CHARTA LEARNING · LESSON NOTES

The building blocks of risk management

FRM Part I

1 A story: the bakery and the loan

- A baker borrows money for a bigger oven
- Outcome unknown: a boom or a rival across the street
- The bank lends to hundreds of bakers like her
- Taking uncertainty on purpose, then controlling it

IN THE LESSON

- Picture a baker who borrows money to buy a bigger oven.
- If the neighbourhood loves her bread, she earns far more than before.
- If a rival opens across the street, she may struggle to repay.
- Nobody knows which will happen, and that not knowing is what risk really is.
- Notice that she accepted it on purpose, because the bigger oven might pay.
- Now picture her bank.
- It lends to hundreds of bakers like her, expecting some to fail and most to succeed.
- The bank does not hide from that uncertainty.
- Its whole business is to take it, price it and keep it under control.

2 What is risk?

- Risk = uncertainty about outcomes that matter
- The spread of possible results around the expected one
- Includes better outcomes as well as worse ones
- Wider spread = more risk

IN THE LESSON

- Here is the definition to learn.
- Risk is uncertainty about outcomes that matter to an organization.
- Think of a spread of possible results around the result you expect.
- The wider the spread, the greater the risk.
- Pay close attention to one point: the spread includes better outcomes as well as worse ones.
- A share that might jump twenty per cent is just as uncertain as one that might fall twenty per cent.
- In everyday speech people say risk when they mean danger, but in the F R M the word is wider than that.
- It describes the whole fan of results, good and bad, that could happen.

3 Risk is not the same as loss

- A certain loss is a cost, not a risk
- A realized loss is the result of risk, not the risk itself
- Risk lives before the event, in the uncertainty
- Risk = the spread; loss = one outcome drawn from it

IN THE LESSON

- Now a distinction that matters.
- Risk is not loss.
- Suppose a firm knows for certain that it will lose a fixed amount every year.
- That is not risk, because there is no uncertainty.
- It is simply a cost, and the firm can budget for it like rent.
- And when a loss does happen, say a borrower fails to repay, that realized loss is the result of risk.
- It is not the risk itself.
- Risk lives before the event, in the uncertainty.
- Loss is what is left after the dice have landed.
- Keep these three words apart: risk is the uncertainty, a cost is a certain loss, and a realized loss is one outcome.

4 Risk taking versus risk management

Two different jobs inside a financial firm

Role	What it does
Risk taking	Accepts exposure for reward
Risk management	Identify, measure, control, monitor
Making a new loan	Adds exposure
Setting a loan limit	Constrains exposure

IN THE LESSON

- These two ideas are easy to mix up, so set them side by side.
- Risk taking is the deliberate acceptance of exposure in return for reward.
- For a bank or a fund, it is the business itself: a bank earns by lending, and a fund earns by holding assets that might rise or fall.
- Risk management is different.
- It is the process of identifying, measuring, controlling and monitoring those exposures.
- Taking risk adds exposure.
- Managing risk keeps the exposure understood, intended and within limits.
- A good test for any action is this: does it add exposure, or does it constrain it?
- Lending adds.
- A limit on lending constrains.

5 The purpose of risk management

- Risks taken are understood
- Risks taken are intended, not accidental
- Risks stay within the stated risk appetite
- Risks are adequately compensated by reward
- The aim is NOT to eliminate risk

IN THE LESSON

- What is risk management for?
- Its purpose is that the risks a firm takes are four things, which the board lists in turn.
- They are understood, so that the firm knows what it holds.
- They are intended, so that nothing was taken by accident.
- They are within the stated appetite, meaning the amount of risk the board has said it is willing to bear.
- And they are adequately compensated, so that the reward is worth the danger.
- Notice what is missing from the list: the word eliminate.
- A firm that removed all risk would remove all return, and with it the reason to exist.
- Risk management is about choosing risk well, not about running away from it.

6 Hedging changes risk, it does not erase it

- Example: an exporter locks in an exchange rate
- Removes rate risk, adds risk of the other side failing
- And gives up the gain if the rate moves her way
- A hedge turns a risk into a smaller or different one

IN THE LESSON

- A common belief is that hedging makes risk disappear.
- Usually it does not.
- Suppose an exporter expects to be paid in a foreign currency in three months and locks in today's exchange rate with a contract.
- She has removed the exchange rate uncertainty.
- But she has taken on new risks.
- The other side of the contract might fail to honour it, and if the foreign currency strengthens she gives up a gain she would have enjoyed.
- So a hedge usually transforms a big risk into a smaller or different one.
- Remember that phrase, because questions often test whether you can see what risk remains after a hedge.

7 Enterprise risk management

- ERM: one firm-wide, integrated framework
- Identify, assess, respond to ALL material risks together
- Linked to strategy and risk appetite
- Contrast: each unit managing its own risks in a silo

IN THE LESSON

- Now to the second theme: how risk is gauged and handled.
- The first idea is enterprise risk management, or E R M.
- It is a firm wide, integrated framework.
- It identifies, assesses and responds to all the material risks together, and it links them to the firm's strategy and risk appetite.
- Compare that with the old way, where each unit managed its own risks in a silo.
- In the silo model the credit team watches credit, the trading desk watches trading, and no one sees how they pile up on the same customer or the same event.
- E R M brings the pieces into a single view, so that the board can decide what the firm as a whole should bear.

8 Qualitative techniques

- Self-assessments, risk registers, heat maps
- Scenario workshops and expert judgement
- Good for risks hard to measure or with no history
- Weakness: subjective and coarse

IN THE LESSON

- Inside E R M, firms use two families of tools.
- The first is qualitative.
- Examples are risk and control self assessments, where managers rate their own areas; risk registers, which list each risk and who owns it; heat maps, which plot likelihood against impact; scenario workshops, where experts imagine how things could go wrong; and plain expert judgement.
- These tools suit risks that are hard to measure or that have no loss history, such as a new type of cyber attack.
- Their weakness is that they are subjective and coarse.
- Two experts can score the same risk differently, and a heat map has only a few boxes.

9 Quantitative measures

- Standard deviation, sensitivities, value at risk
- Expected shortfall and stress tests
- Strength: comparable numbers
- Weakness: models, past data, false precision

IN THE LESSON

- The second family is quantitative.
- It includes standard deviation, which measures the spread of returns; sensitivities, which show how much a value moves when a market factor moves; value at risk, which is the loss that should not be exceeded at a stated confidence level over a stated horizon; expected shortfall, which is the average loss in the cases beyond it; and stress tests, which ask what happens in an extreme but plausible scenario.
- The strength is that these give comparable numbers.
- You can put two desks side by side and compare them.
- The weakness is that the numbers rest on models and on past data, and a tidy number can create false precision.
- A figure with many decimal places can feel exact while hiding the fact that the future may not resemble the past.

10 Which tool fits which risk?

Matching the tool to the risk

Situation	Best tool or approach
New risk, no loss history	Scenario workshop
Trading book with price data	Sensitivities, VaR
Rare, extreme event	Stress test
Likelihood versus impact view	Heat map
Whole firm, one view	ERM (the framework)

IN THE LESSON

- Good practice combines both families, and good practice turns on the judgement of which tool fits which risk.
- Look at the table.
- A new, unusual risk with no loss history, such as a brand new product, calls for scenario workshops and expert judgement.
- A trading book with years of price data suits value at risk and sensitivities.
- A rare extreme event is best explored with stress tests, because history may contain no example.
- To see likelihood against impact on one page, use a heat map.
- And for the firm as a whole, E R M ties the results together.
- Notice the pattern: use numbers where data is rich and the world is stable, use judgement where it is thin, and use both to check each other.

11 Expected loss: the cost of doing business

- Expected loss (EL): the average loss anticipated
- A cost of doing business, like shop theft
- Priced into spreads
- Covered by provisions

IN THE LESSON

- Next we split losses into two kinds.
- Expected loss is the average loss an institution anticipates over a given horizon.
- Think of a shop that knows some goods will be stolen each year.
- The shop does not call that a surprise.
- It builds the average theft into its prices and sets aside money for it.
- A bank does exactly the same.
- Expected loss is a cost of doing business.
- It is priced into the interest rate, the spread, that a bank charges, and it is covered by provisions, which are funds set aside for the losses that the bank expects to occur.

12 Expected loss for credit

EL = PD × LGD × EAD

Probability of default, PD	Chance the borrower fails to pay
Loss given default, LGD	Share lost if it does
Exposure at default, EAD	Amount owed when it fails
Expected loss, EL	The average loss expected

IN THE LESSON

- For credit, expected loss has a simple formula with three ingredients multiplied together.
- P D is the probability of default: how likely it is that the borrower fails to pay.
- L G D is the loss given default: the share of the exposure that is lost if the borrower does default, after anything recovered.
- E A D is the exposure at default: how much is owed at the moment of failure.
- Multiply the three and you get the average loss.
- Why multiply?
- Because a loss needs all three things at once: the borrower must default, you must lose part of what is owed, and there must be an amount at stake.

13 Example 1: expected loss on one loan

Illustration: a $5 million loan over one year

Probability of default, PD	chance of failing	**2%**
Loss given default, LGD	100% − 60% recovered	**40%**
Exposure at default, EAD	amount owed	**$5,000,000**
Expected loss	2% × 40% × 5,000,000	**$40,000**

EL is an average, not what happens in any one year

IN THE LESSON

- A bank lends five million dollars to a company.
- These are illustration numbers.
- The bank estimates a two per cent chance of default over the year, so P D is two per cent.
- If the company defaults the bank expects to recover sixty per cent of the money, so it loses forty per cent, and L G D is forty per cent.
- The exposure at default is five million.
- Multiply: zero point zero two times zero point four times five million is forty thousand dollars.
- That is the expected loss.
- It does not mean the bank loses forty thousand.
- It will either lose nothing or lose two million.
- Forty thousand is the average.

14 Unexpected loss: the variability around the average

- UL = variability of loss around the average
- Measured by std deviation, or high-confidence loss − EL
- Cannot be priced away, so capital is held against it
- EL: pricing and provisions. UL: capital

IN THE LESSON

- Expected loss is only half the story.
- Unexpected loss is the variability of loss around that average.
- It is often measured as the standard deviation of losses.
- Another common measure takes the loss at a high confidence level, say the loss that is exceeded only one year in a hundred, and subtracts the expected loss.
- Why does it matter?
- Because it cannot be priced away.
- A bank cannot charge every borrower enough to cover the worst year, or no one would borrow.
- So the bank holds capital, a cushion of its own money, against unexpected loss.
- Expected loss gets provisions and pricing.
- Unexpected loss gets capital.

15 Example 2: the card issuer in a recession

Illustration: $100 million of card balances

Normal-year write-offs	4% × 100,000,000	**$4 million**
Recession-year write-offs	9% × 100,000,000	**$9 million**
Part that was expected	the normal-year level	**$4 million**
Unexpected part this year	9 − 4	**$5 million**

Not the whole $9 million: only the excess over EL

IN THE LESSON

- A card issuer has one hundred million dollars of balances.
- In a normal year it writes off four per cent, which is four million.
- So its expected loss is four per cent.
- Now a severe recession arrives and write offs jump to nine per cent, or nine million dollars.
- How much of that is unexpected?
- Only the extra: nine minus four is five million.
- The first four million was already expected, priced and provided for.
- Here is the warning from the course: do not call the whole bad year loss unexpected.
- The expected part is still expected.
- Only the excess over it is the unexpected loss that capital must absorb.

16 Example 3: adding up expected and unexpected loss

Illustration: two loans, loss measured at a high confidence level

Loan 1 EL plus loan 2 EL	40,000 + 60,000	**$100,000**
Portfolio loss at high confidence	assumed	**$400,000**
Unexpected loss	400,000 − 100,000	**$300,000**
Sum of stand-alone ULs	not the same as portfolio UL	**usually larger**

EL adds up; UL depends on how the loans move together

IN THE LESSON

- Here is an important difference.
- Expected loss is roughly additive across positions.
- Suppose loan one has an expected loss of forty thousand dollars and loan two has sixty thousand.
- The portfolio expected loss is simply one hundred thousand.
- Unexpected loss is not additive, because it depends on how the loans move together.
- Suppose the portfolio loss at the high confidence level is four hundred thousand.
- Then unexpected loss is four hundred thousand minus the expected one hundred thousand, which is three hundred thousand.
- If the two loans rarely fail together, the portfolio unexpected loss is usually smaller than the two stand alone unexpected losses added.
- That is the benefit of diversification.

17 Risk and reward go together

- Higher expected return for bearing more risk
- CAPM: only systematic risk is priced
- Systematic risk cannot be diversified away
- Idiosyncratic risk can be diversified: no premium

IN THE LESSON

- Why does anyone accept risk?
- Because of reward.
- Investors and firms expect higher expected returns for bearing more risk.
- If a safe government bill pays three per cent, no one would hold a risky share unless it promised more.
- But the full story has a subtle twist.
- Under the capital asset pricing model, and in a frictionless efficient market generally, only systematic risk is priced.
- That is the part linked to the whole market, which cannot be diversified away.
- Idiosyncratic risk, the bad luck of a single company, can be diversified away by holding many shares, so it earns no premium.
- You are paid for risk you cannot escape, not risk you could have avoided.

18 Judging return per unit of risk: the Sharpe ratio

Sharpe ratio = (portfolio return − risk-free rate) ÷ σ

Portfolio return	What the portfolio earned
Risk-free rate	What a safe asset earns
Standard deviation, sigma	Volatility of the excess return
Sharpe ratio	Reward per unit of risk

IN THE LESSON

- Because return comes with risk, we should judge performance per unit of risk.
- One common measure is the Sharpe ratio.
- Take the portfolio return and subtract the risk free rate, the return on a safe asset.
- That gives the excess return, the reward for taking risk.
- Then divide by sigma, the standard deviation of that excess return, which measures how bumpy the ride was.
- A higher Sharpe ratio means more reward for each unit of risk.
- A fund that earned a large return by taking a huge risk may look worse than a modest fund that took very little.
- Other risk adjusted capital measures serve the same purpose.

19 Example 4: which fund was better?

Illustration: risk-free rate 3%

Fund A excess return	9% − 3%	**6%**
Fund A Sharpe ratio	6 ÷ 10	**0.6**
Fund B excess return	12% − 3%	**9%**
Fund B Sharpe ratio	9 ÷ 18	**0.5**
Better per unit of risk	0.6 is more than 0.5	**Fund A**

IN THE LESSON

- Compare two funds over a year, using illustration numbers.
- The risk free rate is three per cent.
- Fund A returned nine per cent with a standard deviation of ten per cent.
- Fund B returned twelve per cent with a standard deviation of eighteen per cent.
- Fund B has the higher return, but look at the Sharpe ratio.
- For A, nine minus three is six, divided by ten, gives zero point six.
- For B, twelve minus three is nine, divided by eighteen, gives zero point five.
- Fund A delivered more reward for each unit of risk.
- Judged on raw return alone, we would have picked the wrong fund.

20 When interests conflict: the agency problem

- Trader paid a share of profits, not charged for losses
- Payoff like a call option: extra risk raises expected pay
- Shareholders of a leveraged firm gain at lenders' expense
- Risk managers paid by the business they police

IN THE LESSON

- The link between risk and reward can be distorted when interests conflict.
- This is the agency problem: the people making risk decisions are not the people who bear the results.
- Take a trader paid a share of profits but never charged for losses.
- Her payoff is like a call option: she gains if things go well and loses nothing if they go badly.
- So extra risk raises her expected pay.
- Or take shareholders of a heavily borrowed firm.
- They can gain at the expense of lenders by shifting to riskier assets, because lenders bear much of the downside.
- And risk managers who report to, or are paid by, the business they police lose their independence.

21 Example 5: a bonus that rewards gambling

Illustration: bonus = 10% of profit, no charge for losses

Safe: firm's expected profit	sure gain	**$10 million**
Safe: trader's bonus	10% × 10	**$1 million**
Gamble: firm's expected profit	½ × 100 − ½ × 100	**$0**
Gamble: trader's expected bonus	½ × 10 + ½ × 0	**$5 million**
Whose interest is served?	trader prefers the gamble	**Not the firm**

IN THE LESSON

- Let us see the distortion in numbers.
- A trader earns a bonus of ten per cent of profit, and nothing is deducted for losses.
- Strategy one is safe and earns a sure profit of ten million for the firm, so the bonus is one million.
- Strategy two is a coin flip: a hundred million profit if heads, a hundred million loss if tails.
- For the firm, the expected result is zero, which is worse than ten million.
- But the trader earns a ten million bonus on heads and nothing on tails, so her expected bonus is five million.
- She prefers the gamble, and the firm bears the downside.
- Her interests and the firm's have parted.

22 Remedies for distorted incentives

- Independent reporting lines for risk managers
- Deferred pay that can be clawed back
- Incentives tied to risk-adjusted results
- Goal: the decision maker shares the downside

IN THE LESSON

- How do firms repair this?
- There are three main remedies.
- First, independent reporting lines: the risk function should report to the board or a chief risk officer, not to the business unit it checks, and its pay should not depend on that unit's profit.
- Second, deferred and clawbackable pay: if part of a bonus is held back for years, and can be taken back if losses appear later, then the trader shares some of the downside.
- Third, incentives tied to risk adjusted results, such as the Sharpe ratio or profit after a charge for capital, so that extra risk is not rewarded unless it earns enough.
- Together they bring the trader's interests closer to the firm's.

23 The main classes of risk

Classify by the root cause

Class of risk	Root cause
Market	Changes in prices and rates
Credit	Counterparty fails to pay
Liquidity	Cannot sell, or cannot raise cash
Operational	Failed processes, people, systems
Business / reputational	Lost revenue, lost trust

IN THE LESSON

- Now we sort risks into classes.
- The classes differ in where they come from, and that root cause is what you use to classify them.
- Market risk is loss from changes in prices and rates, such as equities, interest rates, currencies and commodities.
- Credit risk is loss from a borrower or counterparty failing to pay, and it includes a downgrade.
- Liquidity risk has two forms, which we meet shortly.
- Operational risk is loss from failed processes, people, systems or outside events.
- Business or strategic risk is the erosion of revenue by competition or regulation.
- Reputational risk is the loss of trust.
- Learn the root cause of each, and classify by it.

24 Market risk and credit risk

- Market risk: prices and rates move (equities, FX, rates)
- Credit risk: borrower or counterparty fails to pay
- A downgrade is also credit risk
- Root cause: a price moving versus a party failing

IN THE LESSON

- Take market risk first.
- A fund holds shares and bonds.
- If share prices fall, or interest rates rise and bond prices drop, or a currency weakens, the fund loses value, even though nobody failed to pay anything.
- The cause is simply the movement of a price or a rate.
- Credit risk is different, because it comes from a person or company.
- A borrower may miss payments, or a trading partner may fail to deliver what it promised.
- A downgrade counts too, because a borrower that becomes less creditworthy makes its loan worth less, even before it defaults.
- So market risk comes from prices moving, and credit risk comes from someone failing.

25 Liquidity risk has two forms

- Market liquidity: cannot sell without a big discount
- Funding liquidity: cannot raise cash when due
- Rich in assets can still fail for lack of cash
- The two feed each other

IN THE LESSON

- Liquidity risk is easy to muddle, because it has two forms.
- Market liquidity risk is being unable to sell an asset without a large price concession.
- A house is hard to sell in a week, unless you accept a much lower price.
- Funding liquidity risk is different: it is being unable to raise cash to meet obligations when they fall due.
- A firm can own plenty of valuable assets and still fail, if it cannot find cash on the day a payment is demanded.
- The two feed each other.
- A firm short of cash must sell assets, and if those assets are hard to sell, it sells them cheaply, which makes the cash problem worse.

26 Operational, business and reputational risk

- Operational: failed processes, people, systems, events
- Basel: includes legal risk, not strategic or reputational
- Business risk: revenue eroded by competition or regulation
- Funding failure: days. Business risk: slowly
- Effects: MTM losses, write-offs, fines, lost customers

IN THE LESSON

- Operational risk is loss from inadequate or failed processes, people, systems or external events.
- A data entry error, a rogue employee, a computer outage and a flood all belong here.
- Under the Basel rules, operational risk includes legal risk but excludes strategic and reputational risk.
- Business or strategic risk is slower.
- It is the erosion of revenue by competition or regulation, such as a rival with a better product.
- Reputational risk is the loss of trust, often following another failure.
- Compare their speed too.
- Funding liquidity failure can be fatal within days, while business risk acts slowly.
- Speed is part of judging what effect each could have.
- Typical effects differ too: market risk brings mark to market losses, credit risk brings write offs, operational risk brings direct losses and fines, and reputational risk costs customers and funding.

27 Exercise: name the class by its root cause

Scenario to class

Scenario	Class
Bond prices fall as rates rise	Market
Customer misses loan payments	Credit
Payment sent to wrong account	Operational
Cannot roll over borrowing	Funding liquidity
Rival product takes customers	Business

IN THE LESSON

- Let us practise classification.
- Look at the table.
- A bank's trading book loses value because interest rates rose: the cause is a rate moving, so market risk.
- A corporate customer misses its loan payments: the cause is a party failing, so credit risk.
- A clerk sends a payment to the wrong account: a failed process, so operational risk.
- A bank cannot roll over its short term borrowing tomorrow: cash cannot be raised, so funding liquidity risk.
- A competitor's new product steadily takes customers: business risk.
- The method is always the same.
- Ask what started it, ignore how big the loss was, and name the class from the cause.

28 Risks do not act alone

- Higher rates raise credit losses on floating-rate debt
- Falling prices force fire sales: funding to market loss
- Wrong-way risk: exposure peaks when default is likeliest
- In crises correlations rise and diversification shrinks

IN THE LESSON

- So far we have treated each class as if it stood by itself.
- In real life, risks act on one another.
- Rising interest rates can raise credit losses on borrowers whose loans have floating rates, because their payments jump.
- A fall in asset prices can force fire sales, which turn a funding problem into a market loss, as prices drop further.
- Wrong way risk occurs when exposure to a counterparty is largest just when its probability of default is highest.
- Imagine buying protection from a firm whose fortunes are tied to the very event you fear.
- And in crises, correlations between risk factors tend to rise, so the benefit of diversification shrinks just when you need it.

29 Combining two risk figures

Total = $\sqrt{(A^2 + B^2 + 2 \times \rho \times A \times B)}$

Stand-alone risk figure A	Risk figure for one risk class
Stand-alone risk figure B	Risk figure for another class
Correlation, rho	How closely they move together
Combined total	The combined figure

IN THE LESSON

- Suppose a firm has a market risk figure, A, and a credit risk figure, B, and wants one total.
- Simply adding them assumes the two are perfectly correlated, so they always go wrong together, and that overstates the total unless they really are.
- Assuming they are completely unrelated, with zero correlation, understates it when they are in fact positively correlated.
- The right answer depends on the correlation, written rho.
- Square each figure, add twice rho times A times B, and take the square root.
- When rho is one, the formula gives A plus B.
- When rho is zero, it gives the square root of the two squares added.
- For anything in between, the total falls in between.

30 Example 6: totals for different correlations

Illustration: A = 30 (market), B = 40 (credit), in $ millions

Correlation 1 (just add)	30 + 40	**70**
Correlation 0	√(900 + 1,600)	**50**
Correlation 0.5	√(900 + 1,600 + 1,200)	**about 60.8**
Spread of answers	70 − 50	**20**

One assumption moves it by $20 million

IN THE LESSON

- Take a market risk figure of thirty million and a credit risk figure of forty million.
- These are illustration numbers.
- If the two are perfectly correlated, the total is simply seventy.
- If they are completely unrelated, the total is the square root of nine hundred plus sixteen hundred, which is the square root of twenty five hundred, or exactly fifty.
- With a correlation of one half, we add twice one half times thirty times forty, which is twelve hundred.
- The sum is thirty seven hundred, and its square root is about sixty point eight.
- So for correlations between zero and one, the total ranges from fifty to seventy, depending entirely on one assumption.

31 Why a single total is hard to build

- Different models, horizons and confidence levels
- Some risks, like reputation, are hard to quantify
- Adding overstates; assuming zero correlation understates
- Calm-period correlations fail under stress

IN THE LESSON

- Why not just compute one number?
- Several reasons.
- Market, credit and operational risk are measured with different models, over different time horizons, and at different confidence levels.
- A market figure may cover ten days, while a credit figure covers a year, so the two are not on the same footing.
- Some risks, such as reputation, are hard to quantify at all.
- Simply adding stand alone figures overstates the total, and assuming zero correlation understates it.
- And correlations estimated in calm periods are unreliable under stress, since they rise in a crisis.
- So any total should come with its assumptions, and with stress tests as a check.

32 Four mistakes to avoid

- 1. Risk is uncertainty, not loss; includes good outcomes
- 2. Only the excess over EL is unexpected loss
- 3. A hedge changes a risk; a conflict needs independence
- 4. Classify by root cause, not by the size of the loss

IN THE LESSON

- Before we finish, four mistakes to avoid.
- One.
- Calling risk the same as loss.
- Risk is the uncertainty, and it includes good outcomes.
- Two.
- Calling the whole bad year loss unexpected.
- The expected part is still expected, and only the excess is unexpected.
- Three.
- Thinking a hedge makes risk vanish.
- A hedge changes the risk, and a conflict of interest is not fixed by hiding it, it needs independent reporting lines.
- Four.
- Classifying by the size or the headline of the loss rather than by the root cause.
- A bank run may feel like credit risk, but if the problem is raising cash, it is funding liquidity.
- Keep these four in mind.

33 Recap: the building blocks

- Risk = uncertainty; manage it, do not eliminate it
- ERM: qualitative and quantitative tools together
- EL = PD × LGD × EAD; capital covers UL
- Sharpe = excess return ÷ σ; classify by root cause
- Total = $\sqrt{(A^2 + B^2 + 2\rho AB)}$: depends on correlation

IN THE LESSON

- Let us recap.
- Risk is uncertainty about outcomes that matter.
- Risk taking accepts exposure for reward, and risk management identifies, measures, controls and monitors it, without trying to eliminate it.
- E R M joins every risk into one framework, using qualitative tools for hard to measure risks and quantitative tools for comparable numbers.
- Expected loss is probability of default times loss given default times exposure; unexpected loss is the variability around it, and capital covers it.
- Return should be judged per unit of risk, and conflicts of interest distort risk taking.
- The main classes are market, credit, liquidity, operational, and business and reputational risk.
- And they interact, so adding them needs a correlation.

34 Try three quick questions

1. Normal-year loss 3%; this year 7%. Find EL and the unexpected part.
   Answer: EL = 3%; unexpected = 7% − 3% = 4%.
2. Return 8%, risk-free 2%, σ = 12%. Find the Sharpe ratio.
   Answer: (8 − 2) ÷ 12 = 0.5.
3. A bank cannot raise cash to meet debts when they fall due. Which risk?
   Answer: Funding liquidity risk.

IN THE LESSON

- Now try three quick questions yourself.
- Pause after each one.
- Question one.
- A bank expects to lose three per cent on its card balances in a normal year, but this year it loses seven per cent.
- What is the expected loss and what is the unexpected part?
- The expected loss is three per cent and the unexpected part is the four per cent excess.
- Question two.
- A fund returns eight per cent, the risk free rate is two per cent, and the standard deviation is twelve per cent.
- What is the Sharpe ratio?
- It is zero point five.
- Question three.
- A bank cannot raise cash to pay its debts when due.
- Which class of risk is this?
- Funding liquidity risk.