CHARTA LEARNING · LESSON NOTES

Measures of Financial Risk

FRM Part I

1 Two numbers describe a portfolio

- Each portfolio: expected return and volatility
- Investors want more return and less risk
- Volatility measures the wobble around the average
- Mixing assets does not just average the wobble

IN THE LESSON

- Imagine you must describe an investment to a friend using only two numbers.
- The first is what you expect to earn.
- The second is how much that result might wobble.
- That is the whole idea behind the mean-variance framework.
- Expected return is the average outcome.
- Standard deviation, which we call volatility, measures the wobble.
- The framework assumes people like more return and dislike more risk.
- So, given two portfolios with the same return, you pick the calmer one, and given two with the same volatility, you pick the one that pays more.
- What makes it powerful is that mixing investments does not simply average their wobbles.
- The way they move together matters, and that is where we go next.

2 Risk of a two-asset portfolio

$$\sigma_p^2 = w_1^2 \sigma_1^2 + w_2^2 \sigma_2^2 + 2 \cdot w_1 \cdot w_2 \cdot \rho \cdot \sigma_1 \cdot \sigma_2$$

Weights, w	Share of money in each asset
Volatility, sigma	Standalone risk of each asset
Correlation, rho	How closely the two move together
Portfolio variance	Portfolio risk, squared

IN THE LESSON

- Here is the rule for mixing two assets.
- Take the share of money in each, called the weights, and the volatility of each.
- The portfolio variance, which is volatility squared, has three parts.
- The first is the weight squared times the variance of the first asset.
- The second is the same for the second asset.
- The third is the interesting one: twice the two weights, times the correlation, times both volatilities.
- Correlation runs from minus one to plus one and tells us whether the assets move together.
- If the correlation is one, the third part is as large as it can be.
- If it is below one, the third part shrinks, and the portfolio is calmer than the plain average of the two volatilities.
- This assumes the weights are positive and add to one.

3 Example 1: a half and half mix

A: 11% return, 24% risk. B: 5%, 12%. Correlation 0, half each

Portfolio return	0.5×11% + 0.5×5%	**8%**
Variance	0.0144 + 0.0036 + 0	**0.0180**
Portfolio volatility	√0.018	**13.42%**
Average of the two risks	0.5×24% + 0.5×12%	**18%**

IN THE LESSON

- Let us try it.
- Asset A is a share fund that expects eleven per cent with volatility twenty four per cent.
- Asset B is a bond fund expecting five per cent with volatility twelve per cent.
- The two are unrelated, so the correlation is zero, and we put half our money in each.
- The return is simply the weighted average: half of eleven plus half of five gives eight per cent.
- For the variance, A contributes point zero one four four, B contributes point zero zero three six, and the third part is zero.
- The total is point zero one eight.
- The square root is thirteen point four per cent.
- The plain average of the two volatilities is eighteen per cent, so mixing has cut risk by over four points without touching the return.

4 Correlation drives the benefit

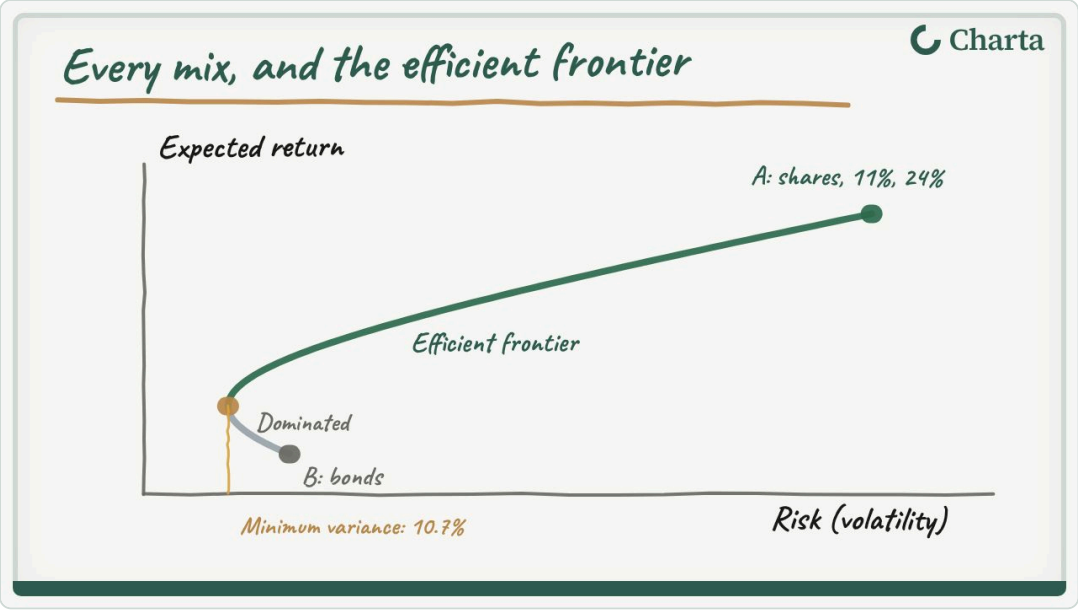
Same half and half mix, same 8% return, different correlation

Correlation	Portfolio volatility
+1.0	18.00%
+0.5	15.87%
0.0	13.42%
−1.0	6.00%

IN THE LESSON

- Now keep the same half and half mix and the same eight per cent return, and change only the correlation.
- When the correlation is plus one, the two assets move in lockstep and there is no help at all: the portfolio volatility is exactly the average, eighteen per cent.
- At plus a half it falls to about fifteen point nine.
- At zero, with no link, it is thirteen point four, our example.
- And at minus one, where one asset always zig zags against the other, the wobble cancels almost entirely and the volatility is only six per cent.
- The lesson is that the gap between the average volatility and the true portfolio volatility is the diversification benefit.
- It exists whenever correlation is below one.
- The most common mistake is to average volatilities and stop there.

5 Every mix, and the efficient frontier



IN THE LESSON

- Now draw every possible mix of our two funds.
- Risk runs along the bottom and expected return up the side.
- Start at the bond fund, B, and slowly move money into the share fund, A.
- At first the curve actually moves left, because a little A reduces risk through diversification.
- It reaches a leftmost point, the global minimum variance portfolio, here about ten point seven per cent risk.
- Beyond that point, more A means more risk and more return.
- The whole left edge is called the minimum variance frontier, but only the upper part, from that leftmost point upwards, is the efficient frontier: the best return available for each level of risk.
- The lower part is a trap, because a portfolio with the same risk and a higher return sits right above it.

6 Reading the frontier

- Dominated: better is available elsewhere
- Efficient frontier: the upper edge only
- It starts at the global minimum variance portfolio
- Your risk appetite picks the point on it

IN THE LESSON

- A few ways to read the picture.
- A portfolio is dominated if another one offers more return for the same risk, or less risk for the same return.
- Everything inside the curved region, or on its lower edge, is dominated.
- Only the upper edge from the minimum variance point is efficient, so be careful: the whole minimum variance curve is not the efficient frontier, only its top half.
- Diversification does not always remove all risk: with zero correlation, even the safest mix here still has volatility of about ten point seven per cent.

7 Add a safe asset: the Sharpe ratio

Sharpe = (E(R) − Rf) ÷ σ

Expected return	Average return expected
Risk-free rate	Return on the safe asset
Volatility, sigma	Total risk of the portfolio

IN THE LESSON

- Suppose you can also hold a risk-free asset, such as a short government bill, paying a known rate.
- Now the best risky mix is no longer a matter of taste.
- Draw a straight line from the risk-free rate on the vertical axis, and swing it up until it just touches the efficient frontier.
- The touching point is the tangency portfolio, and it has the steepest line, which means the highest return per unit of risk.
- That measure is the Sharpe ratio: the expected return above the risk-free rate, divided by volatility.
- Every investor holds some mix of the tangency portfolio and the safe asset, the cautious ones holding more safe asset.
- Beware of comparing returns alone, since a higher return may simply carry much higher risk.

8 Example 2: which portfolio is better?

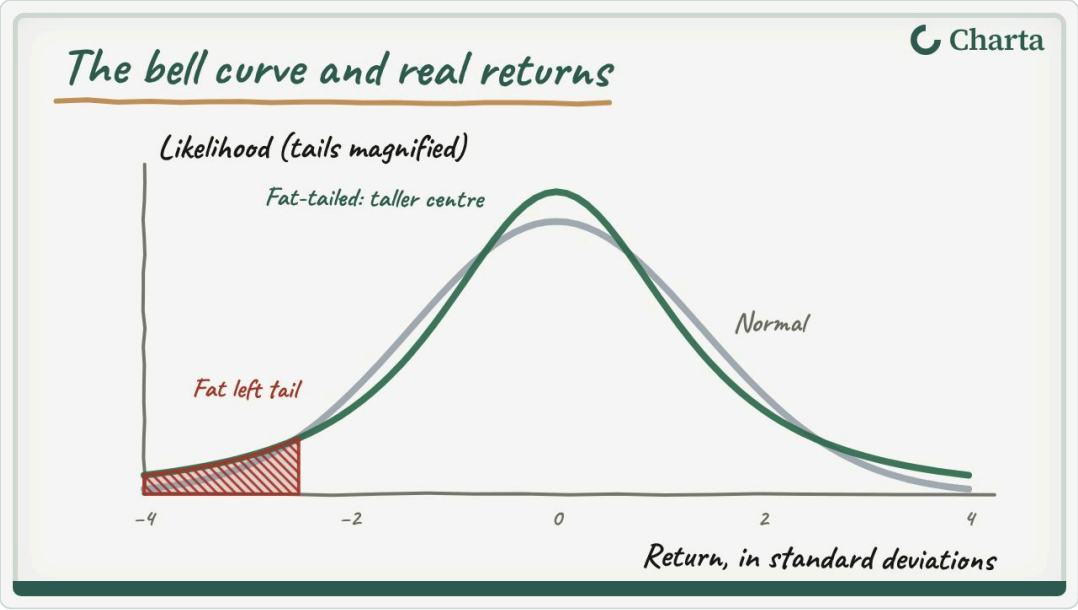
Risk-free rate 2%. P1: 8% return, 12% risk. P2: 11%, 20%. P3: 10%, 16%

Portfolio 1	(8 − 2) ÷ 12	**0.50**
Portfolio 2	(11 − 2) ÷ 20	**0.45**
Portfolio 3	(10 − 2) ÷ 16	**0.50**
Highest return, best Sharpe?	no: Portfolio 2 is lowest	**No**

IN THE LESSON

- Let us compare three portfolios with a risk-free rate of two per cent.
- Portfolio one expects eight per cent with volatility twelve.
- Its excess return is six points, and six divided by twelve gives a Sharpe ratio of point five.
- Portfolio two expects eleven per cent with volatility twenty.
- The excess is nine, and nine over twenty is point four five.
- Portfolio three expects ten per cent with volatility sixteen.
- Its excess is eight, and eight over sixteen is also point five.
- So portfolios one and three tie, and portfolio two, despite the highest return, gives the least reward for each unit of risk.
- The tangency portfolio is the one with the highest Sharpe ratio.

9 The bell curve and real returns



IN THE LESSON

- Many textbook models use the normal distribution, the bell curve.
- It is symmetric, fully described by its average and its variance, with no other freedom.
- Real returns on shares look different.
- The grey curve here is a normal.
- The green one is a fat-tailed distribution with the same average and the same variance.
- Notice two things.
- The green curve is taller and narrower in the middle, so very quiet days are more common.
- And far out on either side it lies above the grey one, so extreme days are more common too.
- A model that assumes the grey curve treats a huge loss as almost impossible.

10 Three ways real returns differ

- Excess kurtosis: fat tails, taller narrow centre
- Negative skewness: heavier left tail of losses
- Volatility clustering: turbulence follows turbulence
- Normal: skewness 0, kurtosis 3, symmetric

IN THE LESSON

- Let us name the differences.
- The first is kurtosis.
- The normal has kurtosis exactly three.
- Real returns often have much more, which we call excess kurtosis, meaning fat tails and a taller, narrower centre.
- The second is skewness.
- The normal is perfectly symmetric, so its skewness is zero.
- Equity returns often show negative skewness: a longer, heavier left tail, where the large moves are losses rather than gains.
- The third is volatility clustering.
- Turbulent days tend to follow turbulent days, and calm days follow calm days, so volatility is not constant through time.
- A normal model misses all three, so it can understate losses far out in the tail.

11 Skewness and kurtosis

Skew = E[(R − μ)³] ÷ σ³
Kurt = E[(R − μ)⁴] ÷ σ⁴

Mean, mu	Average return
Volatility, sigma	Volatility, to the same power
Skewness	Zero for a normal
Kurtosis	Three for a normal

IN THE LESSON

- Here is how the two shape numbers are measured.
- Take each return, subtract the average, and raise the gap to a power.
- For skewness the power is three, so the sign is kept: a symmetric spread cancels to zero, and a longer tail on one side does not.
- We then average and divide by volatility cubed, which makes the answer a pure number without units.
- A negative result means a heavier left tail.
- For kurtosis the power is four, which punishes large gaps heavily, and we divide by volatility to the fourth power.
- A normal distribution scores zero on skewness and three on kurtosis.
- Subtract three and you have excess kurtosis, which is the number people usually quote.
- Positive excess kurtosis means fatter tails than the normal.

12 Example 3: how rare is a four sigma day?

Normal model, one day, one-sided: a fall of 4 standard deviations

Probability of 4 sigma fall	about 0.0000317	**0.00317%**
Typical wait in days	1 ÷ 0.0000316712	**31,574 days**
Trading days in a year	assume	**252**
Wait in years	31,574 ÷ 252	**125 years**

IN THE LESSON

- A move of four standard deviations downwards sounds almost impossible, so let us test it.
- For a normal distribution, the chance of falling four standard deviations or more below the average is about point zero zero zero zero three two, which is roughly one in thirty one thousand five hundred seventy four.
- If there are about two hundred fifty two trading days a year, that is one such day every one hundred twenty five years or so.
- Yet in real data, moves of four or five standard deviations show up far more often than that, within a working lifetime and sometimes several times in one.
- This is evidence that returns are not normal.

13 Example 4: the Jarque-Bera test

JB = n÷6 × (S² + (K − 3)²÷4). n = 1000, S = −0.5, K = 5

Skewness part	(−0.5)²	**0.25**
Excess kurtosis part	(5 − 3)² ÷ 4	**1.00**
Statistic	1000÷6 × 1.25	**208.3**
Cut-off, 2 degrees, 5%	chi-square table	**5.99**

IN THE LESSON

- How do we test whether a sample of returns is normal?
- The Jarque-Bera statistic combines skewness and kurtosis into one number.
- It is n over six, times the sum of skewness squared and a quarter of excess kurtosis squared, where n is the number of observations.
- Take a thousand daily returns with sample skewness of minus a half and sample kurtosis of five.
- Skewness squared is a quarter.
- Excess kurtosis is five minus three, which is two, and two squared over four is one.
- The sum is one and a quarter.
- A thousand over six is about one hundred sixty six point seven, so the statistic is about two hundred eight point three.
- In large samples we compare it with a chi-square distribution with two degrees of freedom, whose five per cent cut-off is about five point nine nine.
- Two hundred eight is far above that, so we reject normality.

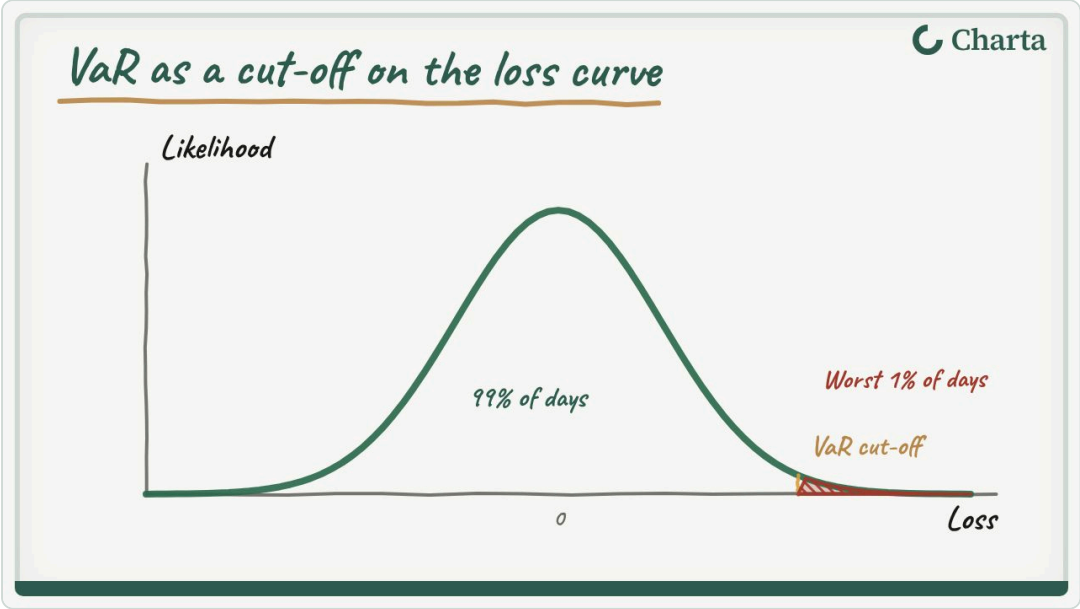
14 What Value at Risk means

- A quantile of the loss distribution
- Needs a holding period and a confidence level
- 99% one-day VaR $5m: 1% chance of a larger loss
- A threshold, not the worst case

IN THE LESSON

- Now to the number risk managers quote daily.
- Value at Risk, or V a R, is a quantile of the loss distribution.
- It is the loss that you do not expect to exceed over a stated holding period, at a stated confidence level.
- Read this sentence carefully: a one-day ninety nine per cent V a R of five million dollars means that, under the model's assumptions, there is a one per cent chance of losing more than five million dollars in a day.
- Notice what it does not say.
- It does not say that the loss will be five million, or that five million is the worst case.
- And the days that break the limit are not scheduled politely, one per hundred days.
- They can cluster, even arrive in a row.
- V a R is a threshold, not a ceiling.

15 VaR as a cut-off on the loss curve



IN THE LESSON

- Picture the likelihood of each possible outcome for tomorrow, with big losses on the right.
- Most of the weight sits in the middle, where small gains and losses live.
- Now choose a confidence level, say ninety nine per cent.
- We cut the curve at the point where ninety nine per cent of the area lies to the left.
- That cut-off is the V a R.
- The small hatched area to the right is the one per cent of days that breach it.
- Here is the important fact about this picture: V a R tells us where the hatched area begins.
- It tells us nothing about how far the curve stretches into it, whether the losses past the line are a little bigger or ten times bigger.

16 Parametric VaR for a normal

VaR = (z × σ − μ) × V

Z score	1.645 at 95%, 2.326 at 99%
Volatility, sigma	One-period volatility, as decimal
Mean, mu	One-period average return
Portfolio value	Value of the portfolio

IN THE LESSON

- When we assume returns are normal, V a R has a simple formula.
- Take the z score for the confidence level: one point six four five for ninety five per cent and two point three two six for ninety nine per cent.
- Multiply it by the volatility of one period's returns, subtract the average return, since a positive average cushions the loss, and multiply by the portfolio value.
- Use decimals throughout, so one and a half per cent is point zero one five.
- The period must match the volatility: use daily volatility for a daily V a R.
- The z score is the number of standard deviations below the average that marks off the tail.
- If the data show fatter tails than the normal, this formula can understate the far tail.

17 Example 5: one-day parametric VaR

Value $10m, daily σ = 1.2%, daily mean = 0.05%

95%: z × σ − μ	1.645×0.012 − 0.0005	**0.01924**
95% VaR	0.01924 × 10m	**$192,400**
99%: z × σ − μ	2.326×0.012 − 0.0005	**0.027412**
99% VaR	0.027412 × 10m	**$274,120**

IN THE LESSON

- Let us work one through.
- A portfolio is worth ten million dollars.
- Daily volatility is one point two per cent, written point zero one two, and the average daily return is point zero five per cent, written point zero zero zero five.
- At ninety five per cent, z is one point six four five.
- Multiply by the volatility to get point zero one nine seven four, subtract point zero zero zero five to get point zero one nine two four, and multiply by ten million: a V a R of one hundred ninety two thousand four hundred dollars.
- At ninety nine per cent z is two point three two six.
- That gives point zero two seven nine one two, minus point zero zero zero five, which is point zero two seven four one two, times ten million: two hundred seventy four thousand one hundred twenty dollars.

18 Example 6: the square root of time

99% one-day VaR $274,120. Horizon: 10 days

One-day VaR	from Example 5	**$274,120**
Square root of 10	√10	**3.1623**
Ten-day VaR	274,120 × 3.1623	**$866,844**
Compared with 10 × one-day	less than $2,741,200	**Smaller**

IN THE LESSON

- What about a ten-day V a R, which is the horizon regulators have often used?
- We do not recompute everything.
- The square-root-of-time rule says V a R over T days equals the one-day V a R times the square root of T.
- Take our ninety nine per cent one-day figure of two hundred seventy four thousand one hundred twenty dollars.
- The square root of ten is about three point one six two.
- Multiplying gives about eight hundred sixty six thousand eight hundred forty four dollars.
- Independent daily moves partly cancel, so risk grows more slowly than time.
- But the rule assumes returns are independent from day to day, volatility stays constant, and the portfolio does not change over the period.
- Clustering and trading both break it.

19 Where VaR falls short

- Blind to the size of losses beyond the cut-off
- Unreliable with fat tails, illiquidity, options
- Not guaranteed subadditive: can punish merging
- A threshold, not a worst case

IN THE LESSON

- V a R is useful, but we must know its limits.
- First, it says nothing about how bad the losses beyond the quantile are.
- Second, it is unreliable when returns have fat tails, when positions are illiquid and cannot be sold at the model price, or when payoffs are non-linear, as with options.
- Third, and most surprising, it is not guaranteed to be subadditive.
- That is a technical word we meet shortly, but in plain terms it means that the V a R of a combined portfolio can exceed the sum of the V a Rs of its parts, so V a R can seem to punish diversification.
- Remember: it is a threshold, not a worst case.

20 Expected shortfall for a normal

ES = σ × φ(z) ÷ (1 − c) × V

Volatility, sigma	One-period volatility
Density at the cut-off	Height of the bell curve at z
Confidence level, c	For example 0.99
Portfolio value	Value of the portfolio

IN THE LESSON

- Expected shortfall, or E S, answers the question V a R ignores: when things go beyond the cut-off, how bad are they on average?
- It is the average loss given that the loss reaches or exceeds V a R.
- Another name is conditional V a R.
- For a normal distribution with a mean of zero there is a neat formula.
- Take the volatility, multiply by the height of the bell curve at the V a R cut-off, divide by one minus the confidence level, and multiply by portfolio value.
- At ninety nine per cent the factor works out to about two point six six five volatilities, against two point three two six for V a R.
- Because it averages losses at or beyond V a R, expected shortfall is always at least as big as V a R at the same confidence.

21 Example 7: VaR and ES compared

Value $10m, daily σ = 1.2%, mean 0, confidence 99%

VaR factor	z	**2.326**
VaR	2.326 × 0.012 × 10m	**$279,120**
ES factor	φ(z) ÷ (1 − c)	**2.665**
ES	2.665 × 0.012 × 10m	**$319,800**

IN THE LESSON

- Let us compute both for the same book.
- The portfolio is worth ten million dollars, the daily volatility is one point two per cent, and we set the mean to zero for simplicity.
- At ninety nine per cent, V a R is two point three two six times point zero one two times ten million, which is two hundred seventy nine thousand one hundred twenty dollars.
- Expected shortfall uses the factor of about two point six six five, so it is two point six six five times point zero one two times ten million, or three hundred nineteen thousand eight hundred dollars.
- The tail average is about fifteen per cent above the cut-off.
- V a R says where the bad one per cent starts, and expected shortfall says what that bad one per cent costs on average.

22 Example 8: expected shortfall from history

100 days, 95%. Worst five losses, $m: 9, 7, 6, 5, 4

Tail size	5% of 100 days	**5 days**
VaR (fifth worst)	convention used here	**$4m**
Sum of worst five	9 + 7 + 6 + 5 + 4	**31**
Expected shortfall	31 ÷ 5	**$6.2m**

IN THE LESSON

- You can also find expected shortfall from data, with no formula.
- Suppose we hold one hundred daily profit and loss outcomes and want ninety five per cent expected shortfall.
- The worst five per cent of one hundred is five outcomes.
- Say the five worst losses, in millions of dollars, are nine, seven, six, five and four.
- Using the convention that the V a R is the fifth worst loss, that is four million.
- The expected shortfall is the average of the worst five: nine plus seven plus six plus five plus four is thirty one, and thirty one over five is six point two million dollars.
- A common error is to average only the outcomes strictly beyond the V a R point, which would give six point seven five, the average of nine, seven, six and five, and understate the answer.
- Include the V a R point itself.

23 VaR and ES side by side

Same loss distribution and confidence. Middle: VaR. Right: ES

What it gives	Cut-off	**Tail average**
Size of tail losses	Ignored	**Counted**
Size at same confidence	Smaller	**At least VaR**
Subadditive	Not guaranteed	**Yes**
Coherent	No	**Yes**

IN THE LESSON

- Let us line the two measures up.
- They are alike in more ways than people expect.
- Both need a holding period and a confidence level.
- Both are read off the same loss distribution, and both depend on how good that distribution is.
- Both are stated in money or as a share of value.
- Now the differences.
- V a R is a single quantile and ignores how large the losses beyond it are.
- Expected shortfall averages those losses, so it responds when the tail gets worse.
- For the same confidence, expected shortfall is never smaller than V a R.
- And expected shortfall is subadditive, while V a R is not guaranteed to be.
- That last difference is the reason regulators and risk managers have taken a growing interest in expected shortfall.

24 The four coherence properties

- Monotonicity: worse in every scenario, riskier
- Positive homogeneity: $\rho(kX) = k \times \rho(X)$
- Translation invariance: $\rho(X + k) = \rho(X) - k$
- Subadditivity: $\rho(X + Y) \leq \rho(X) + \rho(Y)$

IN THE LESSON

- Is there a list of sensible demands we can make of any risk measure?
- Yes, a coherent risk measure, which we write rho, satisfies four properties.
- Monotonicity: if portfolio X loses at least as much as portfolio Y in every scenario, its risk must be at least as high.
- Positive homogeneity: scale a position by a factor k and its risk scales by k, which quietly ignores any liquidity effect.
- Translation invariance: add k of riskless cash and the risk falls by exactly k, when risk is expressed as the capital needed.
- And subadditivity: the risk of two books together is no more than the sum of their separate risks.
- In plain words, diversification cannot increase risk, and merging two books cannot create risk.

25 Example 9: the properties in numbers

A coherent measure gives a position capital of $500,000

Double the position	2 × 500,000	**$1,000,000**
Add $100,000 of cash	500,000 − 100,000	**$400,000**
Merge two $300,000 books	at most 300,000 + 300,000	**≤ $600,000**
Worse in every scenario	risk must not be lower	**Monotone**

IN THE LESSON

- Let us see each property with simple numbers.
- Suppose some risk measure says a position needs five hundred thousand dollars of capital.
- Positive homogeneity: if we double the position, a coherent measure needs twice as much, one million.
- Translation invariance: if we add one hundred thousand dollars of riskless cash, the capital needed falls by exactly that, to four hundred thousand.
- Subadditivity: if book one needs three hundred thousand and book two needs three hundred thousand, then the merged book should need no more than six hundred thousand.
- Monotonicity needs no sum: a book that loses at least as much in every scenario must never be shown as safer.
- A measure that fails any of these can give misleading signals, and that is the practical case for coherence.

26 Which measures are coherent?

The four tests: monotone, homogeneous, translation, subadditive

Measure	Coherent?
Expected shortfall	Yes, passes all four
Value at Risk	No: not always subadditive
Standard deviation	No: fails monotonicity

IN THE LESSON

- Now apply the test to the measures we know.
- Expected shortfall is coherent: it passes all four properties.
- Value at Risk passes monotonicity, positive homogeneity and translation invariance, but it is not guaranteed to be subadditive, so it is not coherent.
- Standard deviation fails monotonicity, because it measures spread, not loss: a portfolio that always loses more than another can have a smaller spread, and so look safer.
- The key fact is that V a R fails on subadditivity, which we demonstrate next.

27 Can merging two books create risk?

- Two independent bonds, each could lose $10m
- Each has only a 4% chance of that loss
- Alone, each looks safe at 95%
- Together, the risk appears out of nowhere

IN THE LESSON

- Here is the puzzle.
- If risk measures are sensible, joining two portfolios should never produce more risk than the two had apart.
- Yet V a R can break that rule.
- The standard demonstration uses two bonds.
- They are independent, and each has a small chance of a large loss.
- At ninety five per cent confidence, a small chance can hide inside the five per cent we are told to ignore, so each bond looks risk free, with a V a R of zero.
- But hold the two together and the small chances add up until they no longer fit in the ignored five per cent, and V a R suddenly jumps.
- The failure arises with skewed, discrete or fat-tailed losses.
- Let us do the numbers.

28 Example 10: two bonds break VaR

Each bond: 4% chance of losing $10m, else 0. Independent. 95%

VaR of each bond	4% sits inside the 5% tail	**$0**
Chance neither defaults	0.96 × 0.96	**0.9216**
Chance of some loss	1 − 0.9216	**7.84%**
Combined VaR	7.84% exceeds 5%: $10m	**$10m > $0**

IN THE LESSON

- Each bond loses ten million dollars with probability four per cent, and nothing otherwise.
- At ninety five per cent we ignore the worst five per cent.
- A four per cent chance fits inside it, so each bond's V a R is zero.
- Zero plus zero is zero.
- Now combine them.
- The chance neither bond defaults is point nine six times point nine six, which is point nine two one six.
- So the chance of at least one loss is one minus that, seven point eight four per cent, which is more than five per cent and no longer fits inside the ignored tail.
- The ninety five per cent V a R of the combined portfolio is therefore ten million dollars.
- Ten million is larger than the zero from adding the stand-alone figures, so subadditivity fails.

29 When VaR does behave: normal returns

VaR(A+B) = √(V a R_A² + V a R_B² + 2·ρ·V a R_A·V a R_B)

Stand-alone VaR of A	Value at Risk of position A
Stand-alone VaR of B	Value at Risk of position B
Correlation, rho	Correlation, at most one

IN THE LESSON

- The failure is not universal.
- For elliptical distributions, the normal being the main example, V a R is subadditive.
- With zero means, the combined V a R of two positions follows a rule that looks just like the variance of two assets.
- Square the two stand-alone V a Rs, add twice the correlation times both of them, and take the square root.
- Since correlation can be at most one, the combined figure is at most the plain sum, with equality only at perfect correlation.
- This is why the problem hides in practice: under normal assumptions V a R happens to be subadditive.
- The trouble comes when the real losses are skewed, discrete or fat-tailed, and then a firm could cut its reported risk simply by splitting into smaller units.

30 Example 11: combining normal V a Rs

A: $4m, B: $3m, normal, zero means. Sum of the two = $7m

Correlation 0.5	√(16 + 9 + 12)	**$6.08m**
Correlation 1	√(16 + 9 + 24)	**$7.00m**
Correlation 0	√(16 + 9 + 0)	**$5.00m**
Compared with the sum	never above $7m	**Subadditive**

IN THE LESSON

- Let us try it.
- Two positions have stand-alone V a Rs of four million and three million dollars, with zero means and normal returns.
- First suppose the correlation is a half.
- Four squared is sixteen, three squared is nine, and twice a half times four times three adds twelve.
- The total is thirty seven, and its square root is about six point zero eight million.
- That is less than the plain sum of seven, so diversification helped.
- At perfect correlation of one, the extra term is twenty four, the total is forty nine, and the root is exactly seven.
- At zero correlation the total is twenty five, and the root is five.
- The combined figure never exceeds seven, so for normal returns subadditivity holds, unlike the bond example.

31 Recap: the main ideas

- Correlation below 1 diversifies; frontier = upper edge
- Real returns: fat tails, negative skew, clustering
- VaR: threshold. ES: average beyond it, never smaller
- Coherent: four properties; ES passes, VaR fails
- VaR fails subadditivity: two bonds, $0 against $10m

IN THE LESSON

- Let us pull it together.
- Portfolios are described by return and volatility, and because correlation is below one, mixing assets reduces risk.
- The efficient frontier is the top half of the minimum variance curve, and the best risky mix with a safe asset has the highest Sharpe ratio.
- Real returns have fat tails, negative skewness and clustering, so normal models understate extreme losses.
- V a R is a quantile, a threshold and not a worst case, and its square-root-of-time scaling needs strong assumptions.
- Expected shortfall averages the losses beyond V a R, is never smaller, and is coherent.
- A coherent measure is monotone, homogeneous, translation invariant and subadditive.
- V a R fails subadditivity, as the two bond example showed.

32 Try three quick questions

1. Two assets, each 10% volatility, correlation 0, half each. Portfolio volatility?
   Answer: Variance = 0.25×0.01 + 0.25×0.01 = 0.005, so volatility = 7.07%.
2. One-day 99% VaR is $2m. Five-day VaR by square-root-of-time?
   Answer: 2m × √5 = $4.47m, if returns are independent and volatility constant.
3. Which coherence property does VaR fail, and what shows it?
   Answer: Subadditivity. Two bonds: each VaR $0, combined VaR $10m.

IN THE LESSON

- Pause the video, think about each question, then listen to the answer.
- Question one.
- Two assets each have ten per cent volatility, with zero correlation and half your money in each.
- What is the portfolio volatility?
- The variance is zero point two five times zero point zero one, twice, which is zero point zero zero five, so volatility is seven point oh seven per cent.
- Question two.
- A one day ninety nine per cent V a R is two million dollars.
- What is the five day V a R by the square root of time rule?
- Two million times the square root of five is four point four seven million dollars, if returns are independent and volatility is constant.
- Question three.
- Which coherence property does V a R fail, and what shows it?
- Subadditivity.
- Two bonds each have a V a R of zero, yet combined the V a R is ten million dollars.