CHARTA LEARNING · LESSON NOTES

Case Study: Model Risk and Validation

FRM Part II

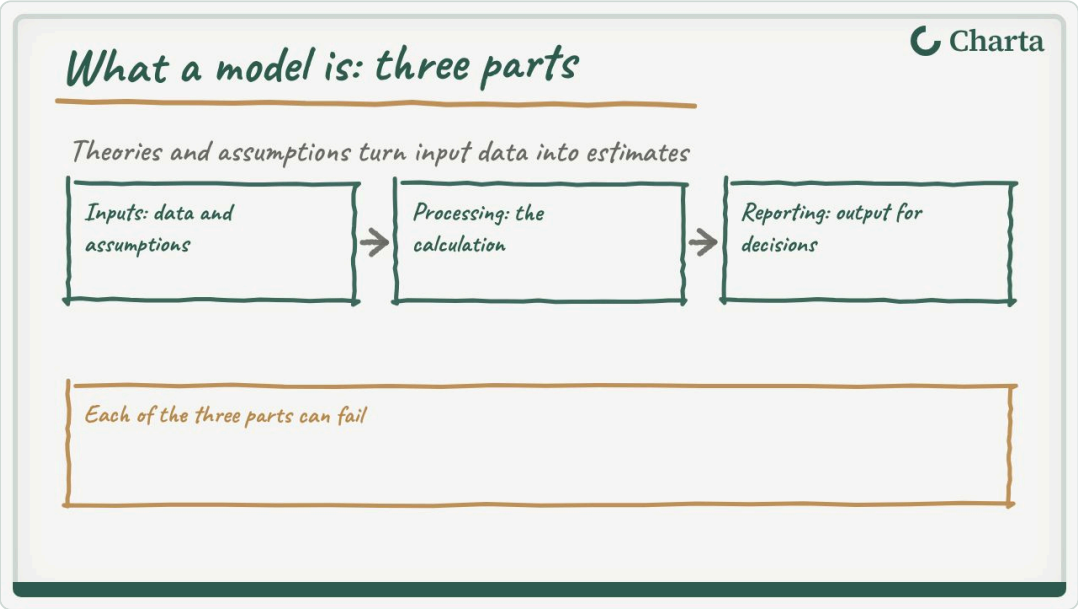
1 A story: the number everyone trusted

- One daily figure: trading value at risk
- Calm for two years; nobody asks how
- Markets turn: losses far bigger
- No fraud: short history, unchecked step
- One number treated as the truth

IN THE LESSON

- Imagine Fernhill Bank on an ordinary Monday.
- Every morning a single figure lands on the desk of the chief risk officer: the trading book's value at risk, a model's estimate of how much the bank could lose on a bad day.
- For two years the figure has been calm, and nobody asks how it is produced.
- Then markets turn.
- Losses arrive that are several times larger than the figure ever suggested.
- The inquiry finds no fraud and no villain.
- It finds a model built on a short, quiet history, a spreadsheet step nobody checked, and managers who treated one number as the truth.
- That is model risk.
- To understand it, we first need to be precise about what a model is.

2 What a model is: three parts



IN THE LESSON

- Here is the definition.
- A model is a quantitative method, system or approach that applies statistical, economic, financial or mathematical theories and assumptions to turn input data into quantitative estimates.
- Look at the three boxes.
- The first is the inputs: the data, and the assumptions that come with it.
- The second is the processing component, where the theory is applied and the calculation is done.
- The third is the reporting component, which translates the result into something a decision maker can use, such as a daily report, a price or a credit score.
- Underneath the boxes is the point to remember.
- Each part can fail.
- Bad data, a flawed calculation or a misleading report can each turn a good idea into a bad decision.

3 Is it a model? Fernhill's examples

Test: do theory and assumptions turn input data into estimates?

Tool at Fernhill	A model?
Trading value at risk engine	Yes: estimates a loss
Mortgage scoring tool	Yes: estimates default
Curve builder for pricing	Yes: theory fills gaps
Report adding up balances	No: plain arithmetic

IN THE LESSON

- Let us apply the definition to some of Fernhill's tools.
- The trading value at risk engine takes market data, assumes how prices move, and estimates a possible loss.
- That is a model.
- The mortgage scoring tool turns an applicant's data into a probability of default, using statistical theory.
- That is a model too.
- The curve builder, which turns market quotes into a full set of interest rates for pricing, rests on financial theory and assumptions about how to fill the gaps.
- Also a model.
- But a report that simply adds up the balances of all loans applies no theory and makes no assumptions.
- It is arithmetic, not a model in this sense.
- The test is whether theory and assumptions turn inputs into estimates.

4 What model risk is

- Adverse consequences of decisions based on
- incorrect or misused model output
- Financial loss: mispricing, thin reserves
- Poor decisions: wrong clients, limits
- Reputational damage: numbers were wrong

IN THE LESSON

- Now the risk itself.
- Model risk is the potential for adverse consequences from decisions based on incorrect or misused model output.
- The adverse consequences come in three kinds.
- There is financial loss, for example when a trade is priced too cheaply or a loss is reserved too thinly.
- There are poor business decisions, such as lending to the wrong customers or setting the wrong limits.
- And there is reputational damage when a bank has to admit that its numbers were wrong.
- Notice that the definition is about decisions.
- A model sitting unused on a server creates no risk.
- Risk arises when people act on what it says.
- That is why the reporting part matters as much as the mathematics.

5 Two sources of model risk

The model is wrong, or a sound model is used for the wrong job

Source	Examples
Fundamental error	Flawed theory, assumptions
Fundamental error	Bad data, coding mistakes
Fundamental error	Assumptions stop holding
Inappropriate use	Sound model, wrong job

IN THE LESSON

- Model risk arises from two broad sources, and the exam expects you to keep them apart.
- The first is fundamental errors.
- The model itself is wrong, even when it is used exactly as intended.
- That covers flawed theory, wrong assumptions, bad data and coding mistakes.
- It also covers assumptions that were true once but stop holding, for example a correlation that was stable for years and then changes.
- The second source is inappropriate use.
- Here the model may be perfectly sound, but it is applied outside the purpose or domain it was built for: a different product, a different kind of customer, or market conditions it never saw.
- Remember this sentence.
- A perfect model used for the wrong purpose still creates model risk.

6 Example 1: a sound model, the wrong job

Mortgage scoring model used on $400m of small business loans; LGD 45%

Default rate: model	predicted for small businesses	**1.5%**
Default rate: actual	small businesses	**6.0%**
Loss expected by model	400 × 1.5% × 45%	**$2.7m**
Actual loss	400 × 6.0% × 45%	**$10.8m**
Shortfall	10.8 − 2.7	**$8.1m**

IN THE LESSON

- Here is inappropriate use at Fernhill.
- The bank has a good scoring model, built and tested on salaried mortgage borrowers.
- A new team, keen to grow, uses it to approve four hundred million dollars of small business loans.
- For those borrowers the model predicts a default rate of one point five per cent.
- The actual default rate turns out to be six per cent.
- Loss given default is forty five per cent.
- The model's expected loss was four hundred million times one point five per cent times forty five per cent, which is two point seven million dollars.
- The actual loss is four hundred million times six per cent times forty five per cent, ten point eight million.
- The shortfall is eight point one million.
- The actual loss is four times what the model forecast.
- Nothing in the model was broken.

7 Example 2: severity grows with use

Same flaw in two models: a valuation error of 0.5%

Model A: one desk	positions valued	**$40m**
Model A: error	0.5% × 40m	**$0.2m**
Model B: curve builder	six desks, positions valued	**$3,000m**
Model B: error	0.5% × 3,000m	**$15m**

Same flaw, 75 times the damage: breadth of use and exposure

IN THE LESSON

- How badly can model risk hurt?
- Its severity grows with how widely a model is used and how much exposure rests on it.
- Suppose two of Fernhill's models each have the same small flaw, a valuation error of half of one per cent.
- Model A is used by one desk to value forty million dollars of positions.
- Its error is half a per cent of forty million, which is two hundred thousand dollars.
- Model B is the shared curve builder.
- It feeds six desks and values three thousand million dollars of positions.
- The same flaw now costs fifteen million dollars, seventy five times as much.
- The flaw is identical.
- What differs is breadth of use and exposure.
- That is why a bank sorts its models into tiers by how much depends on them.

8 Mistakes to avoid: what model risk is not

- Not only a coding problem
- A perfect model, wrong purpose: still a risk
- Assumptions that held can stop holding
- Severity: breadth of use and exposure

IN THE LESSON

- Before we move on, four mistakes to avoid.
- First, model risk is not only a coding problem.
- Code is one of many places a model can go wrong, alongside theory, assumptions, data and use.
- Second, a perfect model can still create model risk if it is used for the wrong purpose, as Fernhill's scoring model showed.
- Third, a model that worked last year can fail this year, because assumptions that held in the past can stop holding.
- And fourth, do not judge a model's risk only by its own quality.
- Ask how widely it is used and how much exposure rests on it.
- A small flaw in a model used everywhere can matter more than a large flaw in a model nobody relies on.

9 Why a model risk management function exists

- Identify, measure and control model risk
- Across the whole institution
- Independent of the people who build models
- and of the people who use them
- Someone with no stake must be able to say no

IN THE LESSON

- So who looks after all this?
- Fernhill's board decides it needs a model risk management function.
- Its purpose is to identify, measure and control model risk across the whole institution.
- Identify means knowing which models exist and where they are used.
- Measure means judging how much could go wrong and how much rests on each model.
- Control means making sure models are approved, checked and limited.
- The vital word is independently.
- The function must be independent of the people who build models and the people who use them.
- Builders are proud of their work, and users want the model's answer, especially when it makes a trade look profitable or a risk look small.
- Someone without that stake must be able to say no.

10 The function's toolkit

Five tools of a model risk management function

Tool	What it does
Complete model inventory	Every model, owner, purpose
Tiering by materiality	Most attention where it matters
Governance and approval	Clear owner; sign off before use
Documentation standards	Others see how it works, its limits
Independent validation	Effective challenge

IN THE LESSON

- The function has a toolkit, and you should be able to list it.
- First, a complete model inventory: a register of every model, its owner, its purpose and its validation status.
- Second, tiering, which means ranking models by materiality, so that the most important models get the most attention.
- Third, governance, with clear ownership and approval before use, so that no model goes live without sign off.
- Fourth, documentation standards, so that someone other than the builder can understand how a model works and where its limits are.
- And fifth, independent validation with effective challenge, which is the heart of the function.
- We now take tiering with numbers, and then validation in detail.

11 Example 3: Fernhill's inventory and tiers

Fernhill policy: Tier 1 from $1,000m; Tier 2 from $100m; else Tier 3

Trading VaR engine	$2,500m rests on it	**Tier 1**
Mortgage scoring model	$1,200m	**Tier 1**
Deposit run-off model	$600m	**Tier 2**
Marketing response model	$30m	**Tier 3**

Validated every 1, 2 or 3 years by tier

IN THE LESSON

- Here are four entries from Fernhill's inventory, with the exposure that rests on each.
- The bank's own policy, which we have made up, puts a model in tier one if one thousand million dollars or more rests on it, tier two from one hundred million up to that, and tier three below one hundred million.
- The trading value at risk engine carries two thousand five hundred million, so tier one.
- The mortgage scoring model carries one thousand two hundred million, also tier one.
- The deposit run off model, which predicts how fast deposits leave, carries six hundred million, tier two.
- The marketing response model carries thirty million, tier three.
- Under the same policy, tier one is validated every year, tier two every two years and tier three every three years.

12 Who does what with a model

Owner, validator, auditor and who is accountable

Role	Who at Fernhill
Owns and uses the model	The business, e.g. trading desk
Validates and challenges	Independent model risk function
Checks the framework works	Internal audit
Ultimately accountable	Board and senior management

IN THE LESSON

- A common exam trap is to confuse who owns a model, who validates it and who is ultimately accountable.
- At Fernhill, the owner is the business that builds and uses the model, for example the trading desk for the value at risk engine.
- The owner is responsible for the model working and for using it properly.
- The validator is the independent model risk function.
- It tests the model and challenges the owner, but it does not own the model.
- Internal audit checks that the whole framework works, including whether validation is done well.
- And ultimate accountability sits with the board and senior management, who set the appetite for model risk, approve policies and receive reports on model risk in aggregate.

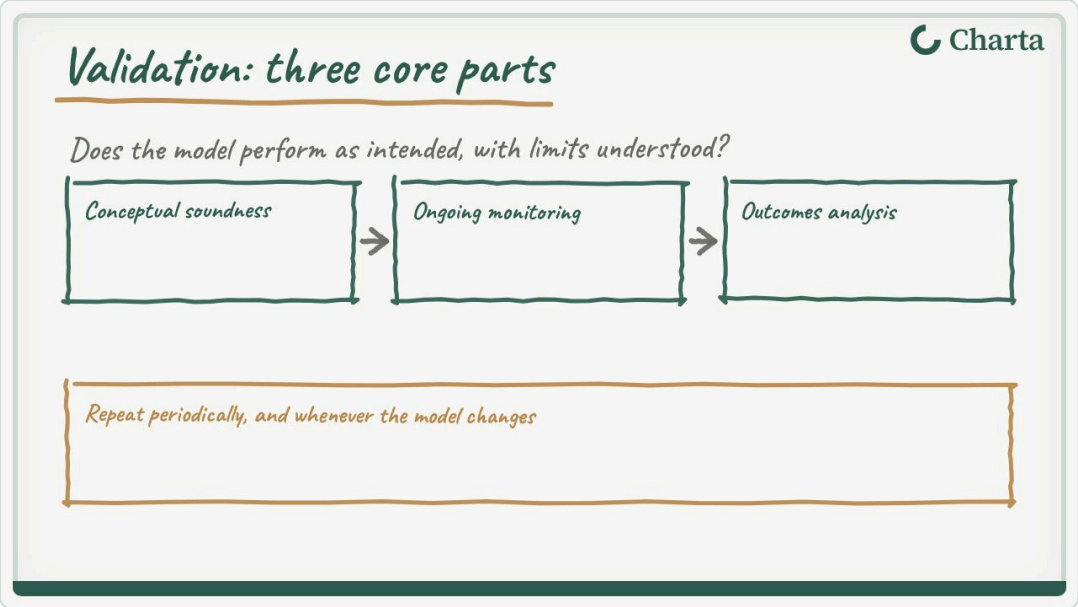
13 Effective challenge

- Critical analysis that can change outcomes
- Competence: maths, data and business
- Standing: senior enough to be heard
- Authority: can question and restrict a model
- Key control: independent of builders, users

IN THE LESSON

- The idea that ties the function together is effective challenge.
- It means critical analysis of a model by people who can actually change the outcome.
- Three things are needed.
- Competence: the validators must understand the mathematics, the data and the business well enough to find real problems.
- Standing: they must be senior and respected enough that their findings are taken seriously.
- Authority: they must be able to question a model and to restrict it, for example by delaying its use, limiting what it is used for or requiring changes.
- Challenge without authority is only an opinion.
- That is why the exam point is that independence from developers and users is the key control.
- A validator who reports to the trading desk cannot challenge it effectively.

14 Validation: three core parts



IN THE LESSON

- Validation is the work of checking that a model performs as intended and that its limits are understood.
- It has three core parts.
- Look at the boxes.
- The first is evaluating conceptual soundness: is the design right for the purpose, before the model is ever used?
- The second is ongoing monitoring: does the model stay appropriate as products and markets change?
- The third is outcomes analysis: when we compare the model's forecasts with what really happened, do they match?
- Underneath is the rule about timing.
- Validation is repeated periodically, and also whenever the model changes.
- It is not something done once, at launch, and then forgotten.
- We now look at each part with Fernhill's value at risk engine.

15 Conceptual soundness

- Review the design and the theory behind it
- Assumptions: reasonable for these positions?
- Data: how long, and does it include stress?
- Evidence: do the builders' tests support it?
- Finding: only three calm years of data

IN THE LESSON

- Conceptual soundness asks whether the model was built on a sound basis.
- The validator reviews the design and the theory behind it.
- For Fernhill's value at risk engine, is the method suitable for the kinds of positions the desk holds?
- Next come the assumptions.
- Does the engine assume price changes are normally distributed, and is that reasonable for these positions?
- Then the data.
- How long is the history, and does it include any stressed periods?
- Last, the developmental evidence: what tests did the builders run, and do the results support their choices?
- Here the validator finds a problem.
- The engine was calibrated on only three years of calm markets.
- That is a limitation, and it must be recorded and dealt with.

16 Ongoing monitoring

- Still fit as products and markets change?
- Process verification: feeds, approved code
- Benchmarking against an alternative model
- Sensitivity: small change, sensible move?
- New products the model was never built for?

IN THE LESSON

- The second part is ongoing monitoring, which confirms that the model stays appropriate as products and markets change.
- One tool is process verification.
- Do the data feeds arrive complete and correct, and does the code in production match the code that was approved?
- Another is benchmarking: compare the model with an alternative, such as a simpler model or a vendor's figure.
- A large, unexplained difference is a warning sign.
- A third is sensitivity analysis: change an input a little and see whether the output moves sensibly.
- If a tiny change in one input swings the answer wildly, the model is fragile.
- Monitoring also asks whether the desk has started trading products the model was never built for.

17 Outcomes analysis: the back-test

Exception: a day's loss bigger than VaR
Expected exceptions = N × (1 − c)

Expected exceptions	If the model is right
Days in the test	N: trading days in the test
Confidence level	c: e.g. 99%
Exceptions observed	Too many: understates risk

IN THE LESSON

- The third part is outcomes analysis, comparing forecasts with actual results.
- The classic example is back-testing a value at risk model.
- An exception is a day on which the actual loss is bigger than the value at risk figure.
- If a model with a ninety nine per cent confidence level is right, losses should exceed it on about one day in a hundred.
- So the expected number of exceptions is the number of days in the test, N, times one minus the confidence level, c.
- Then we count the exceptions observed and compare.
- Far too many exceptions means the model understates risk.
- Far too few can mean it overstates it.

18 Example 4: back-testing Fernhill's VaR

99% VaR, back-tested over the last 250 trading days

Expected exceptions	250 × (1 − 0.99)	**2.5**
Observed exceptions	losses bigger than VaR	**8**
Chance of 8 or more if right	binomial, p = 1%	**about 0.4%**
Verdict	fits the calm-data finding	**Understates risk**

IN THE LESSON

- Let us back-test Fernhill's engine.
- It reports a ninety nine per cent value at risk, and we look at the last two hundred and fifty trading days.
- The expected number of exceptions is two hundred and fifty times one per cent, which is two point five.
- The bank actually saw eight exceptions.
- If the model were right, how likely is it to see eight or more by chance?
- Using the binomial distribution, the probability is about zero point four per cent, roughly four chances in a thousand.
- So bad luck is a very unlikely explanation.
- The evidence says the model understates risk, which fits what conceptual soundness found: three calm years of data.
- Two parts of validation point the same way.

19 Why back-testing alone is not enough

A back-test is one tool, not the whole of validation

Limit of a back-test	What covers it
Calm model passes calm test	Stress tests, soundness
Little on the rare tail	Conceptual soundness
Silent on purpose and use	Governance, use restrictions
Cannot check feeds or code	Process verification

IN THE LESSON

- A back-test is powerful, but on its own it is not sufficient, and the exam likes this point.
- First, a back-test only looks at the past.
- A model calibrated on calm years can pass a back-test run in calm years, and then fail in the first storm; stress tests and conceptual soundness cover that.
- Second, it can only judge outputs we observe often, like daily losses.
- It says little about the rare tail, where the biggest losses live; conceptual soundness covers that.
- Third, a pass says nothing about whether the model is being used for the right purpose; governance and use restrictions cover that.
- And fourth, it cannot check the data feeds or the code; process verification covers that.
- That is why validation needs all three parts, not just the back-test.

20 Example 5: living with a known limitation

Model capital $80m; validator finds the data covers only calm years

Model output: capital	credit model	**$80m**
Conservative adjustment	15% × 80m	**$12m**
Capital held	80 + 12	**$92m**
Use restriction	built on prime borrowers	**Prime only**
Limit	until revalidated	**Book capped**

IN THE LESSON

- Validation rarely ends with a perfect model.
- More often it finds limitations that remain, and the function handles them in three ways: conservative adjustments, limits and use restrictions.
- At Fernhill, a credit model gives a capital figure of eighty million dollars, but the validator knows it was calibrated on calm years.
- So the function requires a conservative adjustment of fifteen per cent, which is twelve million, and the bank holds ninety two million.
- It also restricts use: the model may be used only for prime mortgages, the borrowers it was built on.
- And it sets a limit, so the book it covers cannot grow beyond an agreed size until the model is rebuilt and validated again.

21 Mistakes to avoid: validation

- Not a one-time back-test at launch
- Builders and users never validate their own
- Keep owner, validator, auditor, board apart
- Challenge needs authority to restrict
- A passed back-test is not proof of soundness

IN THE LESSON

- Here are the mistakes to avoid on management and validation.
- First, treating validation as a one time back-test at launch.
- Validation has three parts and is repeated periodically and whenever the model changes.
- Second, letting developers or users validate their own models.
- Independence from developers and users is the key control.
- Third, confusing roles.
- The business owns the model, the independent function validates it, internal audit checks the framework, and the board and senior management are ultimately accountable.
- Fourth, thinking a validator's job is just to write a report.
- Effective challenge needs the authority to restrict a model.
- And fifth, believing that a passed back-test proves a model is sound.
- It is necessary evidence, but it is not sufficient.

22 Three case studies, two ways to fail

- JPMorgan's London Whale, 2012
- Long-Term Capital Management, 1998
- Credit derivatives priced by Gaussian copula
- 1: benign calibration breaks; leverage magnifies
- 2: manual build, rushed, weak validation

IN THE LESSON

- Now to the three case studies.
- The first is J P Morgan's London Whale episode in twenty twelve.
- The second is the collapse of Long Term Capital Management, known as L T C M, in nineteen ninety eight.
- The third is the pricing of credit derivatives with a Gaussian copula.
- They look very different, but they teach recurring lessons, and these fall into two families.
- In the first family, the model's assumptions were calibrated on a short and benign history, and they broke in stress.
- Leverage then made the loss far bigger.
- In the second family, the model was built or changed in a manual, poorly controlled process and approved in a hurry, without proper independent validation.
- Let us take the London Whale first.

23 The London Whale, 2012

- Big credit derivative positions; 2012 losses
- A new VaR model reported much lower risk
- Built in spreadsheets with manual steps
- Rushed approval; cursory review only
- Lesson: a convenient drop needs scrutiny

IN THE LESSON

- The London Whale was the nickname for a trader at J P Morgan whose very large positions in credit derivatives lost several billion dollars in twenty twelve.
- For model risk, the heart of the story is a model change.
- A new value at risk model was introduced, and it reported a much lower risk number for the positions.
- The model had been built in a manual, poorly controlled way, using spreadsheets with figures copied by hand.
- It was approved in a hurry, after only a cursory review, without the independent validation a model of that importance needed.
- So when the reported risk fell, the lower number looked like good news instead of a warning.
- The lesson: a model change that conveniently lowers a risk number needs extra scrutiny, not less.

24 Example 6: the change that lowers the number

Desk VaR limit $100m. Old model and proposed new model, same positions

Old model VaR	above the $100m limit	**$120m**
New model VaR	inside the limit	**$66m**
Drop from the change	(120 − 66) ÷ 120	**45%**
Reconciliation finds	divides by sum, not average	**Error**
New model, corrected	within 2% of old; over limit	**$118m**

IN THE LESSON

- Here is the same pattern at Fernhill, with made up numbers.
- A desk's value at risk under the old model is one hundred and twenty million dollars, above its limit of one hundred million.
- The desk proposes a new model, which shows sixty six million on the same positions, forty five per cent lower, and comfortably inside the limit.
- That convenient drop is the warning sign.
- So the model risk function insists on a parallel run, running old and new side by side, and a reconciliation that explains every difference.
- The reconciliation finds a spreadsheet step that divides by a sum instead of an average.
- Corrected, the new model shows one hundred and eighteen million, within two per cent of the old.
- The desk is still over its limit.

25 LTCM, 1998

- Hedge fund betting small price gaps close
- Borrowed heavily: high leverage
- Assumed stable correlations, liquidity
- Both broke in stress; leverage magnified
- Lesson: stress test the assumptions

IN THE LESSON

- The second case is Long Term Capital Management, a hedge fund that collapsed in nineteen ninety eight.
- It made trades that bet on small price gaps between similar securities closing, and it borrowed heavily to make those small gaps worth the effort.
- For model risk, the heart of the story is calibration.
- The fund's models rested on assumptions drawn from a short and benign history, in particular that correlations between positions would stay stable and that liquidity would stay constant, so that positions could always be sold.
- In the stress of nineteen ninety eight, correlations jumped and liquidity vanished.
- Leverage then magnified the loss.
- The lesson: stress test the assumptions themselves, not just the positions.

26 Example 7: leverage magnifies a broken calibration

Positions $60m. Calm-history worst month 0.5%; stress loss 3%

Calm worst month	0.5% × 60m	**$0.3m**
Stress loss	3% × 60m	**$1.8m**
Equity $2m: 30 times	calm 15%; stress 1.8 ÷ 2	**90% lost**
Equity $6m: 10 times	stress 1.8 ÷ 6	**30% lost**

IN THE LESSON

- Here is a made up fund that shows how leverage turns a broken assumption into ruin.
- It holds sixty million dollars of positions.
- Calibrated on calm history, its worst monthly loss is half of one per cent, which is three hundred thousand dollars.
- In stress, when correlations jump and liquidity disappears, the loss is three per cent, or one million eight hundred thousand dollars.
- With equity of two million, that is thirty times leverage.
- The calm worst case would cost fifteen per cent of equity, but the stress loss wipes out ninety per cent.
- With equity of six million, ten times leverage, the same stress loss is thirty per cent of equity.
- Painful, but survivable.
- Same assumptions, same mistake; only the leverage decides whether the fund survives.

27 The Gaussian copula and credit derivatives

- Copula: joins borrowers' default risks
- Gaussian: one correlation; a convention
- Calibrated on a short, benign history
- Widely used, so treated as truth
- Lesson: a market convention is not the truth

IN THE LESSON

- The third case is the pricing of credit derivatives with a Gaussian copula.
- A copula is a way of joining the default risks of many borrowers into one picture of how they default together.
- The Gaussian version summarises that joint behaviour with a correlation, and it became a market convention for pricing products whose value depends on several borrowers defaulting at once.
- The trouble was twofold.
- The correlation was calibrated on a short and benign history, when defaults were rare and scattered, so joint defaults looked unlikely.
- And because everyone used it, the convention's answer was treated as truth.
- When stress came, defaults clustered far more than the model allowed.
- The lesson: a market convention is not the truth.

28 Example 8: what correlation does to a joint default

Two loans, each PD p = 5%; senior piece loses only if both; rho = default correlation

Both default, correlation 0	5% × 5%	**0.25%**
Default correlation 0.05: benign	p^2 + rho × p × (1 − p)	**0.49%**
Default correlation 0.30: stress	same formula	**1.68%**
Stress versus benign	1.68 ÷ 0.49	**3.4 times**

IN THE LESSON

- Here is why correlation matters so much.
- Take two loans, each with a five per cent chance of default, which we call p.
- A senior piece of a deal loses money only if both default.
- In a simple formula, rho is the default correlation, the correlation between the two default events.
- The copula's own correlation enters differently, but the direction is the same: higher correlation, far more joint defaults.
- With correlation zero, the chance both default is five per cent times five per cent, a quarter of one per cent.
- With a default correlation of zero point zero five, calibrated on benign years, the formula p squared plus rho times p times one minus p gives about zero point four nine per cent.
- In stress the default correlation is zero point three, and the chance rises to about one point six eight per cent.
- That is about three point four times the benign figure.
- The single default chances never changed.
- Only the correlation did, and the senior piece was far riskier than its price said.

29 The common lessons, and the remedies

Remedies: governance, independent validation, stress testing assumptions

What went wrong	Remedy
Benign calibration broke in stress	Stress test the assumptions
Leverage magnified the loss	Limits that allow for model error
Manual build, hurried approval	Independent validation before use
Change lowered a risk number	Parallel run and reconciliation
One figure treated as truth	Management knows model limits

IN THE LESSON

- Now put the three cases together.
- Each failure has a remedy, and the remedies all come back to three themes: governance, independent validation and stress testing of assumptions.
- First, assumptions calibrated on a short, benign history broke in stress, at L T C M and with the copula.
- The remedy is to stress test the assumptions.
- Second, leverage magnified the loss.
- Governance must set limits that assume the model can be wrong.
- Third, models were built or changed in manual, poorly controlled processes and approved in a hurry.
- The remedy is independent validation before use.
- Fourth, a change that lowered a risk number went unchallenged.
- It needs parallel runs and reconciliation.
- Fifth, a convention or a single figure was treated as truth.
- Management must understand model limits.

30 Try three quick questions

1. 99% VaR back-tested over 500 days. Expected exceptions?
   Answer: 500 × 1% = 5 exceptions.
2. A validated pricing model is used for a quite different product. Model risk? Source?
   Answer: Yes: inappropriate use of a sound model.
3. A model change cuts VaR from $80m to $50m. By how much? What next?
   Answer: 37.5% lower: extra scrutiny, parallel run, reconcile.

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- A bank back-tests a ninety nine per cent value at risk model over five hundred trading days.
- How many exceptions should it expect if the model is right?
- Five hundred times one per cent, which is five exceptions.
- Question two.
- A pricing model that is well built and fully validated for one product is used without change to price a quite different product.
- Is there model risk, and from which source?
- Yes.
- It is inappropriate use: a sound model applied outside its purpose.
- Question three.
- A proposed model change cuts a desk's value at risk from eighty million dollars to fifty million.
- By how much, and what should happen?
- It falls by thirty seven and a half per cent, so it needs extra scrutiny, a parallel run and reconciliation with the old model.