CHARTA LEARNING · LESSON NOTES

Market-Driven Scenarios

FRM Part II

1 A story: the stress test that made no sense

- Committee: what if shares fall 15% in a month?
- Ad hoc answer: shock shares, hold everything else still
- But in a sell-off, credit spreads widen and yields move
- A scenario must be plausible and coherent

IN THE LESSON

- Picture the risk committee of the Halden Balanced Fund.
- The chair asks a simple question: what happens to us if shares fall fifteen per cent in a month?
- A junior analyst gives an ad hoc answer.
- Ad hoc means built for the occasion, with no method behind it.
- He shocks the share holdings and holds every other market variable still.
- The answer looks precise, but think about it.
- When shares crash, investors demand more to hold company bonds, so credit spreads widen, and they rush into government bonds, so yields move.
- A month in which only shares move is one that markets have rarely shown.
- What the committee needs is a scenario that is plausible, meaning believable, and coherent, meaning its parts fit together.

2 The market-driven scenario approach

- Shock only a few key variables, e.g. an equity index
- Every other variable moves as history implies
- That move: its conditional expectation given the shocks
- Aim: a plausible, internally coherent scenario
- No crash with spreads still, unless the data show it

IN THE LESSON

- The market driven scenario approach, M D S for short, solves this.
- You shock only a few key market variables, for example an equity index or a credit spread.
- Every other variable then moves by the amount that its historical co-movement with the shocked variables implies.
- In statistical language, that move is the variable's conditional expectation given the shocks, meaning the average value the variable takes, given the shocks, under the historical covariances.
- The aim is a scenario that is plausible and internally coherent.
- Equities do not crash while credit spreads sit still, unless the data show that this has happened.
- So the person designing the test chooses the story and the key shocks, and history fills in the rest.

3 Three ways to build a scenario

Three ways Halden could build a stress scenario

Approach	How it is built	Weakness
Ad hoc	Shock one, hold the rest	Incoherent
Historical replay	Repeat a past episode	Only the past
Market-driven	Shock a few; rest follow	Model-dependent

IN THE LESSON

- There are three broad ways Halden could build a scenario, and the exam asks you to contrast them.
- The ad hoc way shocks one factor and holds the rest still, as our analyst did.
- Its weakness is that it is incoherent, because it ignores how other variables move with that factor.
- The historical replay repeats exactly the moves that every variable made in a real past episode, such as a crisis month.
- It is coherent, since it really happened, but it can only show the past, in the exact form it took.
- The market driven way shocks a few variables and lets the rest follow their historical relationships, using the average pattern across the whole data window rather than copying one episode.
- It is coherent and flexible, but it depends on a model: the estimated covariances.

4 Merits of the MDS approach

- Consistent with observed market relationships
- Explores events that never happened exactly as stated
- Transparent link from the story to the numbers

IN THE LESSON

- What are the merits?
- First, the scenario is consistent with observed market relationships, so its parts fit together the way markets have really behaved.
- Second, it can explore events that never occurred exactly as stated.
- A historical replay can only show the past.
- With M D S, Halden can ask about a fall of fifteen per cent in shares even if no past month looked just like that, and still get a coherent set of moves.
- Third, there is a transparent link between the narrative and the numbers.
- Anyone can see which variables were shocked, by how much, and why the rest moved as they did.

5 Limits of the MDS approach

- Depends on the estimated covariance matrix
- Roughly linear and normal framework
- Plausible is not the same as probable
- Shows the average companion move, not the worst

IN THE LESSON

- The approach also has limits.
- First, everything depends on the estimated covariance matrix, the table of volatilities and correlations taken from history.
- A different data window gives a different scenario.
- Second, the framework is roughly linear and normal.
- It assumes each variable responds in proportion to the shock, with bell shaped randomness around that response.
- Real markets in a crisis are not so tidy.
- Third, plausible is not the same as probable.
- A plausible scenario is one whose moves fit together.
- That says nothing about how likely the whole event is, and a coherent scenario can still be very rare.
- Fourth, because the unshocked variables sit at their conditional expectations, they show the average companion move, not the worst one that could come with the shock.

6 z-scores: moves in standard deviations

- A z-score expresses a move in standard deviations
- Individual z-scores say how large each move is
- Same scale for every variable
- One is an ordinary month; three is a rare one

IN THE LESSON

- To judge a scenario we need a common ruler.
- A fall of fifteen per cent in shares and a rise of forty basis points in a credit spread are in different units.
- A basis point is one hundredth of one per cent.
- The z score solves this by expressing every move in standard deviations.
- Individual z scores then say how large each move is, on the same scale for every variable.
- A z score of one is an ordinary month.
- A z score of three is a rare one.

7 The z-score and the horizon

z = shock ÷ σ over the horizon
σ monthly = σ annual ÷ √12

Shock	The move the scenario assumes
Volatility over the horizon	Volatility for the same horizon
z-score	Size of the move in SDs

IN THE LESSON

- Here is the formula.
- The z score equals the shock divided by sigma, the volatility measured over the same horizon as the scenario.
- Volatilities are often quoted per year, so they must be scaled.
- If returns are independent from month to month, volatility grows with the square root of time.
- Monthly volatility is annual volatility divided by the square root of twelve, which is about three point four six.
- So the shock is the move assumed, sigma is the volatility for the same horizon, and z is the size of the move in standard deviations.
- The common error is to divide a monthly shock by an annual volatility, which makes the move look far smaller than it is.

8 Example 1: Halden's z-scores

Scenario A over one month: shares −10%, credit spread +40bp

Share index: monthly vol	17.32% ÷ √12	**5.0%**
Share z-score	−10% ÷ 5%	**−2.0**
Credit spread z-score	+40bp ÷ 20bp monthly vol	**+2.0**
Wrong: annual vol for shares	−10% ÷ 17.32%	**−0.58**

IN THE LESSON

- Halden's committee proposes Scenario A, over one month: shares fall ten per cent and the credit spread on corporate bonds widens by forty basis points.
- The share index has an annual volatility of seventeen point three two per cent.
- Divided by the square root of twelve, its monthly volatility is five per cent.
- So the share z score is minus ten divided by five, which is minus two.
- The credit spread has a monthly volatility of twenty basis points, so its z score is forty divided by twenty, plus two.
- Both moves are two standard deviations: large, but not extreme.
- Now the trap.
- Dividing the monthly fall of ten per cent by the annual volatility gives minus point five eight, which would make a severe month look almost routine.

9 The volatility z-score, and what it misses

- Volatility z-score: size judged on volatilities alone
- Scenario A: $\sqrt{(2^2 + 2^2)} = \sqrt{8} = 2.83$
- It ignores co-movement between the variables
- Scenario B: shares −10%, spread −40bp: also 2.83
- Same size, but B breaks the usual pattern

IN THE LESSON

- The volatility z score reflects the size of the whole scenario judged on volatilities alone.
- One simple way to build it, which we use here, is to square each z score, add them up, and take the square root.
- For Scenario A that is the square root of four plus four, about two point eight three.
- The trouble is that it ignores co-movement between the variables.
- Consider Scenario B: shares fall ten per cent, but the credit spread tightens by forty basis points.
- Its z scores are minus two and minus two, so its volatility z score is also two point eight three.
- The same size.
- Yet in a sell-off spreads normally widen, so B breaks the usual pattern.
- We need a measure that notices.

10 The Mahalanobis distance

$D = \sqrt{(z' R^{-1} z)}$
Two variables: $Dsq = (z_1^2 + z_2^2 - 2\rho z_1 z_2) \div (1 - \rho^2)$

Correlation matrix, R	Historical correlations
Vector of z-scores, z	One z-score per variable
Correlation, rho	Correlation of the two
Mahalanobis distance, D	Large if moves fight correlation

IN THE LESSON

- That measure is the Mahalanobis distance, written D.
- It uses R, the correlation matrix, which holds the historical correlation between every pair of variables.
- Here z is the vector, meaning the list, of the individual z scores, one for each variable.
- In matrix form, D is the square root of the quantity z transposed, times the inverse of R, times z.
- For two variables there is a simple version.
- The squared distance, written D squared, equals z one squared plus z two squared, minus two rho z one z two, all divided by one minus rho squared.
- Rho is the correlation between the two variables.
- The key is the middle term.
- D penalises combinations that break historical correlations, so it is large when the moves fight the correlation.

11 Example 2: the distance of Scenario A

Scenario A: $z = -2$ for shares, $+2$ for the spread. Correlation $\rho = -0.6$

Sum of squares	$(-2)^2 + 2^2$	**8**
Cross term	$-2 \times (-0.6) \times (-2) \times 2$	**−4.8**
Dsq: divide by $1 - \rho^2$ = 0.64	$(8 - 4.8) \div 0.64$	**5.0**
Distance D	$\sqrt{5}$	**2.24**

IN THE LESSON

- Now compute Scenario A's distance.
- Shares and the credit spread have a historical correlation of minus point six: when shares fall, spreads tend to widen.
- The z scores are minus two and plus two.
- The sum of squares is four plus four, eight.
- The cross term is minus two, times rho of minus point six, times minus two, times two, which is minus four point eight.
- One minus rho squared is one minus point three six, which is point six four.
- So D squared is eight minus four point eight, which is three point two, divided by point six four, which gives five.
- D is the square root of five, about two point two four.
- That is below the volatility z score of two point eight three, because the two moves go together the way history says they should.

12 Example 3: the distance of Scenario B

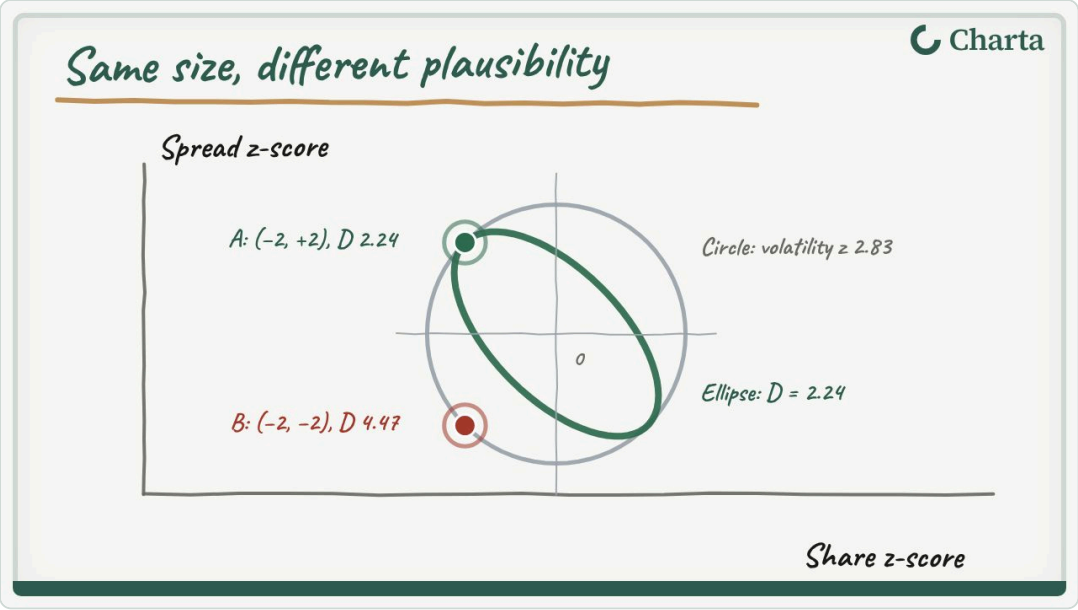
Scenario B: $z = -2$ for shares, -2 for the spread. Same $\rho = -0.6$

Sum of squares	$(-2)^2 + (-2)^2$	**8**
Cross term	$-2 \times (-0.6) \times (-2) \times (-2)$	**+4.8**
Dsq: divide by 0.64	$(8 + 4.8) \div 0.64$	**20.0**
Distance D	$\sqrt{20}$	**4.47**

IN THE LESSON

- Scenario B has individual z scores of the same size: minus two for shares, but minus two for the spread as well.
- The sum of squares is again eight.
- The cross term is minus two, times minus point six, times minus two, times minus two, which is plus four point eight.
- So D squared is eight plus four point eight, which is twelve point eight, divided by point six four, which gives twenty.
- D is the square root of twenty, about four point four seven.
- That is twice Scenario A's distance, and well above the volatility z score of two point eight three.
- Nothing about the size of each move changed.
- What changed is the combination: a share crash with tighter spreads goes against the correlation, so the distance treats it as far more remote.

13 Same size, different plausibility



IN THE LESSON

- This picture shows the two scenarios.
- Along the bottom is the share z score, and up the side is the spread z score.
- The grey circle joins every combination with a volatility z score of two point eight three.
- The green ellipse joins every combination with Scenario A's Mahalanobis distance of two point two four.
- It is tilted, because shares down tends to come with spreads up.
- The two thin lines cross at zero, a month with no unusual move at all.
- Both scenarios sit on the circle, so on volatilities alone they are the same size.
- Scenario A, shares down and spreads up, sits on the ellipse.
- Scenario B, shares down and spreads down, lies far outside it, in a direction that history rarely goes.

14 Three scores and what each tells you

Halden's two scenarios on each score

Score	What it measures	A	B
Individual z	Size of each move	±2	±2
Volatility z	Size, ignoring co-movement	2.83	2.83
Scenario z (D)	Size and fit together	2.24	4.47
Correlation z	Is the mix typical?	Fits	Fights

IN THE LESSON

- Let us line the scores up for Halden's two scenarios.
- The individual z scores measure the size of each move: two in size, in both scenarios.
- The volatility z score measures the size of the scenario on volatilities alone, ignoring co-movement: two point eight three for both.
- Sources define these scores slightly differently, so here we take the scenario z score to be the Mahalanobis distance.
- It measures size and fit to the correlations together: two point two four for A and four point four seven for B.
- Finally, we read the correlation z score as whether D sits below or above the volatility z score, which tells us if the mix of moves is typical.
- For A, D is lower, so the mix fits history.
- For B, D is higher, so the mix fights it.

15 Example 4: how rare is each scenario?

Joint normality: Dsq is chi-square with 2 degrees of freedom

Expected Dsq	= number of variables	**2**
A: chance Dsq is 5 or more	exp(−5 ÷ 2)	**8.2%**
B: chance Dsq is 20 or more	exp(−20 ÷ 2)	**0.0045%**
B is rarer by a factor of	8.2% ÷ 0.0045%	**about 1,800**

IN THE LESSON

- Under joint normality, D squared follows a chi square distribution with as many degrees of freedom as there are variables.
- Its expected value equals the number of variables, so for our two variables an ordinary month has D squared of about two.
- With two degrees of freedom, the chance that D squared is at least some value is e to the power of minus half that value.
- For Scenario A, D squared of five gives e to the minus two point five, about eight point two per cent of months.
- For Scenario B, D squared of twenty gives e to the minus ten, about point zero zero four five per cent.
- So B is rarer by a factor of about one thousand eight hundred.

16 Example 5: the exam trap

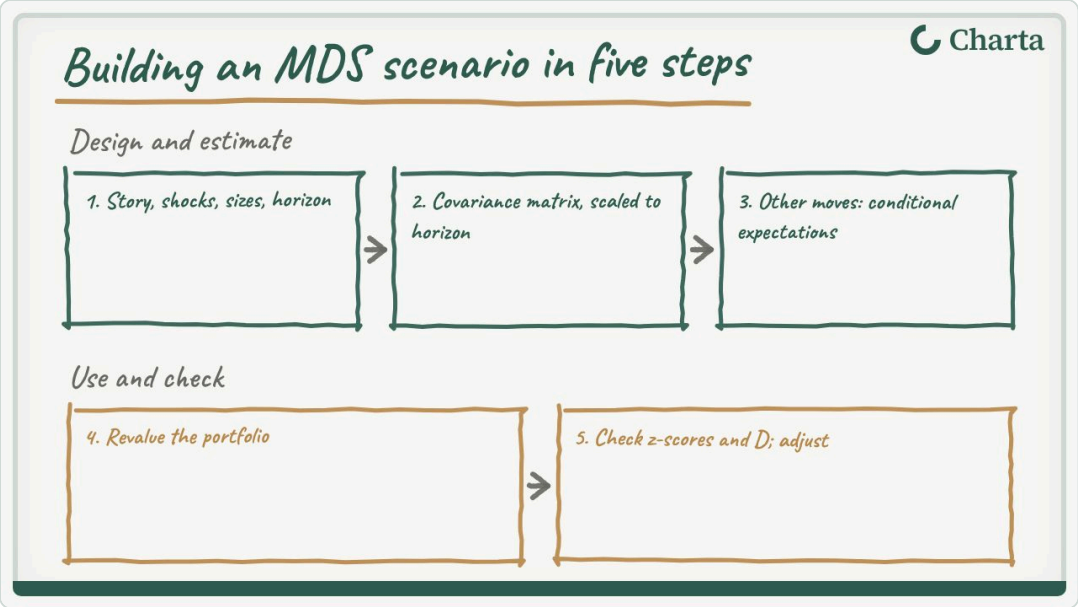
US and European share indices, ρ = 0.9. Moves: z = +1.5 and −1.5

Each z-score alone	modest	**±1.5**
Sum of squares	$1.5^2 + 1.5^2$	**4.5**
Cross term	−2 × 0.9 × 1.5 × (−1.5)	**+4.05**
Dsq: divide by 1 − 0.81	(4.5 + 4.05) ÷ 0.19	**45.0**
Distance D	√45	**6.71**

IN THE LESSON

- Here is the classic exam trap.
- Take two share indices, U S and European, with a correlation of point nine.
- A scenario has U S shares up one and a half standard deviations and European shares down one and a half.
- Each z score alone looks modest.
- The sum of squares is four point five.
- The cross term is minus two, times point nine, times one point five, times minus one point five, which is plus four point zero five.
- One minus point nine squared is point one nine.
- So D squared is eight point five five divided by point one nine, which is forty five, and D is about six point seven one.
- Opposite moves in highly positively correlated assets give a large distance, even when each z score is modest.

17 Building an MDS scenario in five steps



IN THE LESSON

- Now we build a scenario properly.
- The process has five steps, in two groups.
- First, design and estimate.
- Step one: define the narrative and choose the key variables to shock, the size of each shock, and the horizon.
- Step two: estimate the covariance matrix of the risk factors over a suitable window, and scale it to that horizon.
- Step three: derive the unshocked variables as conditional expectations given the shocks.
- Then use and check.
- Step four: revalue the portfolio with the full set of factor moves.
- Step five: assess plausibility with z scores and the Mahalanobis distance, and adjust the shocks if the scenario is too remote or too mild.
- Let us take Halden through all five.

18 Example 6: Halden's inputs, steps one and two

Story: a sudden equity sell-off over one month.
Key shock: shares −15%

Shares: the shocked variable	monthly vol 5%	**z = −3.0**
Credit spread	vol 20bp, ρ −0.6	**widens**
10-year Treasury yield	vol 25bp, ρ +0.4	**falls**
Euro against the dollar	vol 2.5%, ρ +0.3	**falls**

IN THE LESSON

- Step one for Halden.
- The story is a sudden equity sell-off over one month, and the single key shock is a fall of fifteen per cent in the share index.
- With a monthly volatility of five per cent, that is a z score of minus three.
- Step two: Halden estimates monthly volatilities and correlations.
- The credit spread has a monthly volatility of twenty basis points and a correlation of minus point six with shares, so it should widen.
- The ten year Treasury yield has a volatility of twenty five basis points and a correlation of plus point four, so it should fall, as investors buy government bonds.
- The euro, measured in dollars, has a volatility of two point five per cent and a correlation of plus point three, so it should fall too.

19 The conditional expectation of one move

Move in y = ρ × (σy ÷ σx) × shock to x
In z-scores: zy = ρ × zx

Correlation, rho	Correlation of x and y
Volatility of y	Same horizon, units of y
Volatility of x	Of the shocked variable
Shock to x	The move chosen for x

IN THE LESSON

- Step three uses this formula.
- For one shocked variable x, the expected move of another variable y equals rho, the correlation between them, times the ratio of their volatilities, sigma y over sigma x, times the shock to x.
- It is the slope of a regression of y on x, applied to the shock.
- The volatility ratio converts the units: it turns a move in share returns into basis points of spread or per cent of currency.
- In z score terms the formula is even simpler: the z score of y equals rho times the z score of x.
- So a shock of minus three standard deviations to shares implies minus three times rho standard deviations for each other variable.
- The exam trap is using only the correlation, without the volatility ratio.

20 Example 7: propagating the shock, step three

Shares −15%, monthly vol 5%. Move = ρ × (σy ÷ σx) × shock

Credit spread	−0.6 × (20 ÷ 5) × (−15)	**+36bp**
10-year Treasury yield	0.4 × (25 ÷ 5) × (−15)	**−30bp**
Euro against the dollar	0.3 × (2.5 ÷ 5) × (−15)	**−2.25%**
Check in z-scores: spread	−0.6 × (−3) = 1.8; × 20bp	**+36bp**
Trap: correlation only	−0.6 × (−15)	**9: wrong**

IN THE LESSON

- Now propagate the shock.
- For the credit spread: minus point six, times twenty over five, which is four basis points of spread for each one per cent move in shares, times minus fifteen, gives plus thirty six basis points.
- For the ten year yield: point four, times twenty five over five, which is five, times minus fifteen, gives minus thirty basis points.
- For the euro: point three, times two point five over five, which is a half, times minus fifteen, gives minus two point two five per cent.
- Check the spread in z scores: minus point six times minus three is one point eight standard deviations, and one point eight times twenty basis points is again thirty six.
- Now the trap: using only the correlation, minus point six times minus fifteen gives nine, a figure with the wrong units and far too small.

21 Example 8: revaluing Halden, step four

Halden $500m: shares 200, corporate bonds 200, Treasuries 50, euro assets 50

Shares	200 × (−15%)	**−$30.00m**
Corp. bonds, mod. duration 6	−200 × 6 × (0.36% − 0.30%)	**−$0.72m**
Treasuries, mod. duration 8	−50 × 8 × (−0.30%)	**+$1.20m**
Euro assets, unhedged	50 × (−2.25%)	**−$1.125m**
Total	−30.645 ÷ 500 = −6.1%	**−$30.645m**

IN THE LESSON

- Step four revalues the fund with the full set of moves.
- Halden has five hundred million dollars: two hundred million in shares, two hundred million in corporate bonds, fifty million in Treasuries and fifty million in euro assets.
- Shares lose fifteen per cent of two hundred, which is thirty million.
- Corporate bonds have a modified duration of six, and their yield moves by the spread plus the Treasury yield: thirty six minus thirty, a net rise of six basis points.
- Six times point zero six per cent of two hundred million is a loss of point seven two million.
- Treasuries, with a modified duration of eight, gain one point two million as yields fall thirty basis points.
- Assume the euro assets are unhedged, so only the currency move matters: they lose two point two five per cent, one point one two five million.
- The total loss is thirty point six four five million, about six point one per cent.

22 Example 9: is it plausible? Step five

Shares z = −3.0. Spread z = +1.8 (MDS) or 0 (ad hoc). ρ = −0.6

MDS: Dsq	(9 + 3.24 − 6.48) ÷ 0.64	**9.0**
MDS: distance D	√9	**3.0**
Ad hoc, spread unchanged: Dsq	(9 + 0 − 0) ÷ 0.64	**14.06**
Ad hoc: distance D	√14.06	**3.75**

Moves set at their conditional expectations add nothing to D

IN THE LESSON

- Step five checks plausibility, using shares and the credit spread.
- In the market driven scenario, the z scores are minus three and plus one point eight.
- D squared is nine plus three point two four, minus six point four eight for the cross term, all divided by point six four, which is nine.
- So D is exactly three, the same as the share shock alone.
- Now take the ad hoc scenario, with the spread unchanged, a z score of zero.
- D squared is nine divided by point six four, about fourteen point zero six, and D is three point seven five.
- The ad hoc scenario is less plausible, not more.
- The general result: moves set at their conditional expectations add nothing to D.
- If three is too remote or too mild, the share shock is resized and the steps rerun.

23 Example 10: two shocks at once

Shocks: shares z = −3, spread z = +3, ρ = −0.6. Yield: ρ 0.4 and −0.3

Wrong: add the two effects	0.4 × (−3) − 0.3 × 3	**z = −2.10**
Inverse-R weight on shares	(−3 + 0.6 × 3) ÷ 0.64	**−1.875**
Inverse-R weight on spread	(3 − 0.6 × 3) ÷ 0.64	**+1.875**
Right: yield z-score	0.4 × (−1.875) − 0.3 × 1.875	**z = −1.31**
Yield move, vol 25bp	−1.3125 × 25 vs −2.1 × 25	**−32.8bp, not −52.5**

IN THE LESSON

- Now a credit led story shocks two variables: shares, with a z score of minus three, and the spread, with plus three.
- Their correlation is minus point six.
- The yield has a correlation of point four with shares and minus point three with the spread.
- Simply adding the two effects gives point four times minus three, minus point three times three, a z score of minus two point one.
- That double counts, because part of the spread shock is already explained by the share shock.
- The fix is to invert the correlation matrix of the shocked variables.
- The weight on shares is minus three plus point six times three, over point six four: minus one point eight seven five.
- The weight on the spread is three minus point six times three, over point six four: plus one point eight seven five.
- The yield z score is then minus one point three one, a move of minus thirty two point eight basis points, not minus fifty two point five.

24 What the propagated moves are, and are not

- Linear, model-based: built from historical covariance
- Not tail dependence: the average companion move
- Stressed dependence is a separate, explicit choice
- Made by shocking the correlations directly

IN THE LESSON

- Be clear about what these propagated moves are.
- They are a linear, model based approximation built from the historical covariance matrix.
- They are not a measure of tail dependence, which is how variables move together in extreme events specifically.
- Each derived move is the average companion move given the shock, estimated mostly from ordinary months.
- If Halden believes that dependence is stronger in a crisis, that stressed dependence is a separate, explicit choice.
- It is made by shocking the correlations directly, for example by replacing the historical correlation between shares and spreads with a higher crisis value before propagating the shock.
- The scenario is then labelled clearly as using stressed correlations, so the committee knows which choice it is looking at.

25 Challenge 1: which history?

- The covariance matrix is estimated from history
- Window, weighting, calm or stressed periods drive results
- Correlations typically rise in crises
- Calm-period estimates understate co-movement

IN THE LESSON

- Now the practical challenges.
- The first is which history to use.
- The covariance matrix is estimated from past data, so the choices made in estimating it drive the results: the window, meaning how many years of data, the weighting, for example giving recent months more weight, and whether the window spans calm or stressed periods.
- This matters because correlations typically rise in crises.
- Assets that move fairly independently in quiet times start to fall together when investors rush for safety.
- So a matrix estimated in a calm period understates co-movement, and the scenario it produces will look gentler than a real crisis would be.
- Let us measure how much this matters for Halden.

26 Example 11: the same shock, three windows

Shares −15%; spread vol 20bp. Spread loss on $200m of bonds, duration 6

Calm window: ρ = −0.3	−0.3 × 4 × (−15) = 18bp	**loss $2.16m**
Full history: ρ = −0.6	−0.6 × 4 × (−15) = 36bp	**loss $4.32m**
Crisis window: ρ = −0.8	−0.8 × 4 × (−15) = 48bp	**loss $5.76m**
Crisis versus calm	5.76 ÷ 2.16	**2.67 times**

IN THE LESSON

- Keep Halden's share shock of minus fifteen per cent and the spread volatility of twenty basis points, so the volatility ratio is four.
- Look only at the spread effect on two hundred million dollars of corporate bonds with a duration of six.
- With a calm window, the correlation is minus point three.
- The spread widens by eighteen basis points, and the loss is two hundred times six times point one eight per cent, which is two point one six million.
- With the full history, the correlation of minus point six gives thirty six basis points and a loss of four point three two million.
- With a crisis window, minus point eight gives forty eight basis points and a loss of five point seven six million.
- Crisis versus calm is five point seven six over two point one six: the same story gives a loss two point six seven times larger, purely from the window.

27 Challenge 2: fat tails and the distance

- Real returns: fat tails and nonlinear dependence
- Distance and chi-square assume joint normality
- So they overstate how rare a scenario is
- A's 8.2% chance of a month this remote may be too low
- Use the probabilities to rank, not as precise odds

IN THE LESSON

- The second challenge is the shape of real returns.
- Real returns have fat tails, meaning extreme moves happen more often than a normal distribution allows, and nonlinear dependence, meaning variables can be loosely linked in normal times but tightly linked in a crash.
- The Mahalanobis distance and its chi square probabilities assume joint normality.
- So they overstate how rare a scenario is.
- Halden's Scenario A had a chance of about eight point two per cent that a month is at least this remote, and in fat tailed markets moves that remote may come along more often than that.
- The practical lesson is to use the probabilities to rank scenarios against each other, not as precise odds.

28 Challenge 3: too many factors

- Many factors, few observations: a noisy covariance matrix
- Example: 100 factors, 60 months: the matrix is singular
- D and several-shock propagation need its inverse
- Fix: limit the factors or regularise the matrix

IN THE LESSON

- The third challenge is size.
- With many risk factors and few observations, the sample covariance matrix is noisy, and it can be singular, meaning it has no inverse.
- Suppose Halden used one hundred factors and sixty months of data.
- Sixty observations cannot pin down a hundred dimensions, so the matrix is singular.
- Yet both the Mahalanobis distance and the propagation of several shocks need an inverse, so inversion fails or becomes unstable.
- The fix is to limit the number of factors, for example to a handful of key indices, or to regularise the matrix.
- Regularising means adjusting the estimate, for instance shrinking it toward a simpler structure, so that it can be inverted reliably.

29 Challenge 4: options, thresholds and horizons

- Linear averages: options need full revaluation
- Add volatility factors when options matter
- Plausibility threshold: a judgment, not a bright line
- Horizon must match the data frequency
- MDS: a disciplined aid to judgment, not a forecast

IN THE LESSON

- First, conditional expectations are linear averages, but a portfolio with options responds nonlinearly.
- If Halden held put options as protection, their value would depend on the size of the fall and on implied volatility, which usually jumps in a sell-off.
- So option positions need full revaluation, meaning each option is priced again under the scenario, and perhaps volatility factors added to the scenario.
- Second, the plausibility threshold is a judgment, not a statistical bright line.
- No rule says that a distance of three is acceptable and four is not.
- Third, the horizon must match the data frequency: a monthly scenario needs monthly volatilities and correlations.
- The exam tests whether you see M D S as a disciplined aid to judgment, not an objective forecast.

30 Mistakes to avoid

Common errors in market-driven scenario questions

Mistake	Correct approach
Monthly shock ÷ annual vol	Use the vol for the same horizon
Move = ρ × shock	Move = ρ × (σy ÷ σx) × shock
Add effects of several shocks	Invert the shocked correlations
Judge only individual z-scores	Use D for the combination
Plausible means probable	Plausible means coherent

IN THE LESSON

- Here are the mistakes to avoid.
- First, dividing a monthly shock by an annual volatility: use the volatility for the same horizon as the scenario.
- Second, writing the move as rho times the shock: the move is rho, times sigma y over sigma x, times the shock.
- Third, adding up the effects of several shocks: invert the correlation matrix of the shocked variables, so their mutual correlation is not double counted.
- Fourth, judging a scenario only by its individual z scores: use the Mahalanobis distance for the combination, since opposite moves in highly correlated assets can be extreme.
- And fifth, treating plausible as probable: plausible means the moves are coherent with each other, not that the event is likely.

31 Recap: the main ideas

- MDS: shock a few; the rest at conditional expectations
- z = shock ÷ vol for the same horizon
- D: size, and whether the moves fit the correlations
- Steps: story, covariance, propagate, revalue, check
- Halden: shares −15% gives a coherent loss of 6.1%

IN THE LESSON

- Let us recap.
- The market driven scenario approach shocks a few key variables and sets every other variable at its conditional expectation, so the scenario is coherent with history.
- A z score measures each move in standard deviations, using the volatility for the same horizon.
- The Mahalanobis distance combines the z scores through the correlation matrix, so it measures both the size of a scenario and whether its moves fit together.
- A scenario is built in five steps: the story, the covariance matrix, propagation, revaluation and the plausibility check.
- For Halden, a fall of fifteen per cent in shares became a coherent scenario with a loss of about six point one per cent.

32 Try three quick questions

1. Annual vol 20%. Index falls 25% in one quarter. z-score?
   Answer: Quarterly vol = 20% ÷ √4 = 10%. z = −2.5.
2. ρ = 0.5; z-scores +2 and −2. Mahalanobis distance?
   Answer: Dsq = (4 + 4 + 4) ÷ 0.75 = 16. D = 4.
3. Oil +20%, vol 8%. Airline vol 6%, ρ = −0.5. Expected move?
   Answer: −0.5 × (6 ÷ 8) × 20 = −7.5%.

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- A share index has an annual volatility of twenty per cent and falls twenty five per cent over one quarter.
- What is the z score?
- Quarterly volatility is twenty divided by the square root of four, ten per cent, so z is minus two point five.
- Question two.
- Two variables have a correlation of point five, and z scores of plus two and minus two.
- What is the Mahalanobis distance?
- D squared is four plus four plus four, twelve, divided by point seven five, which is sixteen, so D is four.
- Question three.
- Oil rises twenty per cent, with a monthly volatility of eight per cent.
- An airline index has a monthly volatility of six per cent and a correlation of minus point five with oil.
- What is its expected move?
- Minus point five, times six over eight, times twenty: minus seven point five per cent.