99%. Five worst returns on $20m: −6.20, −4.90, −4.35, −3.80, −3.25%

Tail count	500 × (1 − 0.99)	**5**
Losses = returns × $20m, $000	1,240+980+870+760+650	**4,500**
ES: average of the five	4,500,000 ÷ 5	**$900,000**
VaR: the 5th-worst loss	from example 3	**$650,000**
ES minus VaR	900,000 − 650,000	**$250,000**

IN THE LESSON

- Ashby's five worst returns over the five hundred days are minus six point two, minus four point nine, minus four point three five, minus three point eight and minus three point two five per cent.
- At ninety nine per cent the tail count is five hundred times point zero one, five.
- Convert returns to money by multiplying by twenty million: one million two hundred and forty thousand, nine hundred and eighty thousand, eight hundred and seventy thousand, seven hundred and sixty thousand and six hundred and fifty thousand, which add to four million five hundred thousand.
- Expected shortfall is their average, nine hundred thousand dollars.
- V a R is the fifth worst loss, six hundred and fifty thousand, from example three.
- E S is two hundred and fifty thousand larger.

17 Example 9: a tail count that is not whole

Last 250 days only, 99%. Three worst losses: $980k, $870k, $650k

Tail count	250 × (1 − 0.99)	**2.5**
Two worst in full	980,000 + 870,000	**$1,850,000**
Half of the third	0.5 × 650,000	**$325,000**
ES	2,175,000 ÷ 2.5	**$870,000**
Wrong: average of the 2 worst	1,850,000 ÷ 2	**$925,000**

IN THE LESSON

- Suppose the team uses only the last two hundred and fifty days.
- The three worst losses in that window are nine hundred and eighty thousand, eight hundred and seventy thousand and six hundred and fifty thousand dollars.
- At ninety nine per cent the tail count is two hundred and fifty times point zero one: two point five, not a whole number.
- Take the two worst losses in full, one million eight hundred and fifty thousand.
- Add half of the third, three hundred and twenty five thousand.
- That makes two million one hundred and seventy five thousand, and dividing by two point five gives an expected shortfall of eight hundred and seventy thousand dollars.
- Averaging just the two worst would give nine hundred and twenty five thousand, which miscounts the tail.

18 Example 10: ES sees how bad the tail is

What if Ashby's worst day had lost $2,480k instead of $1,240k?

99% VaR: the 5th-worst loss	unchanged	**$650,000**
Sum of the five worst	4,500,000 + 1,240,000	**$5,740,000**
99% ES	5,740,000 ÷ 5	**$1,148,000**
Rise in ES	1,148,000 − 900,000	**$248,000**

ES sees tail severity, VaR does not. Count the tail; ES is not VaR

IN THE LESSON

- Suppose Ashby's worst day had lost twice as much: two million four hundred and eighty thousand dollars instead of one million two hundred and forty thousand.
- The ninety nine per cent V a R is still the fifth worst loss, six hundred and fifty thousand dollars, unchanged.
- The sum of the five worst losses rises by one million two hundred and forty thousand, to five million seven hundred and forty thousand.
- Expected shortfall becomes one million one hundred and forty eight thousand dollars, a rise of two hundred and forty eight thousand.
- Expected shortfall responds to the severity of tail losses, and V a R does not.
- Avoid two errors: averaging the wrong number of observations, and confusing expected shortfall with V a R.

19 Every risk measure comes from quantiles

- Quantile: the value with a given probability below it
- VaR at c = −[the (1 − c) quantile of P&L]
- Same thing: the c quantile of the loss distribution
- Long position: the lower tail of returns
- Short position: the upper tail of returns

IN THE LESSON

- A quantile is the value with a given probability below it, and every standard risk measure is a function of the quantiles of the P and L or loss distribution.
- V a R at confidence c is the one minus c quantile of P and L with the sign reversed: at ninety five per cent, take the five per cent quantile of P and L, a large negative number, and flip its sign.
- Equivalently, it is the c quantile of the loss distribution.
- So first decide which tail matters.
- A long position loses when returns are low, so it uses the lower tail of returns.
- A short position loses when returns are high, so it uses the upper tail.

20 Quantiles, expected shortfall and spectral measures

Quantile = μ + z × σ (normal)
ES = average of VaR over levels from c to 100%
Spectral = Σ w_i × q_i, w_i ≥ 0, Σ w_i = 1

Quantiles, q_i	Ordered better to worse
Weights, w_i	Never fall towards the worse tail
Expected shortfall	Equal weights on the tail

IN THE LESSON

- For a normal, a quantile is the mean plus z standard deviations, with z negative in the lower tail.
- Expected shortfall is the average of the V a Rs at every confidence level from c to one hundred per cent, so we can approximate it by slicing the tail into equal probability slices and averaging the V a R of each.
- More generally, a spectral risk measure is a weighted average of quantiles: the sum of w i times q i.
- Order the quantiles from the better outcome to the worse.
- The measure is coherent if the weights are non negative, sum to one and never fall towards the worse tail.
- Expected shortfall puts equal weights on the tail.

21 Example 11: ES by slicing the tail

Normal: μ = $8k, σ = $250k. Cut the 95% tail into 5 slices; midpoints, one convention

Levels 95.5% and 96.5%	z = 1.695 and 1.812	**$416k, $445k**
Levels 97.5% and 98.5%	z = 1.960 and 2.170	**$482k, $535k**
Level 99.5%	z = 2.576	**$636k**
ES: average of the five	(416+445+482+535+636) ÷ 5	**$503k**
Exact normal ES	250k × φ(1.645) ÷ 0.05 − 8k	**$508k**

IN THE LESSON

- Approximate Ashby's ninety five per cent expected shortfall under the normal model: mean eight thousand dollars, standard deviation two hundred and fifty thousand.
- Cut the five per cent tail into five one per cent slices, and use the middle of each, which is one of several conventions.
- At ninety five and a half and ninety six and a half per cent, z is one point six nine five and one point eight one two, giving V a Rs of about four hundred and sixteen and four hundred and forty five thousand dollars.
- At ninety seven and a half and ninety eight and a half, z is one point nine six and two point one seven: four hundred and eighty two and five hundred and thirty five thousand.
- At ninety nine and a half, z is two point five seven six: six hundred and thirty six thousand.
- Their average is about five hundred and three thousand.
- The exact normal answer, using the density phi, is about five hundred and eight thousand.

22 Example 12: a spectral risk measure

Weights on the five slice VaRs, better to worse: 0.1, 0.1, 0.2, 0.3, 0.3

Weights: ≥ 0, sum to 1, rising	0.1 + 0.1 + 0.2 + 0.3 + 0.3	**1.0**
First three slices	0.1×416 + 0.1×445 + 0.2×482	**$182.5k**
Last two slices	0.3×535 + 0.3×636	**$351.3k**
Spectral measure	182.5 + 351.3	**$533.8k**
Weights falling to the worse tail	e.g. 0.3, 0.3, 0.2, 0.1, 0.1	**not coherent**

IN THE LESSON

- Take the five slice V a Rs from example eleven, ordered better to worse, with weights point one, point one, point two, point three and point three.
- They are non negative, sum to one and rise towards the worse tail: coherent.
- The first three slices give point one times four hundred and sixteen, plus point one times four hundred and forty five, plus point two times four hundred and eighty two: one hundred and eighty two point five thousand dollars.
- The last two give point three times five hundred and thirty five plus point three times six hundred and thirty six: three hundred and fifty one point three thousand.
- The total is five hundred and thirty three point eight thousand.
- Falling weights, such as point three, point three, point two, point one, point one, would not be coherent.

23 Example 13: which tail? A short position

A second desk: short $5m of an index; daily return mean 0.04%, SD 1.25%; 95%

Short loses when returns rise	use the upper tail	**z = +1.645**
Upper quantile of returns	0.04% + 1.645 × 1.25%	**2.09625%**
Short position VaR	2.09625% × $5m	**$104,812.50**
Same-size long: lower tail	0.04% − 1.645 × 1.25%	**−2.01625%**
Long position VaR	2.01625% × $5m	**$100,812.50**

IN THE LESSON

- Suppose a second desk is short five million dollars of an index, with daily mean return point zero four per cent and standard deviation one point two five per cent.
- A short loses when returns rise, so its ninety five per cent V a R uses the upper tail, with z of plus one point six four five.
- The upper quantile is point zero four per cent plus one point six four five times one point two five per cent: two point zero nine six two five per cent.
- On five million the V a R is one hundred and four thousand eight hundred and twelve dollars fifty.
- A long of the same size uses the lower quantile, minus two point zero one six two five per cent, a V a R of one hundred thousand eight hundred and twelve dollars fifty.

24 The standard error of a quantile

V (variance term) = p × (1 − p) ÷ n
SE = √V ÷ f(q)
95% interval: estimate ± 1.96 × SE

Tail probability, p	e.g. 0.05 for 95% VaR
Observations, n	Independent observations
Density, f(q)	How thick outcomes are there
Standard error, SE	Shrinks with √n

IN THE LESSON

- An estimate is itself a random variable, so we judge its reliability with a standard error.
- For a quantile at tail probability p, from n independent observations, first work out a new V, the variance term, equal to p times one minus p, over n.
- The standard error, S E, is the square root of V divided by f of q, the probability density at the quantile.
- A rough ninety five per cent interval is the estimate plus or minus one point nine six standard errors.
- Here p is a tail probability such as point zero five, n counts independent observations, f says how thickly outcomes crowd around the quantile, and the error shrinks with the square root of n.
- For a normal, the density in dollars is phi at z divided by sigma, where phi is the standard normal density, e to the minus z squared over two, over root two pi.

25 Example 14: how precise is Ashby's 95% VaR?

n = 500, p = 0.05; P&L roughly normal, σ = $250,000; HS VaR $398,000

V, the variance term	0.05 × 0.95 ÷ 500	**0.000095**
√V	√0.000095	**0.009747**
SE = √V × σ ÷ φ(1.645)	0.009747 × 250,000 ÷ 0.10311	**$23,632**
1.96 standard errors	1.96 × 23,632	**±$46,319**
95% interval for VaR	398,000 ± 46,319	**$352k to $444k**

IN THE LESSON

- How precise is Ashby's ninety five per cent historical V a R of three hundred and ninety eight thousand dollars?
- With five hundred observations, p of point zero five, and roughly normal P and L with a standard deviation of two hundred and fifty thousand, the variance term V is point zero five times point nine five over five hundred: point zero zero zero zero nine five.
- Its square root is about point zero zero nine seven four seven.
- The standard error is that times two hundred and fifty thousand, divided by phi at one point six four five, point one zero three one one: about twenty three thousand six hundred and thirty two dollars.
- One point nine six standard errors is forty six thousand three hundred and nineteen, so the interval is about three hundred and fifty two thousand to four hundred and forty four thousand dollars.

26 Example 15: deeper in the tail, and more data

99% VaR: p = 0.01, z = 2.326, φ(2.326) = 0.026674; σ = $250,000; n = 500

V at 99%	0.01 × 0.99 ÷ 500	**0.0000198**
SE at 99%	√V × 250,000 ÷ 0.026674	**$41,705**
Against the 95% SE	41,705 ÷ 23,632	**1.76 times**
n = 2,000: 4 times the data	SE ÷ √4	**$20,852**
Cut the SE to a third	9 × 500 days	**n = 4,500**

Tail estimates are far less precise: few observations, low density

IN THE LESSON

- For the ninety nine per cent V a R, p is point zero one, z is two point three two six, and the density there is only point zero two six six seven four.
- V is point zero one times point nine nine over five hundred: point zero zero zero zero one nine eight.
- The standard error is root V times two hundred and fifty thousand, divided by point zero two six six seven four: about forty one thousand seven hundred and five dollars, one point seven six times the ninety five per cent error.
- Four times the data, two thousand days, halves the error to about twenty thousand eight hundred and fifty two.
- A third needs nine times the data, four thousand five hundred days.
- Tail estimates are far less precise, because few observations live out there and the density is low.

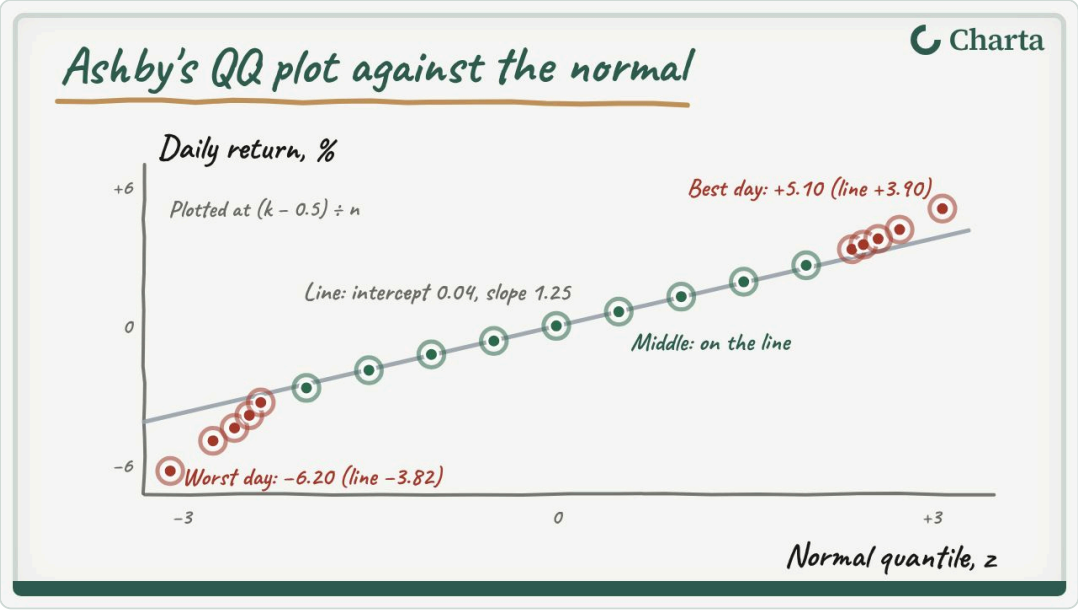
27 Quantile-quantile plots

- Sample quantiles (vertical) vs normal (horizontal)
- Same shape as the reference: a straight line
- Intercept = location; slope = scale
- Intercept 0.2, slope 1.8: normal, mean 0.2, SD 1.8
- Departures from the line reveal the features

IN THE LESSON

- Our last tool is a picture.
- A quantile quantile plot, or Q Q plot, graphs the quantiles of a sample against the quantiles of a reference distribution, usually the standard normal.
- Here the sample goes on the vertical axis and the normal on the horizontal axis.
- If the sample has the same shape as the reference, the points fall on a straight line.
- The line's intercept shows the difference in location, and its slope the difference in scale.
- So a line with intercept point two and slope one point eight, against the standard normal, means a normal with mean point two and standard deviation one point eight.
- When the points leave the line, the departures tell us how the sample differs.

28 Ashby's QQ plot against the normal



IN THE LESSON

- Here is Ashby's Q Q plot, with daily returns in per cent up the side and standard normal quantiles along the bottom.
- Each sorted return is plotted at the normal quantile for k minus a half, over n.
- The grey line has intercept point zero four and slope one point two five, the mean and standard deviation from example four.
- The middle dots hug it.
- At the left end they fall below it: the worst day, minus six point two per cent, sits where the line expects about minus three point eight two.
- At the right end they rise above it: the best day, plus five point one, against about three point nine on the line.
- Below at the left and above at the right means fatter tails than the normal.

29 Reading the shape of a QQ plot

Sample on the vertical axis, normal on the horizontal

What you see	What it means
A straight line	Same shape
Below at left, above at right	Fatter tails
Concave: bends downward	Negative skew
A few stray points	Outliers
Reading the bend backwards	Typical error

IN THE LESSON

- With the sample on the vertical axis: A straight line means the sample has the same shape as the normal, differing only in location and scale.
- Points below the line at the left end and above it at the right end mean fatter tails than the normal: extreme days, bad and good, are more extreme than a normal allows.
- A concave curve, one that bends downward, signals negative skew, a long left tail of losses.
- A few stray points far from the rest signal outliers.
- The typical error is reading the direction of the bend backwards, so always check which axis holds the sample.

30 Example 16: what Ashby's QQ plot means for VaR

Ashby: fitted normal (0.04%, 1.25%) against the sample's own 1% quantile

Normal 1% quantile of return	0.04% − 2.326 × 1.25%	**−2.8675%**
Sample 1% quantile	the 5th-worst return	**−3.25%**
Normal 99% VaR	2.8675% × $20m	**$573,500**
Historical 99% VaR	3.25% × $20m	**$650,000**
Normal understates by	650,000 − 573,500	**$76,500**

At 95% they were close: $403,250 normal against $398,000 historical

IN THE LESSON

- Compare the one per cent quantile of returns under Ashby's fitted normal with the sample's own.
- The normal gives point zero four per cent minus two point three two six times one point two five per cent: minus two point eight six seven five per cent.
- The sample's one per cent quantile, the fifth worst return, is minus three point two five per cent.
- On twenty million dollars, normal V a R is five hundred and seventy three thousand five hundred, and historical V a R six hundred and fifty thousand.
- The normal understates by seventy six thousand five hundred dollars.
- Fat tails mean normal V a R understates high confidence risk.
- At ninety five per cent they were close: four hundred and three thousand two hundred and fifty against three hundred and ninety eight thousand.

31 Mistakes to avoid

- HS: the Tail-th worst loss, not the 95th; state it
- Parametric: −μ, not +μ; scale σ by √T, not by T
- Lognormal for long horizons: VaR never exceeds value
- ES: count the tail correctly, fractions included
- SE: square root, density, σ scale; tails are imprecise

IN THE LESSON

- In historical simulation, V a R is the tail count's worst loss, here the twenty fifth, not the ninety fifth, and you must state the convention.
- In the parametric method, subtract the mean rather than adding it, and scale volatility by the square root of time, not by time.
- Use the lognormal formula when prices are lognormal, especially over long horizons, and remember that its V a R never exceeds the position value.
- For expected shortfall, count the tail correctly, including a fractional observation.
- For standard errors, remember the square root, the density and the volatility scale, and expect tail estimates to be imprecise.

32 Try three quick questions

1. 250 days, 98%. Five worst losses $90k, $80k, $70k, $65k, $60k. HS VaR and ES?
   Answer: Tail 5. VaR $60k; ES = 365 ÷ 5 = $73k.
2. Daily P&L normal: mean $1,000, SD $20,000. 4-day 99% VaR?
   Answer: 2.326 × 40,000 − 4,000 = $89,040.
3. QQ: middle on a line, intercept −0.1, slope 2.5; below at left, above at right?
   Answer: Mean −0.1, SD 2.5, but fatter tails.

IN THE LESSON

- Question one.
- Over two hundred and fifty days at ninety eight per cent, the five worst losses are ninety, eighty, seventy, sixty five and sixty thousand dollars.
- Find V a R and E S.
- The tail count is five, so V a R is sixty thousand, and expected shortfall is three hundred and sixty five over five, seventy three thousand.
- Question two.
- Daily P and L is normal, mean one thousand dollars, standard deviation twenty thousand.
- Find the four day, ninety nine per cent V a R.
- Two point three two six times forty thousand, minus four thousand: eighty nine thousand and forty dollars.
- Question three.
- A Q Q plot's middle follows a line with intercept minus point one and slope two point five, with points below it at the left and above it at the right.
- Mean minus point one, standard deviation two point five, but fatter tails.