CHARTA LEARNING · LESSON NOTES

# Operational Risk

FRM Part I

## 1 A story: the day the machine breaks

- Machine of people, processes and systems
- Loss from failed process, people, systems or outside events
- Basel definition includes legal risk
- Excludes strategic and reputational risk

IN THE LESSON

- Picture a bank as a large machine made of people, procedures and computers.
- Most days it runs quietly.
- Then one day a trader hides a losing position from his managers.
- Another day a storm floods the data centre.
- Another day a salesperson sells a product to a customer who should never have been offered it.
- None of these is a market move or a borrower failing to repay.
- They are failures of the machine itself.
- That is operational risk: the risk of loss from inadequate or failed internal processes, people and systems, or from external events.
- The Basel definition includes legal risk, but leaves out strategic and reputational risk.

## 2 Basel's event types, part one

*Seven event types: the first four*

| Event type | Example |
|---|---|
| Internal fraud | Unauthorized trading by an insider |
| External fraud | Forgery, hacking, theft |
| Employment and safety | Discrimination claim, safety breach |
| Clients, products, practices | Mis-selling, money laundering failure |

IN THE LESSON

- Basel sorts operational losses into seven event types.
- The first four are about wrongdoing and dealings with people.
- Internal fraud needs at least one insider who acts to cheat the firm or break its rules, such as unauthorized trading.
- Barings in nineteen ninety five and Société Générale in two thousand eight are the famous cases.
- External fraud is done by outsiders: forgery, theft, hacking.
- Employment practices and workplace safety covers discrimination claims and breaches of health and safety law.
- Clients, products and business practices covers improper or careless dealings with customers, such as mis-selling products or failing to stop money laundering.
- Intent decides: a deliberate insider act is fraud, while an honest mistake belongs in execution, or in clients and products if a customer was harmed.

## 3 Basel's event types, part two

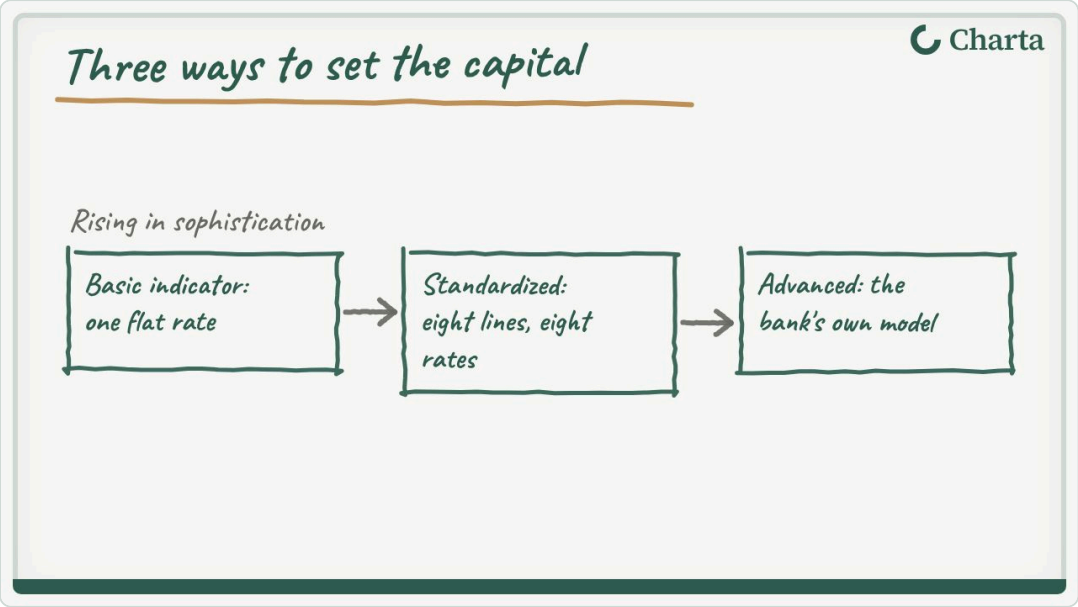
*Seven event types: the last three*

| Event type | Example |
|---|---|
| Damage to physical assets | Fire, flood, vandalism |
| Business disruption, systems | Hardware, software, utility outage |
| Execution, delivery, process | Data entry, settlement, vendor errors |

IN THE LESSON

- The other three types are about the physical world and the machinery.
- Damage to physical assets covers disasters such as fire or flood, and vandalism.
- Business disruption and system failures covers outages of hardware, software and utilities, for example a payments system that goes down for a day.
- Execution, delivery and process management is the everyday category: errors in processing, documentation, settlement and in managing outside vendors.
- A wrong amount keyed into a trade lands here.
- When you classify a loss, ask why it happened.
- Intent and insider involvement decide the fraud types.
- An honest slip belongs in execution, or in clients and products if a customer was harmed.
- Cause decides the box, not the size of the loss.

## 4 Three ways to set the capital




IN THE LESSON

- Now, how much capital should a bank hold against all this?
- Basel Two, the second international banking accord, offered three methods, climbing in sophistication.
- Read the boxes from left to right.
- The basic indicator approach is the crudest: one flat percentage of the bank's income.
- The standardized approach is a little finer: it splits the income into eight business lines and charges each a different rate.
- The advanced measurement approach lets the bank build its own model from its own loss data, subject to supervisors' approval.
- Only that last method responds to the bank's own loss experience and controls.
- The first two simply scale with income, so a bank with excellent controls pays the same as one with poor controls.

## 5 The basic indicator approach

**Capital = 15% × average gross income (last three years, positive years only)**

| | |
|---|---|
| **Capital** | Operational risk capital |
| **Gross income** | Average of the positive years |
| **Fixed rate** | Fifteen per cent, set by Basel |

IN THE LESSON

- Here is the simplest rule.
- Capital equals fifteen per cent of the bank's average annual gross income over the previous three years.
- Gross income is, roughly, net interest income plus net non-interest income such as fees.
- The fifteen per cent is a fixed rate set by the regulator.
- Now the exam trap.
- Only years in which gross income was positive are counted, both in the total and in the number of years you divide by.
- A year with a loss is dropped altogether.
- So if one of three years is negative, you average the other two.
- Banks with more income hold more capital, whether or not they run the business carefully, which is the weakness of this method.

## 6 Example 1: the basic indicator approach

*Gross income, in millions: 400, 500 and -100*

| | | |
|---|---|---|
| **Keep positive years** | 400 and 500 | **2 years** |
| **Average** | (400 + 500) ÷ 2 | **450** |
| **Capital** | 15% × 450 | **67.5** |
| **Wrong way** | (800 ÷ 3) × 15% | **40.0** |

IN THE LESSON

- Take a bank whose gross income was four hundred million in the first year, five hundred million in the second, and a loss of one hundred million in the third.
- The third year is negative, so it is dropped, and we only use two years.
- The sum of the positive years is nine hundred million.
- Divide by two, not by three, to get an average of four hundred and fifty million.
- Multiply by fifteen per cent.
- The capital charge is sixty seven point five million.
- The common mistake is to include the loss year, giving a total of eight hundred million divided by three, which understates the charge.
- The rule is to exclude, not to net.

## 7 Example 2: the standardized approach

*Trading 18%, retail 12%, commercial 15%. Incomes in millions*

| | | |
|---|---|---|
| **Year 1: 200, 300, 100** | 36 + 36 + 15 | **87** |
| **Year 2: -250, 200, 100** | -45 + 24 + 15 = -6, floor | **0** |
| **Year 3: 100, 300, 200** | 18 + 36 + 30 | **84** |
| **Capital** | (87 + 0 + 84) ÷ 3 | **57** |

IN THE LESSON

- The standardized approach splits income into eight business lines, each with its own beta between twelve and eighteen per cent.
- Trading and sales is eighteen, commercial banking fifteen, retail banking twelve.
- Take a bank with these three lines.
- In year one trading earns two hundred, retail three hundred and commercial one hundred million, so the charge is thirty six plus thirty six plus fifteen, eighty seven.
- In year two trading loses two hundred and fifty million while retail earns two hundred and commercial one hundred, and the negative offsets the others, giving minus six.
- A negative yearly total is set to zero, but that year still counts in the average.
- Year three, with one hundred, three hundred and two hundred, gives eighty four.
- The average of eighty seven, zero and eighty four, over three, is fifty seven million.

## 8 The three approaches side by side

*Basel II: three ways to set the capital*

| Approach | Basis and trap |
|---|---|
| Basic indicator | 15% of gross income; drop negative years |
| Standardized | 8 betas, 12-18%; floor year at zero |
| Advanced (AMA) | Own model, 99.9% one-year loss |
| Only AMA | Responds to own losses, controls |

IN THE LESSON

- Let us put the three approaches next to each other, with the traps.
- The basic indicator approach uses one rate of fifteen per cent of average gross income, and drops negative years from the average.
- The standardized approach uses a beta for each of eight business lines, and floors a negative yearly total at zero but keeps that year in the three year average.
- The advanced measurement approach uses the bank's own model, built from internal and external loss data, scenario analysis and control factors, to estimate the loss that will not be exceeded in a year with ninety nine point nine per cent probability.
- It needs supervisory approval.
- Only the advanced approach reflects the bank's own losses and controls.

## 9 Why Basel moved to the standardized measurement approach

- AMA: complex, hard to compare, open to manipulation
- Older standard methods: scale only with income
- SMA: simple, comparable base
- Plus the bank's own loss history

IN THE LESSON

- Why did Basel change course?
- The Basel Committee replaced all three older methods with one, the standardized measurement approach, or S M A.
- There were two problems.
- The advanced approach produced complicated capital figures that were hard to compare between banks, and open to manipulation through modelling choices.
- Two banks with similar risks could report very different capital.
- The older standardized methods had the opposite fault: they were simple, but not risk sensitive, because they scaled only with gross income.
- The S M A tries to keep a simple, comparable base and then adds the bank's own loss history on top.
- Simple where it can be, sensitive where it matters.

## 10 The S M A formula

**Capital = BIC × ILM**
**ILM = ln( e - 1 + (LC ÷ BIC)^0.8 )**

| | |
|---|---|
| **BIC** | Size charge, by marginal slices |
| **Loss component** | 15 × ten-year average loss |
| **Loss multiplier** | Above 1 if losses are high |

IN THE LESSON

- The S M A has three steps.
- First compute the business indicator, a measure of size built from interest, services and financial components.
- Second, turn it into the business indicator component, or B I C, using marginal rates: twelve per cent on the part up to one billion euros, fifteen per cent on the part between one and thirty billion, and eighteen per cent above thirty billion.
- Like a tax schedule, each rate applies only to its own slice.
- Third, multiply by the internal loss multiplier, which uses the loss component, fifteen times the bank's average annual operational losses over ten years.
- When the loss component equals the B I C, the multiplier is exactly one.

## 11 Example 3: the business indicator component

*Business indicator of EUR 10 billion. Slice by slice*

| | | |
|---|---|---|
| **First 1bn at 12%** | 0.12 × 1 | **0.12** |
| **Next 9bn at 15%** | 0.15 × 9 | **1.35** |
| **BIC** | 0.12 + 1.35 | **1.47** |
| **Wrong: 15% on all** | 0.15 × 10 | **1.50** |

IN THE LESSON

- Suppose a bank has a business indicator of ten billion euros.
- Apply the marginal rates like a tax schedule.
- The first billion is charged at twelve per cent, which is zero point one two billion.
- The remaining nine billion falls in the middle band, at fifteen per cent: one point three five billion.
- Add them: one point four seven billion euros.
- Notice we did not charge fifteen per cent on all ten billion, which would give one point five.
- That is the common mistake.
- The marginal rates apply slice by slice.
- A bank of forty billion would pay twelve per cent on the first billion, fifteen on the next twenty nine, and eighteen on the last ten, six point two seven billion in all.

## 12 Example 4: the loss multiplier at work

*BIC 1.47bn. Capital = BIC × ILM*

| | | |
|---|---|---|
| **Avg loss 98m: LC = BIC** | ILM = 1.00 | **1.47bn** |
| **Avg loss 40m: LC 0.6bn** | ILM = 0.79 | **1.16bn** |
| **Avg loss 200m: LC 3.0bn** | ILM = 1.25 | **1.84bn** |

IN THE LESSON

- Now add the bank's own loss history.
- Our B I C is one point four seven billion.
- Case one: average annual losses of ninety eight million.
- Fifteen times that is one point four seven billion, equal to the B I C, so the ratio is one and the multiplier is exactly one.
- Capital equals the B I C, one point four seven billion.
- Case two, a well controlled bank with average losses of forty million: the loss component is zero point six billion, the ratio is zero point four one, and the multiplier works out at zero point seven nine, so capital falls to about one point one six billion.
- Case three, average losses of two hundred million: the multiplier is one point two five, and capital rises to about one point eight four billion.

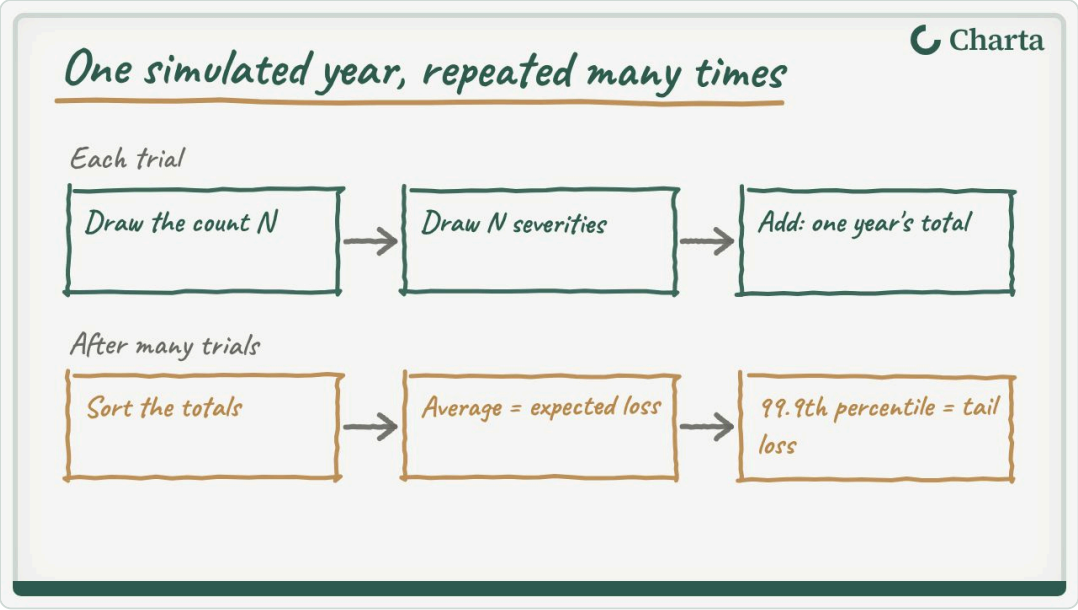
## 13 Monte Carlo: frequency and severity

- Frequency: how many losses a year (often Poisson)
- Severity: how large each loss is (often lognormal)
- Estimated separately, assumed independent
- A year's loss = sum of that many severities

IN THE LESSON

- Now to the models banks use to measure operational loss themselves.
- The idea is to build a year of losses from two separate pieces.
- The first is frequency: how many losses happen in a year.
- This is commonly modelled with a Poisson distribution, which has one parameter, the average count, called lambda.
- The second is severity: how big each loss is.
- This is commonly a lognormal distribution, which has a long right tail, so most losses are modest and a few are huge.
- The two are estimated separately and assumed to be independent.
- A year's total loss is the sum of as many severities as the frequency says.
- Combining them by simulation gives the shape of a whole year of risk.

## 14 One simulated year, repeated many times




IN THE LESSON

- Here is the Monte Carlo recipe.
- Read the boxes from the top row down.
- In every trial, first draw the number of losses for the year from the frequency distribution.
- Second, draw that many severities, one for each loss, independently, from the severity distribution.
- Third, add them up: that is the total loss for one simulated year.
- Then repeat this many thousands of times.
- In the bottom row, sort the totals from smallest to largest.
- They form the simulated loss distribution.
- Their average estimates the expected annual loss.
- The ninety nine point ninth percentile estimates the one year ninety nine point nine per cent loss.
- A frequent error is to multiply the average count by a single severity draw, which ignores the randomness in both pieces.

## 15 Example 5: one trial, then the tail

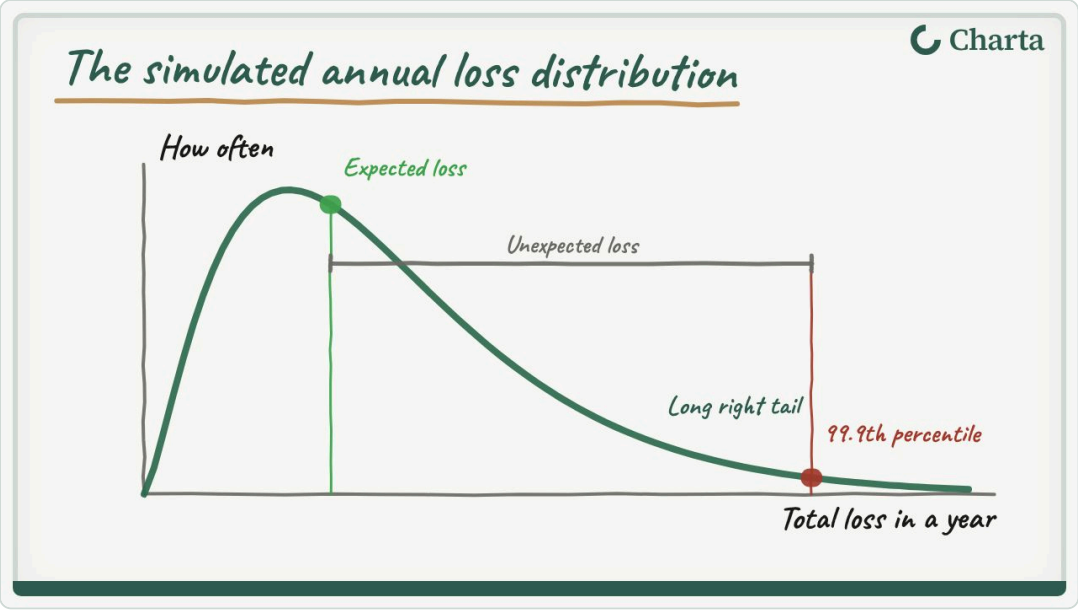
*Poisson mean 4 a year. Average severity 2.5m. 10,000 trials*

| | | |
|---|---|---|
| **Expected loss** | 4 × 2.5m | **10m** |
| **One trial: 4 losses** | 1.2 + 0.8 + 4.5 + 1.5 | **8.0m** |
| **99.9th of 10,000** | 10th largest, say | **90m** |
| **Unexpected loss** | 90m - 10m | **80m** |

IN THE LESSON

- Let us do it by hand with made up numbers.
- Say the frequency is Poisson with an average of four losses a year, and the average severity is two point five million.
- Expected annual loss is four times two point five, ten million.
- Now one trial: suppose the draw gives four losses, of one point two, point eight, four point five and one point five million.
- The total for this simulated year is eight point zero million.
- Repeat ten thousand times.
- With ten thousand trials, the ninety nine point ninth percentile is the tenth largest outcome.
- Suppose that is ninety million.
- Unexpected loss is that percentile minus expected loss: eighty million.

## 16 The simulated annual loss distribution




IN THE LESSON

- Here is what the simulation produces, drawn schematically.
- Along the bottom is the total loss for a year; up the side is how often each total occurs.
- Most years are small, so the curve is high on the left.
- But it stretches a long way to the right, because a few years contain a huge loss.
- The average, marked in green, is the expected loss, the cost of doing business, which a bank should price and provision for.
- Far out in the tail sits the ninety nine point ninth percentile, in red.
- The gap between the two is the unexpected loss, and that gap is what capital is held to absorb.
- The tail is where nearly all of the danger is.

## 17 Why loss data mislead

- Few extreme losses in a bank's own history
- Collection threshold: severity too high, count too low
- Public data skewed to large, newsworthy events
- Late discovery makes recent years look quiet
- Short histories, changing businesses, pooled unlike units

IN THE LESSON

- A model is only as good as its data, and operational loss data are poor in several ways.
- A bank's own history has few or no extreme losses, so the tail is guessed from almost nothing.
- Outside data from consortia or public sources help, but they bring problems.
- Databases usually record only losses above a collection threshold.
- Fit severity to them without adjusting and you overstate the average loss, because small ones are missing.
- Count only the recorded events and you understate the frequency.
- Public data lean toward large, newsworthy events, pushing severity up.
- They come from banks of different sizes, so need scaling.
- Losses are found late, so recent years look too quiet.
- Short histories, changing businesses and pooling unlike units weaken the estimates further.

## 18 Example 6: correcting for a threshold

*Threshold 10,000. 12 recorded a year. P(loss above threshold) = 0.6*

| | | |
|---|---|---|
| **Recorded count** | above the threshold | **12** |
| **True frequency** | 12 ÷ 0.6 | **20** |
| **Understated by** | (20 - 12) ÷ 20 | **40%** |
| **Severity fitted raw** | small losses missing | **too high** |

IN THE LESSON

- Suppose a database only records losses above ten thousand dollars.
- A bank sees twelve recorded losses a year.
- Does that mean twelve events happen?
- No.
- Many small ones fall below the line and are never written down.
- Suppose the fitted severity says that a loss exceeds the threshold with probability zero point six.
- Then the true frequency is the recorded count divided by that probability: twelve over zero point six is twenty events a year.
- Using twelve would understate frequency by forty per cent.
- The same missing small losses push the average recorded severity up, so a severity fitted without allowing for the threshold is too high.
- Both estimates need adjusting before the model is run.

## 19 Scenario analysis when data are thin

- Experts estimate how often and how large
- Turned into a frequency and a severity distribution
- Supplements internal and external loss data
- Risks: availability bias, anchoring, overconfidence

IN THE LESSON

- When history is thin, especially for rare and severe events such as a major cyber attack or a rogue trading loss the bank has never had, experts fill the gap.
- In structured workshops, managers describe plausible scenarios and estimate how often each might happen and how large the loss would be, often giving a typical case and a worst case.
- These judgements become a frequency distribution, such as a Poisson whose mean is the estimated events per year, and a severity distribution, such as a lognormal fitted to the quantiles the experts give.
- The weakness is subjectivity: availability bias, anchoring on the first number, and overconfidence in narrow ranges.
- Structured elicitation, challenge sessions and external data comparison limit these biases.

## 20 Example 7: from an expert's view to a Poisson mean

*Expert view: a serious cyber attack once in 25 years*

| | | |
|---|---|---|
| **Events a year (Poisson mean)** | 1 ÷ 25 | **0.04** |
| **Chance of at least one** | 1 - e^(-0.04) | **3.92%** |
| **Next: severity** | typical and worst case | **lognormal** |
| **Then** | join the loss model | **simulate** |

IN THE LESSON

- A workshop judges that a serious cyber attack happens once in twenty five years.
- How do we turn that into a frequency for the model?
- The mean number of events per year is one divided by twenty five, which is point zero four.
- That becomes the mean of a Poisson distribution.
- Under it, the chance of at least one attack in a given year is one minus e to the power of minus point zero four, about three point nine per cent: a little below the simple point zero four, because two attacks in a year are possible.
- Then the experts give a typical loss and a worst case, from which the severity distribution is fitted.
- The two together join the loss model.

## 21 Causes, education, RCSA and key risk indicators

- Causal analysis: what drives losses? Test the mechanism
- Correlation alone does not prove cause
- Education: staff report losses and near misses
- RCSA and KRIs: self-assessment and warning metrics

**IN THE LESSON**

- Measuring losses is not enough; a bank wants to know why they happen.
- Causal analysis looks for the factors that drive losses, such as trade volume, staff turnover or system changes, using statistics on the loss history.
- But correlation does not prove cause, so the mechanism must be understood and tested.
- Education matters too: staff who understand operational risk report losses and near misses, and those reports build the data and the culture on which everything else rests.
- Two tools then help manage the risks.
- Risk and control self-assessment, or R C S A, has business managers identify and score their own risks and controls.
- Key risk indicators, or K R I s, are measurable warning metrics.

## 22 R C S A against K R I

*Two tools that complement each other*

| | RCSA \| KRI |
|---|---|
| Who | Managers \| Tracked metrics |
| Strength | Breadth \| Timely warning |
| Weakness | Subjective snapshot \| Needs thresholds |
| Example | Likelihood, impact \| Failed trades |

**IN THE LESSON**

- Compare the two tools.
- In an R C S A, managers of each business unit identify their risks and controls, and score them, typically by likelihood and impact.
- It gives breadth, covering every part of the business, but it is subjective and only a snapshot in time.
- A key risk indicator is a number tracked regularly, such as failed trades, staff turnover, system downtime or overtime hours, with thresholds that trigger escalation when crossed.
- A good indicator leads losses rather than reflecting them: rising overtime today may warn of errors next month.
- The two complement each other.
- Self-assessment gives the breadth, and indicators give timely warning, so a bank wants both.

## 23 Allocating capital to business units

**Allocation = stand-alone × bank capital ÷ sum of all stand-alone capitals**

| | |
|---|---|
| **Stand-alone capital** | The unit's own figure |
| **Bank capital** | Diversified, bank-wide figure |
| **Sum of stand-alone** | Adds up to more than the bank |

**IN THE LESSON**

- Once the bank has a capital figure, it must be shared among the business units, so that each manager faces a charge for the risk they create and has a reason to reduce it.
- Allocating by gross income is simple but ignores how well a unit controls its risks, so improvement earns no reward.
- Better schemes use the unit's own loss experience and scorecards from R C S A and K R I results, so that stronger controls lower the charge.
- The units' stand-alone capital usually adds up to more than the bank figure, because of diversification.
- So each unit is scaled pro rata: its stand-alone capital times the bank capital, divided by the sum of the stand-alone figures.

## 24 Example 8: scaling three units

*Stand-alone: retail 60, trading 80, corporate 40. Bank total 144*

| | | |
|---|---|---|
| **Scaling factor** | 144 ÷ 180 | **0.8** |
| **Retail** | 60 × 0.8 | **48** |
| **Trading** | 80 × 0.8 | **64** |
| **Corporate** | 40 × 0.8 | **32** |

IN THE LESSON

- Three units have stand-alone operational risk capital of sixty, eighty and forty million, retail, trading and corporate.
- They add to one hundred and eighty million.
- But the bank, allowing for diversification, needs only one hundred and forty four million.
- The scaling factor is one hundred and forty four over one hundred and eighty, which is point eight.
- Retail gets sixty times point eight, forty eight million.
- Trading gets eighty times point eight, sixty four million.
- Corporate gets forty times point eight, thirty two million.
- Check: forty eight plus sixty four plus thirty two is one hundred and forty four, the bank total.
- One caution: allocating purely on reported losses tempts units to hide losses and near misses.

## 25 The power law

- P(loss > x) = K × x^(-α)
- Log-log plot is a straight line with slope -α
- Smaller α means a fatter tail
- Mean finite only if α > 1, variance only if α > 2

IN THE LESSON

- Back to the heavy tail.
- Operational losses are well described by a power law.
- The probability that a loss exceeds some size x is a constant K times x raised to the power of minus alpha.
- Alpha is called the tail index.
- Take logs of both sides and you get a straight line on a log log graph, with slope minus alpha.
- So alpha can be found from two points on the line.
- A smaller alpha means a fatter tail.
- In the Pareto form, the mean loss exists only if alpha is above one, and the variance only if alpha is above two.
- A very fat tail therefore means that averages themselves can be unreliable, and a model must respect that.

## 26 The power law formulas

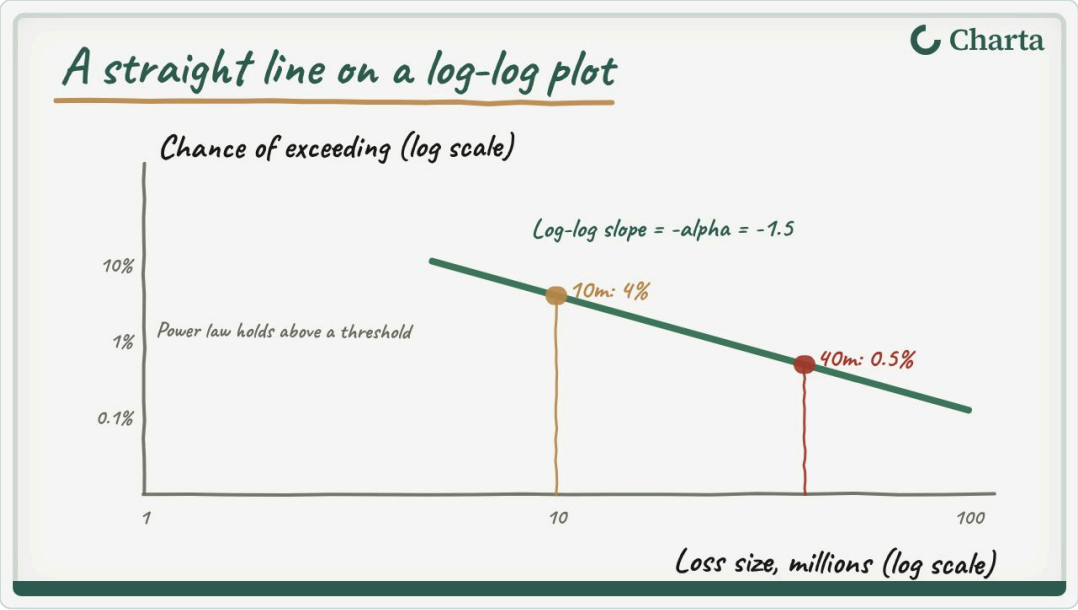
**P(loss > x) = K × x^(-α)**
**α = ln(p₁ ÷ p₂) ÷ ln(x₂ ÷ x₁)**

| | |
|---|---|
| **Scale, K** | Scale constant |
| **Tail index, alpha** | Tail index; slope on log-log plot |
| **Loss size, x** | Loss size |

IN THE LESSON

- Here are the formulas.
- The first line gives the probability that a loss exceeds x.
- The second gives alpha from two points: the natural log of the first probability over the second, divided by the natural log of the second size over the first.
- Notice the sizes are inverted compared with the probabilities, which keeps alpha positive.
- To get a one year loss, combine the power law with a frequency.
- If lambda losses exceed a base size per year, losses above x occur at lambda times x over the base, to the power of minus alpha, per year.
- Set this rate to about one in a thousand and solve for x: that gives the ninety nine point nine per cent one year loss.
- Typical errors are using plus alpha, or forgetting the frequency.

## 27 A straight line on a log-log plot




IN THE LESSON

- Here is the power law on a log log graph.
- Along the bottom is the size of the loss in millions of dollars, on a log scale, from one to one hundred.
- Up the side is the probability of exceeding that size, also on a log scale.
- The result is a straight line sloping down.
- Two points from our example sit on it: the chance of a loss above ten million is four per cent, and the chance above forty million is half of one per cent.
- The slope of the line is minus alpha, here minus one point five.
- A flatter line would mean a smaller alpha and a fatter tail, with big losses much more likely than the bank might hope.

## 28 Example 9: tail index and the one year loss

*P(>10m) = 4%, P(>40m) = 0.5%. 3 losses a year above 10m*

| | | |
|---|---|---|
| **Alpha** | ln 8 ÷ ln 4 | **1.5** |
| **Rate above x** | 3 × (x ÷ 10)^-1.5 | **= 0.001** |
| **Solve** | x = 10 × 3000^(2/3) | **x ≈ 2,080m** |
| **Read** | 99.9% one-year loss | **≈ $2.1bn** |

IN THE LESSON

- First the tail index.
- The chance of a loss above ten million is four per cent, and above forty million it is half of one per cent.
- The ratio of the probabilities is eight; the ratio of the sizes is four.
- So alpha is the natural log of eight over the natural log of four, which is exactly one point five.
- It is above one but below two: the mean is finite, the variance is not.
- Now the frequency.
- Suppose three losses a year exceed ten million.
- The rate of losses above x is three times x over ten, to the power of minus one point five.
- Set it to point zero zero one and solve: x is ten times three thousand to the power of two thirds, about two thousand and eighty million.

## 29 Insurance: moral hazard and adverse selection

- Moral hazard: after buying, less care
- Adverse selection: before buying, riskiest buy most
- Fixes: deductibles, coinsurance, limits, exclusions
- Fixes: disclosure, audits, risk-based pricing

IN THE LESSON

- Insurance can move operational losses off a bank's own books, but it creates two classic problems for the insurer.
- Moral hazard comes after the policy is bought.
- A protected firm has less reason to keep its controls tight, so losses rise.
- Insurers respond with deductibles, coinsurance, where the firm bears a share of each loss, policy limits and exclusions, keeping some skin in the game.
- Adverse selection comes before purchase.
- Firms that know their controls are weak are the most eager buyers, and the insurer cannot see this.
- If it charges one average premium, careful firms find it dear and leave, the pool worsens, and premiums rise.
- Insurers answer with disclosure, audits and risk based pricing.

## 30 Example 10: a policy and a pool

*Loss 5m. Deductible 1m, insurer pays 80%, limit 3m*

| | | |
|---|---|---|
| **After deductible** | 5m - 1m | **4m** |
| **Insurer, capped** | min(80% × 4m, 3m) | **3m** |
| **Firm bears** | 5m - 3m | **2m** |
| **Pool: average, then true** | (1m + 5m) ÷ 2, then 5m | **3m, then 5m** |

IN THE LESSON

- First, how policy terms curb moral hazard: a one million deductible, the insurer paying eighty per cent above it, and a three million limit.
- A loss of five million arrives.
- Take off the deductible: four million remains.
- The insurer covers eighty per cent, three point two million, but the limit caps it at three.
- So the insurer pays three, and the firm bears two.
- Now adverse selection.
- Half the firms expect losses of one million, half expect five million, so the fair average premium is three million.
- Low risk firms see three for a one million risk and leave.
- Only the five million firms remain, and the insurer must charge five.
- Hidden action after purchase is moral hazard; hidden information before it is adverse selection.

## 31 Recap: operational risk in one table

*Operational risk: the key facts*

| Topic | Key fact |
|---|---|
| Types | Seven; classify by cause |
| Old capital | BIA 15%; TSA betas; AMA model |
| SMA | BIC × ILM; marginal 12/15/18% |
| Loss model | Poisson count + lognormal severity |
| Tail, insurance | Power law; moral hazard, adverse selection |

IN THE LESSON

- Let us pull the lesson together.
- Operational risk is loss from failed processes, people, systems or external events, and Basel sorts it into seven types, classified by cause.
- Capital was set by three Basel Two methods, then replaced by the standardized measurement approach, which multiplies a size charge by a loss multiplier.
- Banks model losses by simulating frequency and severity, with care about biased data, and scenarios fill the gaps.
- Controls rest on education, self-assessment and key risk indicators, and capital is allocated by pro rata scaling.
- The power law describes the fat tail.
- And insurance brings moral hazard, which is after purchase, and adverse selection, which is before it.

## 32 Try three quick questions

1. Gross income 300, 200, -50 (m). Basic indicator capital?
   **Answer:** Positive years only: (300 + 200) ÷ 2 = 250; × 15% = 37.5m.
2. A loss happens once in 20 years. Poisson mean per year?
   **Answer:** 1 ÷ 20 = 0.05 events a year.
3. Stand-alone 50, sum 250, bank 200. Allocation?
   **Answer:** 50 × 200 ÷ 250 = 40m.

**IN THE LESSON**

- Now try three quick questions.
- Pause the video after each one, think, then listen.
- Question one.
- A bank had gross income of three hundred million, two hundred million and minus fifty million over three years.
- What is the basic indicator capital?
- Use the positive years only: three hundred plus two hundred, over two, is two hundred and fifty, and fifteen per cent of that is thirty seven point five million.
- Question two.
- A loss happens once in twenty years.
- What is the Poisson mean per year?
- One over twenty, which is point zero five events a year.
- Question three.
- A unit has stand alone capital of fifty million, the sum of stand alone figures is two hundred and fifty, and the bank figure is two hundred.
- What is its allocation?
- Fifty times two hundred over two hundred and fifty, which is forty million.