CHARTA LEARNING · LESSON NOTES

Risk Measurement and Assessment

FRM Part II

1 A story: the loss nobody recorded

- A failed payment file is reloaded: half is sent twice
- 12 duplicates of $15,000 each: $180,000 in all
- Each one sits below the $20,000 recording threshold
- No loss event is recorded at all
- Loss data are the factual base for measurement

IN THE LESSON

- At Fernhill Bank, a file of customer payments fails halfway, and an operator reloads it.
- Nobody notices that the first half is now sent twice.
- Twelve payments of fifteen thousand dollars each go out a second time, one hundred and eighty thousand dollars in all.
- Fernhill records operational losses only above a threshold of twenty thousand dollars.
- Each duplicate, looked at on its own, sits below that line.
- So no loss event is recorded at all, and the risk team never learns that the payments process has a hole.
- The lesson starts here because loss data are the factual base for measurement, so how losses are gathered matters as much as the models built on them.

2 What a loss record holds

One firm-wide taxonomy for every event

Field	Fernhill's duplicate payments
Event type	Execution and process error
Business line	Retail payments
Root cause	Failed file reloaded, no check
Three dates	Occurred, discovered, booked
Amounts	Gross, recoveries, net: kept apart

IN THE LESSON

- Good practice records every event in one firm wide taxonomy, a shared list of categories, so that losses from different desks can be added up and compared.
- Here is what Fernhill's record of the duplicate payments should hold.
- The event type is an execution and process error.
- The business line is retail payments.
- The root cause is the underlying reason, here a failed file reloaded without any check.
- Then come three separate dates.
- The occurrence date is when the event happened, the discovery date is when the bank found it, and the accounting date is when the loss was booked.
- These can be weeks or months apart.
- Finally the amounts.
- Gross loss, recoveries and net loss are stored separately, never as one netted figure.

3 Example 1: group by root cause first

12 duplicates of $15,000 from one reload. Threshold $20,000

Each fragment alone	$15,000 < $20,000	**not recorded**
Grouped: one root cause	12 × $15,000	**$180,000 gross**
Recoveries	payments recalled	**$120,000**
Net loss	180,000 − 120,000	**$60,000**
Grouped event vs threshold	180,000 > 20,000	**recorded**

IN THE LESSON

- Tested one by one, each fragment of fifteen thousand dollars is below the twenty thousand dollar threshold, so nothing is recorded.
- But the twelve duplicates share one root cause, the reloaded file.
- Losses from one root cause are grouped into a single event before any threshold is applied.
- So the event is twelve times fifteen thousand, a gross loss of one hundred and eighty thousand dollars.
- Fernhill then recalls some payments from the receiving banks and recovers one hundred and twenty thousand dollars.
- The net loss is one hundred and eighty thousand minus one hundred and twenty thousand, which is sixty thousand dollars.
- Now test the grouped event against the threshold.
- One hundred and eighty thousand is well above twenty thousand, so the event is recorded.

4 Beyond booked losses

- Near miss: could have caused a loss, but did not
- Non-accounted impact: a cost not booked as a loss
- Both warn of control weaknesses
- Fernhill: a second reload, caught by a clerk
- Capture them in the same taxonomy

IN THE LESSON

- Not every useful event costs money.
- A near miss is an event that could have caused a loss but did not, this time.
- A non accounted impact is a cost that never appears in the general ledger as a loss, such as staff overtime spent fixing an error.
- Why capture them?
- Because both warn of control weaknesses before the big loss arrives.
- Suppose that a month later another Fernhill operator reloads a failed file, and a sharp clerk spots the duplicates before they are released.
- No money is lost.
- Yet it shows the same hole as before.
- Good practice captures near misses and non accounted impacts in the same taxonomy as booked losses.

5 What supervisors expect

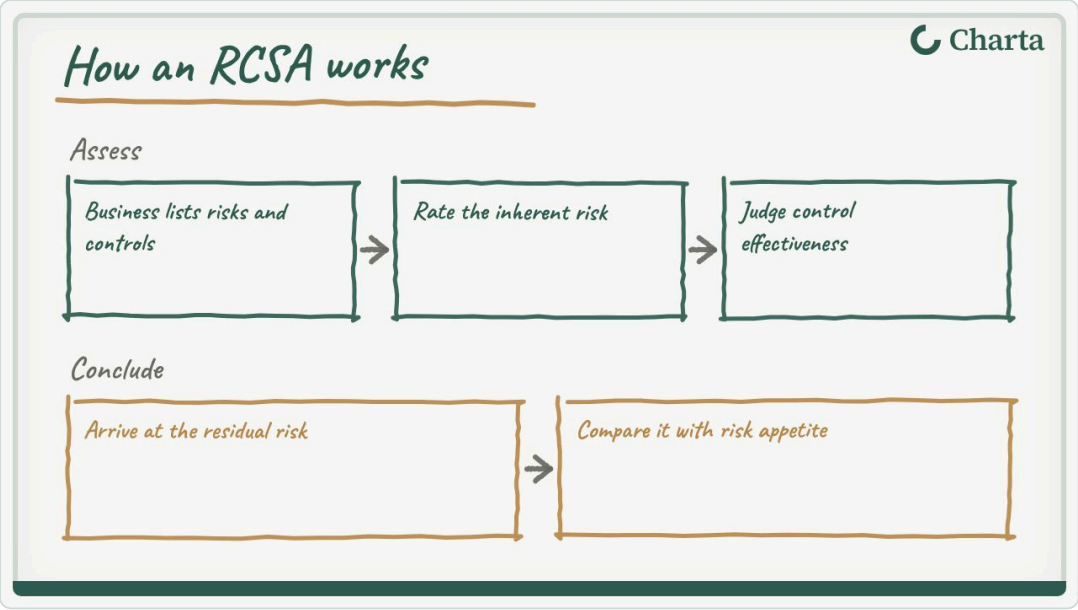
Data-quality expectations, and the usual failure

Expectation	What it means
Completeness	Every event, all units
Accuracy, reconciliation	Ties back to the general ledger
Independent validation	Checked outside the unit
Timely escalation	Big events reach senior management
Usual failure	Under-reporting: blame, fragments

IN THE LESSON

- What do supervisors expect from a loss data process?
- First, completeness: every event is captured, across all business units.
- Second, accuracy and reconciliation: the recorded losses must tie back to the general ledger, so that nothing booked in the accounts is missing from the loss database.
- Third, independent validation: the data are checked by people outside the unit that reported them.
- Fourth, timely escalation: significant events reach senior management quickly.
- The usual failure is under reporting, and it has two typical causes.
- One is a blame culture, where staff hide errors because they fear punishment.
- The other is the one we just met: tests applied to fragments of a larger event, so that each piece slips under the threshold.

6 How an RCSA works



IN THE LESSON

- Now from losses that happened to risks that might.
- A risk and control self assessment, or R C S A, asks the business itself to assess its operational risks.
- Read the boxes from the top row, assess.
- First, the business lists its risks and the controls meant to manage them.
- Second, it rates the inherent risk, meaning the risk before any controls, as if they did not exist.
- Third, it judges how effective the controls are.
- In the bottom row, conclude: that judgement leads to the residual risk, the risk that remains after controls.
- Finally, the residual risk is compared with the risk appetite, the amount of risk the board is willing to accept.
- Above appetite, the business must act.

7 Rating scales anchored on numbers

Fernhill's illustrative five-point scales

Score	Likelihood	Impact
5	Over 12 times a year	Over $5m
4	1 to 12 times a year	$1m to $5m
3	Once in over 1 to 5 years	$250k to under $1m
2	Once in over 5 to 20 years	$50k to under $250k
1	Rarer than 1 in 20 years	Under $50k

IN THE LESSON

- Ratings use two scales, one for likelihood and one for impact.
- Vague words such as rare or likely make assessors rate inconsistently.
- So likelihood is best anchored on frequencies, and impact on money bands.
- Here are Fernhill's illustrative five point scales, reading from the top.
- A five means more than twelve times a year, or an impact over five million dollars.
- A four means one to twelve times a year, or one to five million.
- A three means under once a year, down to once in five years, or two hundred and fifty thousand dollars to just under one million.
- A two means rarer, down to once in twenty years, or fifty thousand to just under two hundred and fifty thousand.
- A one means rarer than once in twenty years, or under fifty thousand dollars.

8 Example 2: an RCSA for duplicate payments

Batch sent twice. Zones: red if both scores 3+; amber if either 4+; else green

Inherent: before controls	likelihood 4, impact 4	**red zone**
Controls	duplicate check, daily recon	**rated effective**
Residual: after controls	likelihood 2, impact 4	**amber zone**
Appetite for payments	green; amber needs a plan	**action plan**

IN THE LESSON

- The payments team assesses the risk that a whole payment batch is sent twice.
- Fernhill's zones are simple: red when both scores are three or more, amber when either is four or more, and green otherwise.
- Without controls, this could happen one to twelve times a year, a likelihood of four, and a full batch sent twice could cost one to five million dollars, an impact of four.
- That inherent rating sits in the red zone.
- Two controls exist: a system check that blocks duplicates, and a daily reconciliation.
- The team rates both as effective.
- Controls of this kind mainly cut the likelihood, so the team rates residual likelihood at two, about once a decade.
- Impact stays at four, because the batch is just as big.
- That sits in the amber zone.
- Fernhill's appetite for payments is green, and amber is accepted only with an action plan, so the team must write one.

9 Heatmaps, and their weakness

- Heatmap: likelihood across, impact up
- Colour zones show priority: green, amber, red
- Ratings are ordinal: an order, not an amount
- Multiplying ranks does not give an expected loss
- Equal scores can hide very different risks

IN THE LESSON

- The results are usually shown on a heatmap.
- It plots likelihood along the bottom and impact up the side, and colour zones show priority: green for low, amber for watch, and red for act now.
- A heatmap is quick to read, and boards like it.
- But its weakness is that the ratings are ordinal.
- Ordinal means they give an order, not an amount.
- A four is worse than a two, but it is not twice as bad in dollars.
- So multiplying a likelihood rank by an impact rank does not give an expected loss, even though many firms print such a score.
- And because of that, equal scores can hide very different risks.
- The next example shows how.

10 Example 3: the same score, different risks

Both risks score 5 as a product of ranks. Translate into money

Risk A: small card fraud	likelihood 5 × impact 1	**score 5**
Risk B: one rogue payment	likelihood 1 × impact 5	**score 5**
A in money	30 a year × $10,000	**$300,000 a year**
B in money	$20m ÷ 25 years	**$800,000 a year**
B in its bad year	one event	**$20m**

IN THE LESSON

- Fernhill has two risks with the same score.
- Risk A is small card fraud.
- Likelihood five times impact one gives a score of five.
- Risk B is one huge rogue payment.
- Likelihood one times impact five also gives five.
- Now translate into money.
- Risk A happens about thirty times a year at around ten thousand dollars each, so about three hundred thousand dollars a year.
- Risk B happens about once in twenty five years, but costs twenty million dollars when it does.
- Its yearly average is twenty million divided by twenty five, eight hundred thousand dollars.
- Worse, in the bad year Fernhill loses the full twenty million at once.
- Same score, very different risks.
- Use a heatmap to rank attention, not to measure loss.

11 Self-assessment leans optimistic

- Self-assessment tends toward optimism
- Second line challenges ratings with evidence
- Loss data: 3 duplicates in 2 years, same control gaps
- Indicators: reconciliations often signed off late
- Residual likelihood raised from 2 to 4: red

IN THE LESSON

- One more weakness.
- An R C S A is done by the business about itself, and self assessment tends toward optimism.
- So the second line, the independent risk function, should challenge the ratings against evidence.
- At Fernhill, the risk manager pulls the loss data and finds three duplicate payment events in two years, including the reload we started with.
- They were smaller than a whole batch, but each slipped past the same controls, and that control failure, not the event size, drives likelihood.
- One and a half failures a year is a likelihood of four, not two.
- The indicators tell the same story, because the daily reconciliations are often signed off late.
- So the control rated effective is not working as claimed.
- After the challenge, the residual likelihood is raised from two to four, and the risk moves back to red.

12 Three kinds of indicator

Classify by what is measured, not by subject

Indicator	Measures	Example
KRI: key risk	A driver of exposure	Turnover, failed trades
KPI: performance	Business objectives met	Throughput, cost
KCI: key control	Whether a control works	Recs on time
All three	Against thresholds	Green, amber, red

IN THE LESSON

- Indicators are numbers tracked regularly, between assessments.
- There are three kinds, and they differ by what they measure.
- A key risk indicator, or K R I, tracks a driver of exposure, such as staff turnover or failed trades.
- It is meant to warn before losses occur.
- A key performance indicator, or K P I, measures how well the business meets its objectives, for example throughput or cost per payment.
- A key control indicator, or K C I, measures whether a control is operating effectively, such as the share of reconciliations signed off on time.
- All three are watched against thresholds, often green, amber and red, and crossing a threshold triggers escalation.
- The headline rule is written at the top: classify by what is measured, not by subject.

13 Example 4: Fernhill's indicator dashboard

Failed trades: amber 30, red 50. Recs on time: amber below 98%, red below 95%

Failed trades today: 42	KRI: driver of exposure	**amber**
Recs on time: 93%	KCI: is the control working?	**red**
Cost per payment: $0.40	KPI: target $0.45	**green**
Red result	escalate	**risk committee**

IN THE LESSON

- Here is part of Fernhill's dashboard.
- For failed trades, amber starts at thirty a day and red at fifty.
- For reconciliations signed off on time, amber is below ninety eight per cent and red is below ninety five.
- Failed trades today are forty two.
- That is a K R I, a driver of exposure, and it sits between thirty and fifty, so it is amber.
- Reconciliations on time are ninety three per cent.
- That is a K C I, telling us whether the control works, and it is below ninety five, so it is red.
- Cost per payment is forty cents, against a target of forty five cents.
- That is a K P I, and it is green.
- The red result is escalated to the risk committee.

14 Two traps with indicators

- Trap 1: classifying by subject area
- IT: downtime KRI; patches on time KCI; tickets KPI
- Trap 2: a performance target alone can raise risk
- Payments per clerk raised from 400 to 500 a day
- Balance it: a KRI (error rate), a KCI (checks done)

IN THE LESSON

- Two traps.
- The first is to classify an indicator by its subject area instead of by what it measures.
- Take three metrics from the I T department.
- Hours of system downtime is a K R I, because downtime drives exposure.
- The share of security patches applied on time is a K C I, because it shows whether a control works.
- Tickets closed per hour is a K P I, because it measures efficiency.
- The second trap is to forget that a performance target alone can raise risk.
- Suppose Fernhill raises the target for payments per clerk from four hundred to five hundred a day.
- Rushed clerks make more errors and skip checks.
- So balance the target with a K R I, such as the error rate, and a K C I, such as the share of four eyes checks completed.

15 A factor-based model

Loss = b_0 + b_1 × Volume + b_2 × Turnover + b_3 × Incidents

Intercept, b_0	Base level of loss
Sensitivities, b_1 to b_3	Loss per unit of each driver
Drivers	Volume, turnover, incidents

IN THE LESSON

- Now to models.
- A factor based model explains operational loss, or capital, through measurable drivers, such as transaction volume, staff turnover, system incidents and control scores.
- It is typically a regression or a scoring formula.
- Here is a simple regression form.
- Loss equals a base level, b nought, plus b one times volume, plus b two times staff turnover, plus b three times system incidents.
- The b coefficients are the sensitivities, estimated from the bank's history: how many dollars of loss come with one more unit of each driver.
- The appeal is that management can see how a change in the drivers moves the estimate.
- If turnover rises, the model says how much the expected loss should rise.

16 Example 5: Fernhill's factor model

Loss ($m) = 1.0 + 0.4 × volume (m trades) + 0.25 × turnover % + 0.05 × incidents

Base plus volume: 10m trades	1.0 + 0.4 × 10	**$5.0m**
Turnover: 12%	0.25 × 12	**$3.0m**
Incidents: 40	0.05 × 40	**$2.0m**
Estimate for the year	5 + 3 + 2	**$10.0m**
Turnover rises to 20%	+0.25 × 8 = +2	**$12.0m**

Limits: needs reliable factor data; assumes stable relationships

IN THE LESSON

- Losses are in millions of dollars.
- The base is one million.
- Volume carries point four per million trades, turnover point two five per percentage point, and incidents point zero five each.
- This year Fernhill processes ten million trades, so base plus volume is one plus four, five million.
- Staff turnover is twelve per cent, adding point two five times twelve, three million.
- There are forty system incidents, adding two million.
- The estimate for the year is ten million dollars.
- Now the what if.
- If turnover rises to twenty per cent, eight points more, the estimate rises by two million to twelve million.
- The limits are in the note: the model needs reliable data on every factor, and it assumes the relationships stay stable.

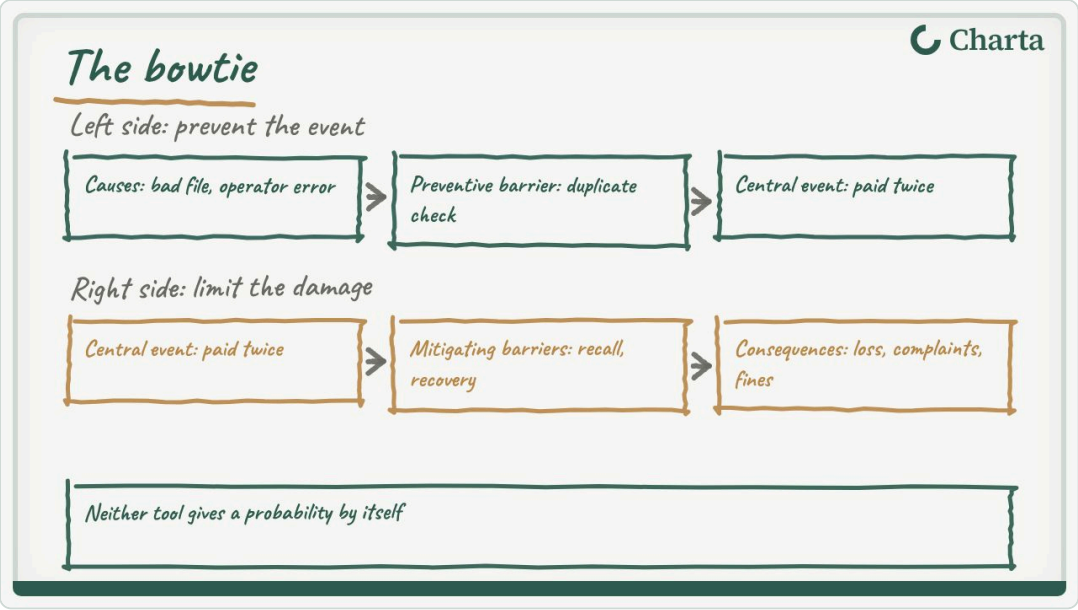
17 The Swiss cheese model

- Each defence is a slice of cheese with holes
- Holes: latent conditions and active failures
- A loss occurs when the holes line up
- Fernhill: check switched off, file reloaded, recon late
- Review every layer; it gives no probability

IN THE LESSON

- Two qualitative tools support the analysis.
- The first is the Swiss cheese model.
- Picture each defence as a slice of cheese.
- No slice is perfect, and each has holes.
- The holes are of two kinds.
- Latent conditions are weaknesses that sit quietly in the system, such as a badly designed process.
- Active failures are unsafe acts at the front line, such as a mistaken keystroke.
- A loss occurs only when the holes in every slice line up.
- At Fernhill, the duplicate check had been switched off during an earlier system upgrade, a latent condition.
- The operator reloaded the file, an active failure.
- And the reconciliation was signed off late.
- So the review must look at every layer.
- The model gives no probability by itself.

18 The bowtie



IN THE LESSON

- The second tool is the bowtie.
- A central event sits in the middle.
- Read the top row, the left side of the bow.
- The causes, such as a bad file or an operator error, lead toward the central event, but preventive barriers stand in the way, such as the duplicate check.
- If they fail, the central event happens: payments are sent twice.
- Now the middle row, the right side.
- From the same central event, mitigating barriers limit the damage, such as recalling the payments and pursuing recovery.
- Behind them sit the consequences: a financial loss, customer complaints, and possibly fines.
- And remember the line at the bottom: neither tool gives a probability by itself.

19 Fault trees: the two gates

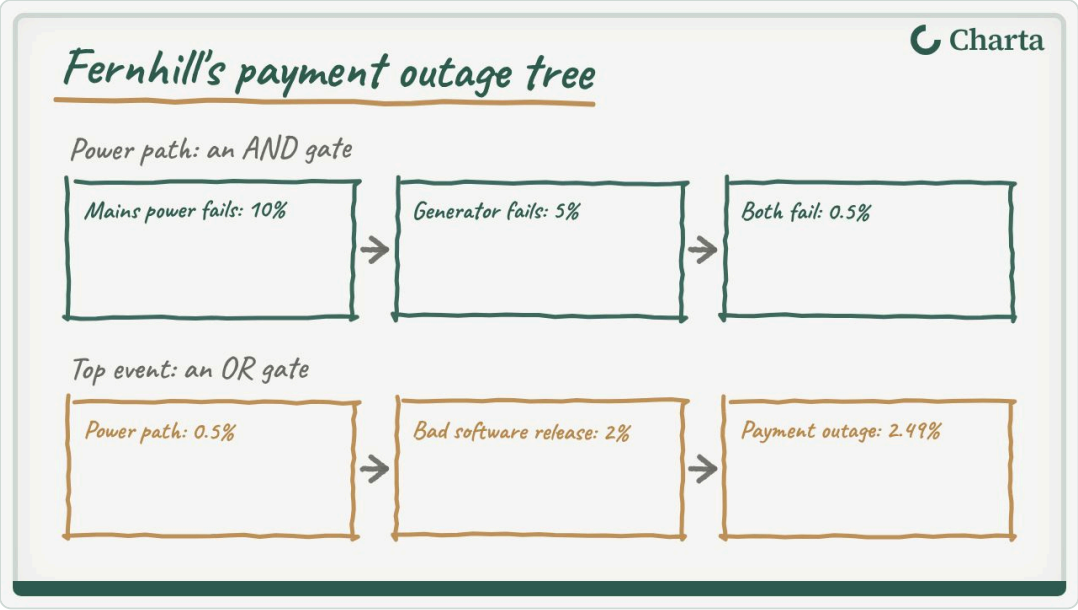
AND: $P = p_1 \times p_2$
OR: $Q = 1 - (1 - p_1) \times (1 - p_2)$

AND gate, P	All inputs must occur
OR gate, Q	Any one input is enough
Inputs, p_1 and p_2	Basic failure probabilities
Independence	Multiply only if independent

IN THE LESSON

- When we do want a probability, we use a fault tree.
- It starts from an undesired top event, such as a payment outage, and works down through logic gates to basic failures.
- There are two gates.
- At an AND gate, all inputs must occur for the output to occur.
- For independent events, the output probability, P, is p one times p two.
- At an OR gate, any one input is enough.
- For independent events, the output probability, Q, is one minus the product of one minus p one and one minus p two.
- The OR result is less than the sum, p one plus p two, because the events can overlap.
- And multiplication is valid only if the events are independent.

20 Fernhill's payment outage tree



IN THE LESSON

- Here is the fault tree for Fernhill's top event, a day long outage of its payment service.
- Work from the bottom up, applying each gate in turn.
- Start with the top row, the power path.
- Each year, mains power to the data centre fails with probability ten per cent, and the backup generator fails with probability five per cent.
- The power path fails only if both fail, so it is an AND gate, and both fail with probability half of one per cent.
- Now the bottom row.
- The outage happens if the power path fails, or if a bad software release brings the system down, which has a probability of two per cent.
- That is an OR gate, and the payment outage has a probability of two point four nine per cent.

21 Example 6: outage probability and expected loss

Independent: mains 10%, generator 5%, release 2%. Expected severity given outage $4m

AND: power path	0.10 × 0.05	**0.5%**
OR: payment outage	1 − 0.995 × 0.98	**2.49%**
Wrong: add them	0.5% + 2%	**2.50%**
Expected loss	2.49% × $4m	**$99,600**

IN THE LESSON

- The AND gate: point one zero times point zero five, half of one per cent.
- Then the OR gate.
- The chance that the power path does not fail is point nine nine five, and the chance that the release does not fail is point nine eight.
- Their product is point nine seven five one, so the outage probability is one minus that, two point four nine per cent.
- Adding the two would give two point five per cent, slightly too high, because it double counts the years when both happen.
- Finally, multiply by the expected severity given an outage, four million dollars.
- Expected loss is two point four nine per cent times four million, ninety nine thousand six hundred dollars a year.
- If each outage costs a fixed amount, that amount is the severity.

22 Example 7: when failures share a cause

Shared cause: Pjoint = P(A) × P(B given A). P(generator fails given mains fails) = 30%

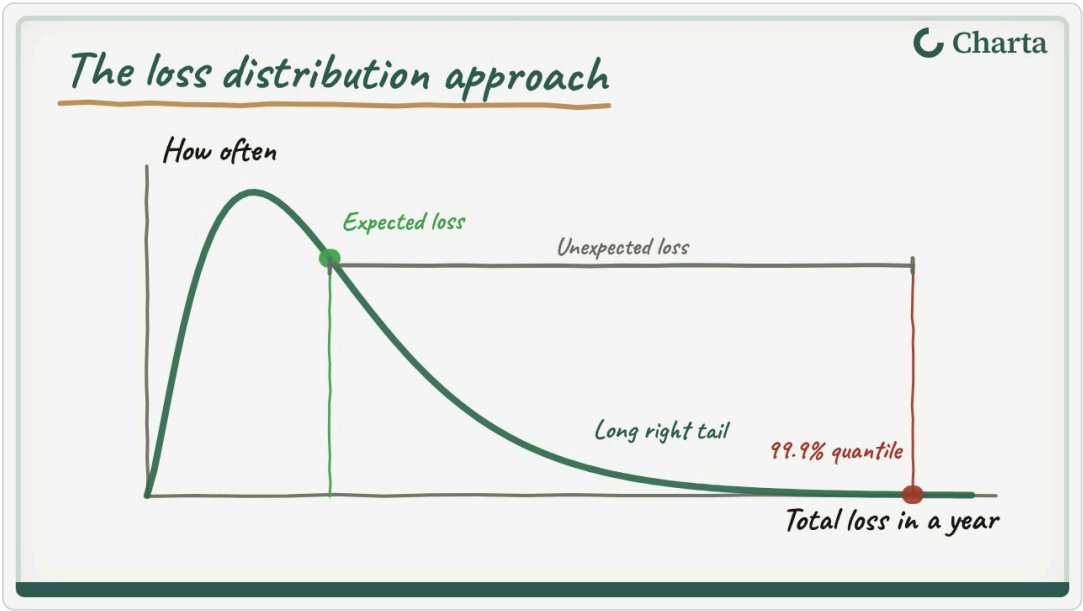
Joint, shared cause	0.10 × 0.30	**3.0%**
Independent figure	0.10 × 0.05	**0.5%**
Payment outage	1 − 0.97 × 0.98	**4.94%**
Expected loss	4.94% × $4m	**$197,600**

Typical errors: adding, the wrong gate, ignoring dependence

IN THE LESSON

- Mains power and the generator sit in the same basement, so a flood can stop both.
- The joint probability then uses the conditional probability: the chance of the first event, times the chance of the second given the first.
- The chance that the generator fails, given that mains power has failed, is thirty per cent.
- So the joint probability is point one zero times point three zero, three per cent, against half a per cent under independence, six times higher.
- Redo the OR gate: one minus point nine seven times point nine eight gives four point nine four per cent.
- Expected loss nearly doubles, to one hundred and ninety seven thousand six hundred dollars.
- The typical errors are in the note: adding probabilities, using the wrong gate, and ignoring dependence.

23 The loss distribution approach



IN THE LESSON

- Now economic capital, held to survive a very bad year.
- The first method is the loss distribution approach.
- It combines a frequency distribution, often Poisson, for how many losses occur in a year, with a severity distribution, often lognormal, for how large each loss is.
- A simulation draws many years and adds up each year's losses.
- This chart shows the result.
- Along the bottom is the total loss in a year, and up the side is how often each total occurs.
- The average, marked in green, is the expected loss.
- The curve has a long right tail of rare, severe years.
- Capital is read from a high quantile, such as the ninety nine point nine per cent level, marked in red.
- The gap between the two is the unexpected loss.

24 Example 8: reading capital from the simulation

Poisson mean 25 a year; mean severity $80,000; 100,000 simulated years

Expected annual loss	25 × $80,000	**$2.0m**
99.9% quantile	100th worst of 100,000	**$15.0m**
Capital: unexpected loss	15.0 − 2.0	**$13.0m**
The tail	far beyond the sample	**unstable**

Suits areas with plentiful data, such as payment errors

IN THE LESSON

- Fernhill applies the method to its retail payments losses.
- Frequency is Poisson with a mean of twenty five losses a year, and the mean severity is eighty thousand dollars.
- The expected annual loss is twenty five times eighty thousand, two million dollars.
- Fernhill simulates one hundred thousand years.
- The ninety nine point nine per cent quantile is the hundredth worst of them, and suppose it is fifteen million dollars.
- The part beyond expected loss is fifteen minus two, thirteen million, the unexpected loss.
- Economic capital usually covers just that, because expected loss is already covered by provisions or pricing.
- If not, capital is the full fifteen million.
- Now the weakness.
- With perhaps ten years of data, the quantile describes a one in a thousand year loss.
- The tail is extrapolated far beyond the sample, so it is unstable.
- The method suits areas with plentiful data, such as payment errors.

25 Three ways to size capital

Match the method to the data situation

Method	Where it fits	Shortcoming
Loss distribution	Plentiful loss data	Unstable, extrapolated tail
Scenario-based	Severe, data-poor events	Subjective, anchoring
Scorecard	Changing businesses	Subjective, weakly calibrated

IN THE LESSON

- There are other methods, and the exam expects you to match each method with its data situation.
- The loss distribution approach fits areas with plentiful loss data, and its shortcoming is the extrapolated, unstable tail.
- Scenario based methods use expert judgement about severe events.
- They fit data poor areas.
- Their shortcoming is that they are subjective and open to anchoring bias, where experts stick too closely to the first number they hear.
- Scorecards rate risk drivers and the quality of controls.
- They fit businesses that are changing, and they are forward looking, because they move as soon as the drivers or controls change.
- But they are subjective and weakly calibrated to actual losses.
- So the first question is always: what data do we have?

26 Combining two units

Capital = √(a^2 + b^2 + 2ρab)
Simple sum a + b assumes ρ = 1

Stand-alone capital, a	First unit's own capital
Stand-alone capital, b	Second unit's own capital
Correlation, ρ	Correlation between losses

IN THE LESSON

- Once each unit has a capital figure, the firm must combine them.
- Let ay and b be the stand alone capital of two units, meaning the capital each would need on its own, and let rho be the correlation between their losses.
- Combined capital follows the square root rule.
- Capital equals the square root of ay squared, plus b squared, plus two times rho times ay times b.
- Simply adding the two figures, ay plus b, assumes perfect correlation, a rho of one, where the bad years of both units always arrive together.
- Whenever rho is below one, the square root result is lower than the sum.
- The difference is the diversification benefit.

27 Example 9: Fernhill's combined capital

Payments a = $30m, trading operations b = $40m, ρ = 0.25

Combined, ρ = 0.25	√(900 + 1,600 + 600)	**$55.7m**
Simple sum	30 + 40, as if ρ = 1	**$70.0m**
Diversification benefit	70.0 − 55.7	**$14.3m**
If ρ = 0	√(900 + 1,600)	**$50.0m**

IN THE LESSON

- Fernhill's payments unit has stand alone capital of thirty million dollars, and its trading operations unit has forty million.
- The correlation between their losses is point two five.
- Square each: nine hundred, and one thousand six hundred.
- The cross term is two times point two five times thirty times forty, which is six hundred.
- The total under the root is three thousand one hundred, and its square root is about fifty five point seven million dollars.
- The simple sum is seventy million, which is what you would get with a correlation of one.
- So the diversification benefit is about fourteen point three million.
- And if the losses were uncorrelated, the root of two thousand five hundred gives fifty million.
- Lower correlation means lower combined capital.

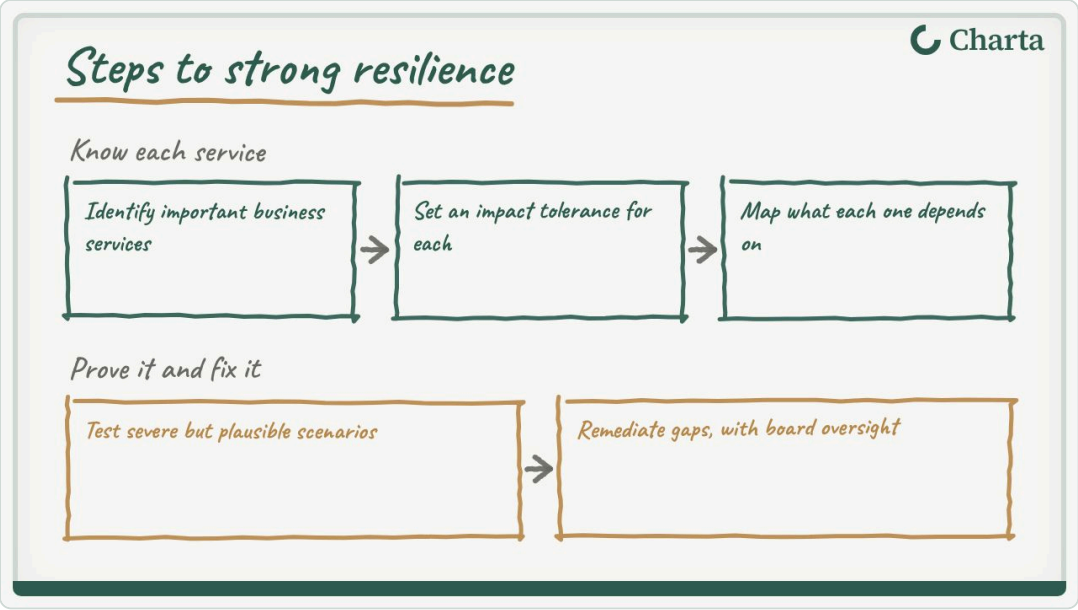
28 Operational resilience

- Keep delivering critical services through disruption
- Prevent, adapt, respond and recover
- Impact tolerance: the maximum tolerable disruption
- It assumes the disruption happens
- Risk appetite concerns likelihood; tolerance does not

IN THE LESSON

- The last outcome turns from measuring losses to staying in business.
- Operational resilience is the ability to keep delivering critical services through disruption.
- It works by preventing disruption where possible, adapting to it, responding when it strikes, and recovering afterwards.
- The key idea is the impact tolerance.
- It is the maximum disruption a service can tolerate.
- Notice the assumption: an impact tolerance assumes the disruption happens.
- It does not ask how likely the outage is.
- That is the difference from risk appetite, which concerns likelihood, meaning how much risk of loss the firm will accept.
- A firm can have a very low chance of an outage and still breach its tolerance when one occurs.

29 Steps to strong resilience



IN THE LESSON

- The steps follow a clear order, starting with the top row, know each service.
- First, identify the important business services, the services whose failure would seriously harm customers or markets, such as Fernhill's retail payments.
- Second, set an impact tolerance for each.
- Third, map what each service depends on: the people, processes, technology, facilities, information and third parties behind it.
- Mapping reveals single points of failure, such as one supplier that every service relies on.
- In the bottom row, prove it and fix it.
- Test the services against severe but plausible scenarios, such as the loss of a data centre or a cyber attack.
- Finally, remediate the gaps that the tests reveal, with board oversight, so that fixes are funded and tracked to the end.

30 Example 10: measuring disruption end to end

Retail payments: impact tolerance 4 hours.
Scenario: primary data centre lost

Detect the outage	alerts reach on-call staff	**0h 30m**
Decide to fail over	incident call	**0h 45m**
Fail over to backup site	technical recovery	**2h 00m**
Checks until usable	reconcile data, reopen	**1h 00m**
End to end	tolerance 4h 00m	**4h 15m: breach**

Failover alone took 2 hours and looked fine

IN THE LESSON

- Fernhill sets an impact tolerance of four hours for retail payments, and tests the loss of its primary data centre.
- Disruption is measured end to end, from detection to a usable service.
- Detection takes thirty minutes, until alerts reach the staff on call.
- Deciding to fail over takes forty five minutes on the incident call.
- The technical failover to the backup site takes two hours.
- Then the checks until the service is usable, reconciling data and reopening to customers, take one hour.
- End to end, that is four hours and fifteen minutes, against a tolerance of four hours: a breach.
- Yet the technology team reported only the failover, which took two hours and looked comfortably inside.
- The breach is a gap to remediate.

31 Ways to test resilience

Passing only planned scenarios gives false comfort

Test	What it proves
Tabletop exercise	Decisions and communication
Simulation	People and processes, in real time
Live failover	Technical recovery really works
Cyber exercise	Response and recovery from attack

IN THE LESSON

- Tests range from light to heavy.
- A tabletop exercise gathers decision makers around a table to talk through a scenario.
- It tests decisions and communication: who decides, and who tells customers and the supervisor.
- A simulation goes further, playing the scenario out in real time, so that people and processes are tried under pressure.
- A live failover actually switches the service to the backup site, proving that technical recovery really works.
- A cyber exercise tests the response to an attack and the recovery from it, end to end.
- Then a warning, written at the top.
- Passing tests built only on planned scenarios gives false comfort.

32 Mistakes to avoid

Five traps from this lesson

Mistake	Instead
Threshold on fragments	Group by root cause first
Multiplying heatmap ranks	Ranks are ordinal, not money
Indicators by subject	Classify by what is measured
Adding OR-gate probabilities	1 − product of (1 − p)
Adding stand-alone capital	Square-root rule with ρ

IN THE LESSON

- Before the quiz, here are five traps from this lesson.
- First, testing the recording threshold on fragments of an event.
- Instead, group losses by root cause first, then apply the threshold.
- Second, multiplying heatmap ranks and treating the score as money.
- Ranks are ordinal, an order and not an amount.
- Third, classifying indicators by subject area.
- Classify them by what is measured: exposure, performance or control.
- Fourth, adding probabilities at an OR gate.
- Use one minus the product of one minus each probability, and check for dependence before multiplying anything.
- Fifth, adding stand alone capital figures, which assumes perfect correlation.
- Use the square root rule with the correlation instead.

33 Try three quick questions

1. Independent failures of 20% and 10% feed an OR gate. Top-event probability?
   Answer: 1 − 0.8 × 0.9 = 0.28, so 28% (not 30%).
2. Stand-alone capital $50m and $50m, ρ = 0.5. Combined capital?
   Answer: √(2,500 + 2,500 + 2,500) = √7,500 ≈ $86.6m.
3. Average days to approve a mortgage: KRI, KPI or KCI?
   Answer: KPI: it measures a business objective, speed.

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- Two independent basic failures, with probabilities of twenty per cent and ten per cent, feed an OR gate.
- What is the probability of the top event?
- One minus point eight times point nine is point two eight, so twenty eight per cent, not the thirty per cent that adding them would give.
- Question two.
- Two units each have stand alone capital of fifty million dollars, and their correlation is point five.
- What is the combined capital?
- Two thousand five hundred, three times over, is seven thousand five hundred, and its square root is about eighty six point six million dollars.
- Question three.
- Average days to approve a mortgage: is it a K R I, a K P I or a K C I?
- It is a K P I, because it measures a business objective, speed.