CHARTA LEARNING · LESSON NOTES

Integrated Risk Management

FRM Part II

1 A story: three silos and one bad year

- Lending, trading, payments: each manages its own risk
- A fine appetite statement, but bonuses on revenue alone
- Staff who raise concerns are told not to slow things
- One bad year hits all three at once
- Lesson: manage risk across the whole enterprise

IN THE LESSON

- Imagine Kestrel Bank a few years ago.
- Its lending, trading and payments businesses each managed their own risks, and nobody added them up.
- The board had approved a risk appetite statement, a fine document that said the bank would avoid large losses.
- But traders were paid bonuses on revenue alone.
- Junior staff who raised concerns were told not to slow the business down.
- Then one bad year hit all three businesses at once.
- The losses were far bigger than any single team had imagined, because each had looked only at its own piece.
- The lesson Kestrel learned is the subject of this module: risk must be managed across the whole enterprise, as one.

2 An ERM framework: three parts that work together

- ERM: all risks, one framework, one view of appetite
- Governance: who is accountable
- Appetite: how much risk, of which types, for the strategy
- Culture: shared values that decide if rules are followed
- Steering, speed limit, the driver's habits: need all three

IN THE LESSON

- Enterprise risk management is one framework that covers every risk the firm takes, managed against one view of how much risk is acceptable.
- It draws on three elements that only work together.
- Risk governance sets who is accountable for what.
- Risk appetite sets how much risk, and which types, the firm is willing to accept to pursue its strategy.
- Risk culture is the shared values and behaviour that decide whether the rules are actually followed.
- Think of a car.
- Governance is the steering and the people allowed to drive.
- Appetite is the speed limit.
- Culture is whether the driver respects the limit when nobody is watching.
- You need all three, and removing any one breaks the framework.

3 Risk governance: who is accountable

Approve, implement, then separate risk-taking, oversight and assurance

Who	Role in the ERM framework
Board	Approves strategy and risk appetite
Senior management	Implements them: policies, limits
1st line: business units	Take and own the risk
2nd line: risk, compliance	Independent oversight, challenge
3rd line: internal audit	Independent assurance to the board

IN THE LESSON

- Start with governance.
- At the top, the board of directors approves the strategy and the risk appetite.
- Senior management implements them: it turns the appetite into policies and limits, and runs the bank day to day.
- Below them, the three lines of defence separate the jobs.
- The first line is the business units, such as Kestrel's lenders and traders.
- They take the risk and own it.
- The second line is independent oversight by risk management and compliance, which set and monitor limits and challenge the first line.
- The third line is internal audit, which gives the board independent assurance that the whole system works.

4 Risk appetite, risk capacity and limits

- Appetite: aggregate level and types of risk accepted
- Capacity: max risk before capital or liquidity minimums
- Appetite sits below capacity: a safety margin
- Appetite is turned into business-line limits
- Chain: capacity, then appetite, then limits

IN THE LESSON

- Next, risk appetite.
- It is the aggregate level and the types of risk the firm is willing to accept in order to pursue its strategy.
- It must sit below risk capacity.
- Risk capacity is the maximum risk the firm could bear before it breaches its capital or liquidity minimums.
- Because appetite sits below capacity, there is a safety margin, so that a bad surprise does not push the bank straight to the edge.
- Appetite is no use as a sentence in a report.
- It must be turned into limits for each business line, so that a trader or a lender knows exactly how much risk they may take.
- The chain runs from capacity, to appetite, to limits.

5 Example 1: from capacity to limits

Kestrel: CET1 capital $7.5bn, risk-weighted assets $100bn

CET1 capital today	best-quality capital	**$7.5bn**
Minimum CET1	4.5% × $100bn	**$4.5bn**
Risk capacity	7.5 − 4.5	**$3.0bn**
Risk appetite: severe year	50% of capacity	**$1.5bn**
Business-line limits	0.6 + 0.6 + 0.3	**$1.5bn**

Retail $0.6bn, corporate $0.6bn, markets $0.3bn

IN THE LESSON

- Let us put numbers on the chain.
- Kestrel holds common equity tier one capital, C E T one for short, of seven point five billion dollars.
- That is its best quality capital, mostly shareholders' equity.
- Its risk-weighted assets, a measure of how risky its assets are that we meet properly later, are one hundred billion dollars.
- The minimum C E T one is four point five per cent of that, so four point five billion.
- Kestrel could therefore lose three billion dollars before breaching the minimum.
- That is its risk capacity.
- The board sets the appetite at half of capacity: in a severe year, losses should not exceed one point five billion.
- Senior management splits that into business line limits: point six billion for retail lending, point six for corporate lending and point three for markets, which add back to one point five.

6 Risk culture: will the rules be followed?

- Shared values and behaviour: are rules followed?
- Tone from the top: what leaders say and do
- Incentives: what people are paid for
- Escalation: bad news travels up, safely
- Challenge: risk can question a profitable business

IN THE LESSON

- The third element is risk culture, the shared values and behaviour across the firm that decide whether the rules are followed when it matters.
- It has four parts you should know.
- Tone from the top: what the board and senior managers say and, more importantly, what they do.
- Incentives: what people are paid for.
- If bonuses reward revenue alone, people will find ways round the limits.
- Escalation: whether bad news travels up quickly, and whether staff feel safe raising a concern.
- Challenge: whether the risk function, and anyone else, can question a profitable business without being overruled.
- Rules describe what should happen.
- Culture decides what does happen.

7 The exam point: appetite fails without culture

- Trap: a strong appetite statement fails without culture
- Markets limit $0.3bn; desk's severe-year loss $0.36bn
- That is 20% over the limit, yet nobody escalates
- Fix: risk-adjusted pay, prompt escalation
- Fix: a second line with real authority to say no

IN THE LESSON

- Here is the point examiners like most.
- A strong risk appetite statement fails without the right culture and incentives.
- Go back to Kestrel.
- Its markets limit is point three billion dollars.
- A trading desk paid on revenue alone builds a position that would lose point three six billion in a severe year.
- That is twenty per cent over the limit, yet nobody escalates.
- In practice, the desk plays down the size, the junior risk analyst who notices stays quiet, and the head of trading, a big earner, is never challenged.
- The fix is not a better document.
- It is pay tied to risk-adjusted results, prompt escalation of every breach, and a second line with real authority to say no.

8 Two kinds of capital

- Regulatory capital: supervisors' minimum, prescribed rules
- Economic capital: own estimate for unexpected losses
- Confidence level matches the target credit rating
- E.g. aiming for AA: about 99.97% over one year
- Regulatory: a floor. Economic: limits, pricing, performance

IN THE LESSON

- Now to capital, the cushion that absorbs losses.
- Regulatory capital is the minimum set by supervisors, using prescribed rules.
- Economic capital is the bank's own estimate of the capital it needs to absorb unexpected losses, at a confidence level that matches its target credit rating.
- For example, a bank aiming for a double A rating might choose a one year confidence level of about ninety nine point nine seven per cent, because a double A firm has roughly a three in ten thousand chance of defaulting within a year.
- Keep the purposes apart.
- Regulatory capital is a floor the bank must meet.
- Economic capital is a management tool, used to set limits, price products and measure performance.

9 Basel III minimum capital ratios

Each minimum is capital ÷ risk-weighted assets (RWA)

Requirement	Share of RWA
CET1 minimum	4.5%
Tier 1 minimum (CET1 + AT1)	6.0%
Total capital minimum (+ Tier 2)	8.0%
Conservation buffer, in CET1	+2.5%
Countercyclical, systemic	Vary by jurisdiction and bank

IN THE LESSON

- Under Basel three, minimum capital is a percentage of risk-weighted assets, R W A for short.
- Risk-weighted assets scale each exposure by its riskiness, and they include charges for market and operational risk.
- The baseline minima are four point five per cent for C E T one, and six per cent for tier one capital, which adds additional tier one instruments, A T one, to C E T one.
- Total capital, which adds tier two, has a minimum of eight per cent.
- On top sits a capital conservation buffer of two point five per cent, which must be met with C E T one.
- Then come any countercyclical buffer, set by each jurisdiction, and surcharges for systemically important banks, which depend on each bank's size and importance.
- So the requirement actually applied varies by jurisdiction and by bank.

10 Capital ratios and the 12.5 multiplier

Capital ratio = Capital ÷ RWA
Op risk RWA = 12.5 × op risk capital

Capital	CET1, Tier 1 or total capital
Risk-weighted assets	Credit, market and op risk
Op risk capital charge	A dollar charge; 12.5 = 1 ÷ 8%

IN THE LESSON

- Every capital ratio has the same shape: the capital, divided by risk-weighted assets.
- Credit risk produces risk-weighted assets directly.
- But operational risk is first measured as a capital charge, a dollar amount.
- To fit it into the same ratio, the charge is converted into risk-weighted assets by multiplying by twelve point five.
- This holds for operational risk under any framework.
- Why twelve point five?
- Because it is one divided by eight per cent.
- If a bank needs one dollar of capital for operational risk, multiplying by twelve point five gives twelve and a half dollars of risk-weighted assets, and eight per cent of twelve and a half is exactly one dollar again.

11 Example 2: Kestrel's capital ratios

$bn: credit RWA 75, market RWA 10, op capital 1.2; CET1 7.5, AT1 1.5, Tier 2 2.0

Op risk RWA	12.5 × 1.2	**$15bn**
Total RWA	75 + 10 + 15	**$100bn**
CET1 ratio	7.5 ÷ 100	**7.5%**
Tier 1 ratio	(7.5 + 1.5) ÷ 100	**9.0%**
Total capital ratio	(9.0 + 2.0) ÷ 100	**11.0%**

With the 2.5% buffer: need 7.0%, 8.5%, 10.5%. All met

IN THE LESSON

- Credit risk-weighted assets are seventy five billion dollars, and market risk-weighted assets ten billion.
- The operational risk capital charge is one point two billion, so its risk-weighted assets are twelve point five times one point two, which is fifteen billion.
- Total risk-weighted assets are seventy five plus ten plus fifteen, one hundred billion.
- C E T one of seven point five billion gives a ratio of seven point five per cent.
- Add one point five billion of additional tier one, and tier one is nine billion, a ratio of nine per cent.
- Add two billion of tier two, and total capital is eleven billion, eleven per cent.
- With the conservation buffer, Kestrel needs seven, eight point five and ten point five per cent.
- All three are met.

12 IRB credit capital: unexpected loss only

K = LGD × (WCDR − PD)
Capital = K × EAD

Loss given default	Share of exposure lost
Worst-case default rate	99.9% one-year default rate
Probability of default	Expected default probability
Capital per unit of exposure	Charge for unexpected loss only

IN THE LESSON

- Now credit risk under the internal ratings based approach, I R B.
- The capital charge per unit of exposure, written K, equals the loss given default times the worst case default rate minus the default probability.
- Take the parts in turn.
- Loss given default, L G D, is the share of the exposure lost.
- The worst case default rate, W C D R, is the one year default rate exceeded only with a probability of one in a thousand, that is, at ninety nine point nine per cent confidence.
- P D is the expected default probability.
- Subtracting it is the key step: the expected part of the loss is covered by provisions and pricing, so K charges capital for the unexpected part only.
- Capital for a loan is then K times its exposure at default.

13 Example 3: capital for one corporate loan

Loan $50m to Harrow Shipping. PD 1%, LGD 45%, WCDR given as 12%

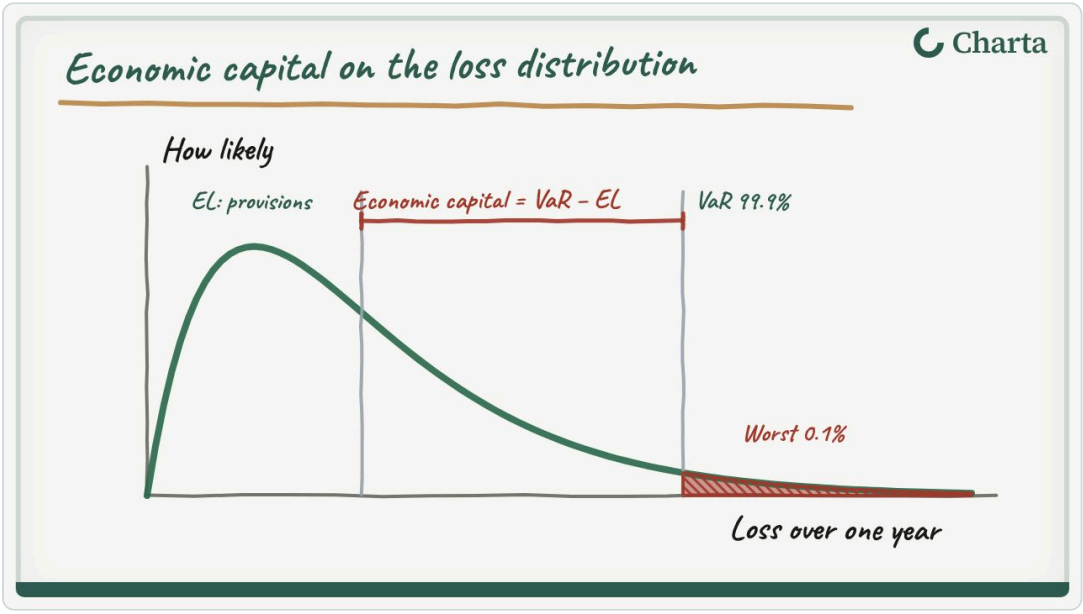
Capital per $1, K	45% × (12% − 1%)	**4.95%**
Capital for the loan	4.95% × $50m	**$2.475m**
Risk-weighted assets	12.5 × 2.475	**$30.9m**
Expected loss	1% × 45% × $50m	**$225,000**
Wrong: 45% × 12%	charges for EL too	**5.40%**

Simplified: WCDR stated, no maturity adjustment

IN THE LESSON

- Kestrel lends fifty million dollars to Harrow Shipping.
- The default probability is one per cent, the loss given default forty five per cent, and the worst case default rate, given here, is twelve per cent.
- K is forty five per cent times twelve minus one per cent, which is four point nine five per cent.
- Capital for the loan is four point nine five per cent of fifty million, two point four seven five million dollars.
- Multiply by twelve point five and the risk-weighted assets are about thirty point nine million.
- The expected loss is one per cent times forty five per cent times fifty million, two hundred and twenty five thousand dollars, which provisions cover.
- The common error is to multiply forty five per cent by twelve per cent, giving five point four per cent, which charges capital for the expected loss too.
- This is the simplified form: the worst case rate is stated and there is no maturity adjustment.

14 Economic capital on the loss distribution



IN THE LESSON

- Along the bottom is the loss over one year.
- Up the side is how likely each loss is.
- The curve is a typical loss distribution, drawn as a sketch and not to scale: small losses are common, and a long tail to the right holds rare, very large losses.
- The first grey line marks the expected loss, the average, which is covered by provisions and pricing.
- The second marks the value at risk, or V a R, at ninety nine point nine per cent: the loss exceeded only one year in a thousand.
- The bracket between them is economic capital, V a R minus expected loss, which is the unexpected part.
- The shaded tail, the worst point one per cent of years, lies beyond even that cushion.

15 Example 4: economic capital by risk type

Kestrel's own model, one year, 99.9%. Amounts in $bn

Credit: VaR − EL	3.0 − 0.4	**$2.6bn**
Market: VaR − EL	0.9 − 0 (EL assumed nil)	**$0.9bn**
Operational: VaR − EL	1.3 − 0.2	**$1.1bn**
Sum, no diversification	2.6 + 0.9 + 1.1	**$4.6bn**
Wrong: VaR not net of EL	3.0 + 0.9 + 1.3	**$5.2bn**

EL is covered by provisions and pricing, not by capital

IN THE LESSON

- Kestrel's own model estimates the ninety nine point nine per cent one year V a R for each risk type, then subtracts the expected loss.
- For credit risk, V a R is three billion dollars and expected loss point four billion, so economic capital is two point six billion.
- For market risk, V a R is point nine billion and we assume the expected loss is nil, so economic capital is point nine billion.
- For operational risk, V a R is one point three billion and expected loss point two billion, so economic capital is one point one billion.
- Added up with no allowance for diversification, the total is four point six billion dollars.
- Forget to subtract expected loss, and you get five point two billion.
- Expected loss belongs to provisions and pricing, not to capital.

16 Example 5: diversification across risk types

Correlation between each pair of risk types assumed 0.5. In $bn

Add, Basel style	2.6 + 0.9 + 1.1	**$4.6bn**
Sum of squares	2.6^2 + 0.9^2 + 1.1^2	**8.78**
Cross terms, ρ = 0.5	2 × 0.5 × (2.34 + 2.86 + 0.99)	**6.19**
Diversified EC	√(8.78 + 6.19)	**$3.87bn**
Diversification benefit	4.6 − 3.87	**$0.73bn**

Benefit about 16%. An approximation; correlations can rise

IN THE LESSON

- The Basel rules simply add the charges for each risk type.
- An economic capital model can recognise diversification.
- Suppose Kestrel assumes a correlation of point five between each pair.
- Adding Basel style gives four point six billion dollars.
- Instead, combine the three like standard deviations.
- The sum of squares is two point six squared, plus point nine squared, plus one point one squared, which is eight point seven eight.
- The cross terms are two, times point five, times the sum of the three products, two point three four, two point eight six and point nine nine, giving six point one nine.
- The square root of the total, fourteen point nine seven, is about three point eight seven billion.
- The diversification benefit is point seven three billion, about sixteen per cent.
- This variance covariance shortcut assumes each capital figure scales like a standard deviation, so it is only an approximation.
- And correlations can rise in a crisis, so the benefit needs care.

17 Mistakes to avoid: capital

- Expected loss: provisions. Unexpected loss: capital
- IRB: K uses WCDR − PD, not WCDR alone
- Economic capital = VaR − EL, not VaR alone
- Basel adds risk types; EC models may diversify
- Op risk capital × 12.5 = RWA. Never divide

IN THE LESSON

- Here are the mistakes that cost marks.
- First and most important, do not mix up expected and unexpected loss.
- Expected loss is the average, covered by provisions and pricing.
- Unexpected loss is the deviation above that, covered by capital.
- Second, in the I R B formula, subtract the default probability from the worst case default rate before multiplying by loss given default.
- Third, economic capital is V a R minus expected loss, not V a R alone.
- Fourth, the Basel rules add the charges for each risk type, while an economic capital model may recognise diversification, so the two totals differ.
- Fifth, multiply an operational risk capital charge by twelve point five to turn it into risk-weighted assets.
- Never divide.

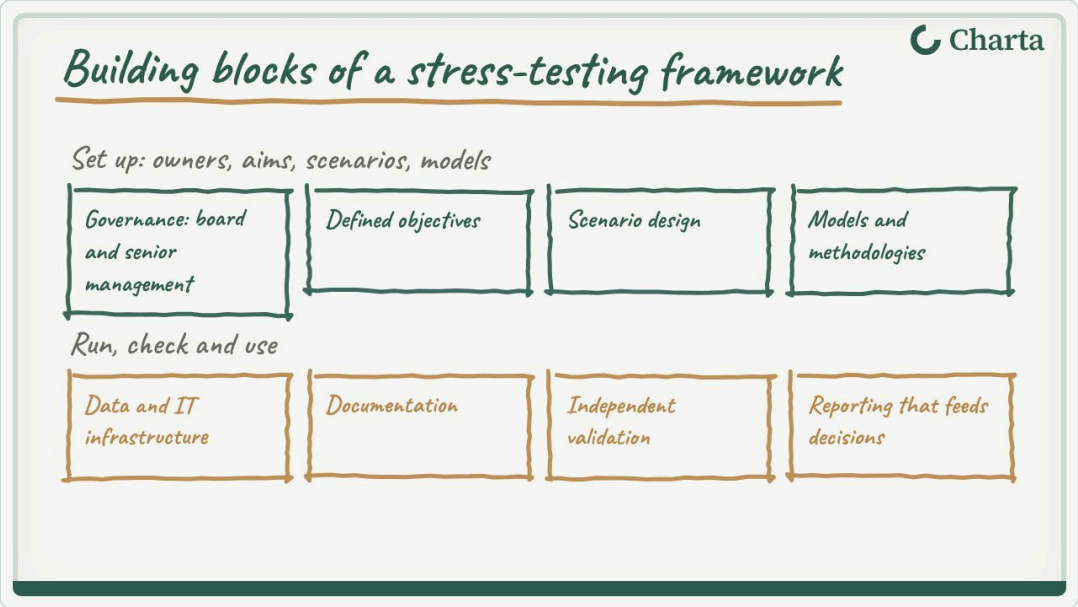
18 Why stress test?

- VaR: built from recent history and normal relationships
- Stress test: what happens under a specific severe event?
- Complements VaR; does not replace it
- Explores the tail; joins risks into one story
- Purpose: inform capital, liquidity, limits, strategy

IN THE LESSON

- Kestrel had relied on its V a R model.
- But V a R is built from recent history and normal relationships between markets, so it says little about what happens when the world changes.
- A stress test asks a different question: what would happen to us under a specific, severe event?
- It does not replace V a R.
- It complements it.
- It explores the part of the tail that history has not shown, and it joins risks together in one story.
- Its purpose is practical: to inform decisions about capital, liquidity, limits and strategy.
- So what does a robust stress testing framework need?

19 Building blocks of a stress-testing framework



IN THE LESSON

- A robust framework has eight building blocks.
- The top row sets the programme up: its owners, its aims, its scenarios and its models.
- Governance: the board and senior management own the programme, approve the scenarios and act on the results.
- Defined objectives: the bank states what each test is for, such as capital planning, liquidity or setting limits.
- Scenario design: the events to be tested.
- Models and methodologies: the tools that turn a scenario into losses, revenues and capital.
- The bottom row runs, checks and uses the tests.
- Data and I T infrastructure: the bank must gather exposures across the whole firm, quickly.
- Documentation: assumptions and choices are written down.
- Independent validation: people who did not build the models check them.
- And reporting that feeds decisions: results reach the people who can act on them.

20 Ways to design a scenario

Scenarios should be severe but plausible

Approach	What it does
Sensitivity analysis	Shocks one risk factor alone
Scenario analysis	Moves many factors in one story
Historical	Replays a past crisis
Hypothetical	An event that has not happened
Reverse stress test	Starts from failure, works back

IN THE LESSON

- Scenarios should be severe but plausible: bad enough to hurt, but not so fantastic that management can dismiss them.
- There are two ways to shock the bank.
- Sensitivity analysis shocks one risk factor alone, such as interest rates, holding everything else fixed.
- Scenario analysis moves many factors together, consistently, within a narrative, such as a recession in which unemployment rises, property prices fall and credit spreads widen.
- Scenarios themselves come in three kinds.
- A historical scenario replays a past crisis on today's portfolio.
- A hypothetical scenario describes an event that has not happened yet, but could.
- And a reverse stress test starts from an outcome, such as breaching capital minimums or failure, and works back to the events that could cause it.

21 Example 6: sensitivity versus scenario

Kestrel: CET1 $7.5bn, RWA $100bn. A one-year recession scenario

Sensitivity: rates +2 points	one factor only	**– $0.3bn**
Scenario: credit, market, op	1.4 + 0.5 + 0.1	**– $2.0bn**
CET1 after the scenario	7.5 − 2.0	**$5.5bn**
RWA after downgrades	assumed	**$105bn**
CET1 ratio after	5.5 ÷ 105	**5.2%**

Above the 4.5% minimum, below 7.0% with buffer: act now

IN THE LESSON

- A sensitivity test raises interest rates by two percentage points, with nothing else changing, and Kestrel's bond portfolio loses point three billion dollars.
- Now a recession scenario, where many factors move together.
- Credit losses are one point four billion, market losses point five billion and operational losses point one billion, two billion in all.
- C E T one falls from seven point five to five point five billion.
- At the same time borrowers are downgraded, so risk-weighted assets rise to an assumed one hundred and five billion.
- The C E T one ratio becomes five point five divided by one hundred and five, about five point two per cent.
- That is above the four point five per cent minimum, but below the seven per cent that includes the buffer.
- The decision follows: cut the dividend and slow property lending now.

22 Example 7: a reverse stress test

Start from failure: CET1 falls to the 4.5% minimum. Work back

CET1 capital now	Kestrel today	**$7.5bn**
Failure point: CET1 minimum	4.5% × $100bn	**$4.5bn**
Loss that gets there	7.5 − 4.5 = risk capacity	**$3.0bn**
Story 1: property book alone	3.0 ÷ 10 book	**30% loss**
Story 2: property and cyber	20% × 10 + 1.0	**$3.0bn**

Is either story plausible? If so, act before it happens

IN THE LESSON

- A reverse stress test runs the other way.
- Kestrel starts from the outcome it must avoid: C E T one falling to the four point five per cent minimum.
- With risk-weighted assets held at one hundred billion, that point is four point five billion of capital.
- From seven point five billion today, a loss of three billion gets there.
- Notice that this is exactly the risk capacity we met earlier.
- Now work back to stories that could cause it.
- Story one: the commercial property book of ten billion dollars alone loses thirty per cent.
- Story two: property loses twenty per cent, which is two billion, and in the same year a cyber attack costs one billion.
- Is either story plausible?
- If so, the bank should act before it happens.

23 Recommended practices

- Cover the whole firm and all material risks
- Capture interactions and feedback effects
- Be realistic about management actions under stress
- Tie results to capital, liquidity, limits, strategy
- Judge it by whether decisions change

IN THE LESSON

- What separates a good stress test from a box-ticking one?
- First, it covers the whole firm and all material risks.
- Second, it captures interactions and feedback effects.
- In a crisis, losses force banks to sell assets, which pushes prices down, which causes more losses, while funding dries up at the same time.
- Third, it is realistic about management actions.
- In stressed markets, everyone tries to sell the same assets or raise capital at once, so a plan to do so may not work.
- Fourth, results are tied to decisions about capital, liquidity, limits and strategy.
- Finally, judge a stress test by whether it changes what the bank does.
- If it never does, it is decoration.

24 Why operational risk is hard to stress

- Harder than stress-testing market or credit risk
- Severe losses are rare and heavy-tailed
- Tail data are scarce and unstable
- Drivers: people, processes, systems, external events
- Weak link to the economy, shifting as controls change

IN THE LESSON

- Now the hardest case.
- Kestrel wants to stress test its operational risk, the risk of loss from failed people, processes and systems, or from external events.
- This is harder than stress testing market or credit risk.
- Severe operational losses are rare and heavy tailed.
- Heavy tailed means the biggest losses are vastly larger than the typical one: a bank may record thousands of small errors and, once in decades, a rogue trader or a giant fine.
- So tail data are scarce, and estimates of the tail are unstable, since one new large loss can change them a lot.
- The losses are driven mainly by people, process, system and external failures, so their link to the economy is weak.
- And that link shifts as controls change.

25 Data problems

- Internal data truncated at a collection threshold
- Losses recognised years after the event
- External data must be scaled to the firm
- Peer: $20bn revenue, $200m loss. Kestrel ($5bn): $50m
- The scaling rule is itself uncertain

IN THE LESSON

- Even the data Kestrel has are imperfect, in three ways.
- First, internal loss databases are truncated: losses below a collection threshold, say twenty thousand dollars, are never recorded.
- Second, losses are often recognised years after the event that caused them.
- So recent years look cleaner than they are.
- Third, to fill the tail, banks use external data, meaning losses suffered by other firms, but these must be scaled to the firm.
- Suppose a peer with twenty billion dollars of revenue lost two hundred million.
- Kestrel's revenue is five billion, so scaled in proportion, the loss becomes fifty million.
- But losses need not grow in proportion to size, so the scaling rule is itself uncertain.

26 Example 8: the truncation trap

Collection threshold $20,000. One year of true losses at Kestrel

Losses below $20,000	700 × $5,000, not recorded	**$3.5m**
Losses above $20,000	300 × $100,000, recorded	**$30.0m**
Frequency: 300 recorded v true	700 + 300 = 1,000	**understated**
Average: $100,000 v true	$33.5m ÷ 1,000 = $33,500	**overstated**

Fit severity conditional on the threshold, not as if complete

IN THE LESSON

- See how truncation distorts a model.
- Last year Kestrel really suffered seven hundred losses averaging five thousand dollars, three point five million in all, and none were recorded, because each fell below the twenty thousand dollar threshold.
- It also suffered three hundred losses averaging one hundred thousand dollars, thirty million in all, and these were recorded.
- Fit a model to the recorded data as if they were complete, and the frequency is three hundred events a year, while the true count is seven hundred plus three hundred, one thousand.
- So frequency is understated.
- The recorded average loss is one hundred thousand dollars, while the true total of thirty three point five million over one thousand losses gives thirty three thousand five hundred dollars.
- So the average is overstated.
- The cure is to fit severity conditional on the threshold, treating the data as truncated.

27 Scenario workshops and their biases

Expert judgement fills the gaps, so it needs structure and challenge

Bias in the workshop	Remedy
Anchoring on past losses	Structured elicitation: ranges, steps
Availability: vivid events loom	Internal and external loss data
Overconfidence: narrow ranges	Ask for 1-in-10 and 1-in-100 losses
A senior voice dominates	Independent challenge, documented

IN THE LESSON

- Because data are thin, banks run scenario workshops, in which experts estimate how often a severe event could happen and how much it would cost.
- These depend on expert judgement, which is exposed to biases.
- The main one is anchoring on past losses: experts fix on the largest loss they remember and adjust too little.
- The remedy is structured elicitation, asking in a set, careful sequence, for ranges and in steps.
- Availability bias makes vivid or recent events loom too large.
- The remedy is evidence from internal and external loss data.
- Overconfidence makes ranges too narrow.
- The remedy is to ask separately for the one in ten year and the one in a hundred year loss.
- And a senior voice can dominate the room.
- The remedy is independent challenge, with the reasoning documented.

28 Example 9: a common cause

Cyber loss chance 2% a year in each line: $150m, $100m, $80m

Stand-alone losses together	150 + 100 + 80	**$330m**
All three hit, if independent	2% × 2% × 2%	**1 in 125,000**
99% loss, assumed independent	retail 2%; >$150m only 0.12%	**$150m**
99% loss, one common cause	2% chance all are hit	**$330m**
Understatement	330 − 150	**$180m**

IN THE LESSON

- One cyber attack can hit several business lines at once.
- At Kestrel, retail, payments and markets each face a two per cent chance a year of a cyber loss: one hundred and fifty million dollars, one hundred million and eighty million.
- Together that is three hundred and thirty million.
- If the model assumes independence, all three strike in the same year with a chance of two per cent cubed, one in one hundred and twenty five thousand.
- Now the ninety nine per cent loss.
- Retail alone is hit with a two per cent chance, more than one in a hundred, so that loss is at least one hundred and fifty million.
- Anything above needs two lines hit together, only about a point one two per cent chance.
- So it is just one hundred and fifty million.
- But if one attack hits all three, a two per cent chance of losing three hundred and thirty million makes that the ninety nine per cent loss.
- Independence understated it by one hundred and eighty million.

29 Example 10: stress frequency and severity

Kestrel's process losses: 40 events a year, average $0.5m

Base case	40 × 0.5	**$20m**
Frequency up 50%	60 × 0.5	**$30m**
Severity up 50%	40 × 0.75	**$30m**
Both up 50%	60 × 0.75	**$45m**

99.9% loss: 1 year in 1,000, so no backtest. Benchmark, review

IN THE LESSON

- Two final points.
- A stress must move both frequency, how often losses happen, and severity, how large each one is.
- Kestrel's process losses run at forty events a year averaging half a million dollars, twenty million in total.
- Raise frequency by fifty per cent, to sixty events, and the total is thirty million.
- Raise severity by fifty per cent instead, to point seven five million, and it is also thirty million.
- Raise both and it is forty five million, more than either alone, because the effects multiply.
- The second point is validation.
- A ninety nine point nine per cent loss is exceeded once in a thousand years on average, so a bank cannot backtest it against its own history.
- Validation relies instead on benchmarking against peers and other models, and on independent review.

30 Mistakes to avoid: operational risk stress tests

- Treating truncated data as complete
- Leaning on macro links that are weak and shifting
- Accepting workshop numbers without challenge
- Assuming loss types are independent
- Stressing frequency or severity, not both

IN THE LESSON

- Here are the traps for operational risk stress testing.
- Do not treat truncated data as complete, or frequency will look too low and the average loss too high.
- Do not lean heavily on links to the economy, because operational losses are driven by people, processes and systems, and the link moves as controls change.
- Do not accept workshop estimates without structured challenge, because anchoring on past losses is strong.
- Do not assume loss types are independent, since a single cause such as a cyber attack can strike several business lines at once, and independence understates the total.
- And do not stress only frequency or only severity: stress both.

31 Recap

- ERM: governance, appetite and culture work together
- Regulatory floor v economic capital (VaR − EL)
- Basel III: 4.5%, 6%, 8% of RWA + 2.5% buffer; op × 12.5
- Stress tests: governance to reporting; severe but plausible
- Op risk: scarce, truncated data; bias; dependence

IN THE LESSON

- Let us gather the ideas.
- An E R M framework works only when governance, appetite and culture work together: who is accountable, how much risk, and whether people follow the rules.
- Regulatory capital is a floor set by prescribed rules, while economic capital is the bank's own estimate, V a R minus expected loss, at a confidence level tied to its target rating.
- Basel three sets minima of four point five, six and eight per cent of risk-weighted assets, plus a two point five per cent buffer, and operational risk capital becomes risk-weighted assets when multiplied by twelve point five.
- A robust stress testing framework runs from governance to reporting, with severe but plausible scenarios.
- And operational risk is the hardest to stress, because of scarce, truncated data, biased judgement and hidden dependence.

32 Try three quick questions

1. Exposure $20m, PD 2%, LGD 40%, WCDR 10%. IRB capital?
   Answer: K = 40% × (10% − 2%) = 3.2%. Capital $640,000.
2. One-year 99.9% VaR $900m, expected loss $150m. Economic capital?
   Answer: EC = 900 − 150 = $750m.
3. Cyber losses in three lines modelled as independent. Effect?
   Answer: Understates the total: one attack hits all.

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- An exposure of twenty million dollars has a default probability of two per cent, a loss given default of forty per cent, and a worst case default rate of ten per cent.
- What is the I R B capital?
- K is forty per cent times ten minus two per cent, which is three point two per cent, so capital is six hundred and forty thousand dollars.
- Question two.
- A one year ninety nine point nine per cent V a R is nine hundred million dollars, and expected loss is one hundred and fifty million.
- What is economic capital?
- Nine hundred minus one hundred and fifty, which is seven hundred and fifty million dollars.
- Question three.
- A model treats cyber losses in three business lines as independent.
- What does that do to the stressed total?
- It understates it, because one attack can hit all three lines at once.