CHARTA LEARNING · LESSON NOTES

Operational Risk Mitigation

FRM Part II

1 A story: three problems, one budget

- Card fraud: many small losses every year
- Data centre on a flood plain: rare but costly
- A new product planned for next spring
- Failed internal processes, people, systems; external events
- Mitigation: choose a response for each risk, one budget

IN THE LESSON

- Imagine Ashcombe Bank at its annual risk review.
- Three problems are on the table.
- First, card fraud: hundreds of small losses every year, each one a few thousand dollars.
- Second, a data centre that sits on a flood plain.
- A flood is rare, but it would be very expensive.
- Third, a plan to launch a new product next spring.
- All three are operational risks.
- Operational risk is the risk of loss from inadequate or failed internal processes, people and systems, or from external events.
- Risk mitigation is the set of choices a firm makes to deal with that risk.
- The risk committee has one budget and cannot spend it on everything, so it must choose a response for each risk.
- This lesson is about making those choices well.

2 Four responses to a risk

No response removes risk completely

Response	How	At Ashcombe
Avoid	Do not enter, or leave	Drop the new product
Reduce	Better controls, processes	Fraud-detection rules
Transfer	Insurance, outsourcing, indemnity	Flood insurance
Accept	Fund from earnings, provisions, capital	Budget for small frauds

IN THE LESSON

- No response removes risk completely, but a firm has four basic responses to an operational risk exposure.
- It can avoid the risk, by not entering an activity, or by leaving it.
- Ashcombe could simply drop the product it is planning.
- It can reduce the risk, by improving its controls and processes, for example with better fraud detection rules.
- It can transfer the risk, through insurance, outsourcing, or a contractual indemnity, which is a promise by another party to cover certain losses.
- Flood insurance is an example.
- Or it can accept the risk, and fund the expected losses from earnings, provisions or capital, as when Ashcombe budgets for its small fraud losses each year.
- Every response has a cost, and every response leaves something behind.

3 Weighing the responses

Weigh cost against benefit, the residual risk, and the risk appetite

Loss profile	Best-suited responses
Frequent, small losses	Reduce, or accept and budget
Rare, severe losses	Transfer, or avoid
Any choice	Check the residual risk left

IN THE LESSON

- How does a firm choose?
- Each response is weighed on three tests.
- Cost against benefit: does the response save more than it costs?
- Residual risk: what is left once the response is in place?
- And risk appetite: is that residual risk within the amount of risk the board is willing to take?
- The loss profile, meaning how often losses happen and how big they are, points to the answer.
- Frequent, small losses suit reduction, or acceptance with a budget, because they are predictable and controls can pay for themselves.
- Rare, severe losses suit transfer, or avoidance, because one event could exceed what the firm can absorb.
- Whatever the choice, some residual risk remains, and the firm must check that it fits the appetite.

4 Example 1: frequent, small fraud losses

400 frauds a year, avg $5,000, none near $250,000. Upgrade: $0.6m, halves losses

Expected loss a year	400 × $5,000	**$2.0m**
Saving from the upgrade	50% × $2.0m	**$1.0m**
Net benefit of reducing	1.0 − 0.6	**$0.4m**
Residual loss, accepted	funded from earnings	**$1.0m**
Insurance, $250,000 deductible	every loss is inside it	**Pays $0**

IN THE LESSON

- Start with card fraud, a frequent, small loss.
- Ashcombe suffers about four hundred card frauds a year, averaging five thousand dollars each, and none comes anywhere near a quarter of a million dollars.
- So the expected loss is two million dollars a year.
- A control upgrade, better fraud detection, costs six hundred thousand dollars a year and is expected to halve the losses.
- The saving is half of two million, one million dollars.
- The net benefit of reducing the risk is one million minus six hundred thousand, four hundred thousand dollars, so the upgrade pays for itself.
- The remaining one million dollars of losses is the residual risk.
- Ashcombe accepts it and funds it from earnings, as a cost of doing business.
- What about insurance?
- The deductible is the part of each loss the firm keeps.
- A policy with a deductible of two hundred and fifty thousand dollars per loss would pay nothing, because every fraud falls inside it.

5 Example 2: a rare, severe loss

Flood: 2% chance a year, $40m loss. Appetite: no single loss above $15m

Expected loss a year	2% × $40m	**$0.8m**
Worst case against appetite	$40m against $15m	**Outside**
Insurer pays, if flooded	$30m cover above $5m	**$30m**
Ashcombe keeps	5 deductible + 5 above 35	**$10m: inside**
Premium vs expected payout	1.2 against 2% × 30	**1.2 vs 0.6**

Costs more than it returns on average, but cuts the tail; or relocate

IN THE LESSON

- Now the flood.
- Ashcombe judges there is a two per cent chance a year that the data centre floods, causing a loss of forty million dollars.
- The expected loss is only eight hundred thousand dollars a year.
- But the worst case, forty million, is far outside the board's appetite, which allows no single loss above fifteen million.
- An insurer offers cover of thirty million dollars above a deductible of five million.
- If the flood happens, the insurer pays thirty million and Ashcombe keeps ten million, inside its appetite: the five million deductible, plus the five million above the thirty five million where the cover ends.
- The premium is one point two million a year, against an expected payout of two per cent times thirty million, which is six hundred thousand.
- So the insurance costs more than it returns on average, yet it is worth buying, because it cuts the tail, the rare and very large loss.
- Relocating the data centre, which avoids the risk, is the other choice.

6 Transfer moves money, not accountability

- Transfer shifts the money, not the accountability
- Customers and regulators still hold the firm responsible
- New risks: vendor failure, an insurer that does not pay
- Exam: match the response to the loss profile
- Exam: name the residual risk that remains

IN THE LESSON

- One warning about transfer.
- It shifts the money, meaning the financial consequence of a loss, but usually not the accountability.
- If Ashcombe outsources its card processing and the vendor fails, customers and regulators still hold Ashcombe responsible.
- Transfer can also add new risks of its own: the vendor may fail, or the insurer may not pay when the claim arrives.
- For the exam, two skills are tested.
- First, match the response to the loss profile: frequent and small points to reduce or accept, rare and severe to transfer or avoid.
- Second, recognise the residual risk that remains after any response, because none of them brings the risk to zero.

7 Four kinds of internal control

Controls are classified by when and how they act

Kind	When it acts	What it does
Preventive	Before the event	Stops error or fraud
Detective	After the event	Finds the problem
Corrective	Once it is found	Limits damage, restores
Directive	All the time	Guides behaviour

IN THE LESSON

- An internal control is any procedure or system that keeps a process on track.
- Controls are classified by when and how they act.
- A preventive control acts before the event, and stops an error or a fraud from happening at all.
- A detective control acts after the event, and finds the problem.
- A corrective control acts once the problem has been found, and limits the damage and restores the position.
- A directive control works all the time, by guiding behaviour towards the outcome the firm wants.
- Think of a house.
- A lock is preventive, a burglar alarm is detective, the locksmith who repairs a forced door is corrective, and the family rule to lock up at night is directive.

8 Examples of each kind

Any of these can be manual or automated

Kind	Examples
Preventive	Segregation of duties, access, approval limits, hard stops
Detective	Reconciliations, exception reports, audit, surveillance
Corrective	Reverse an entry, restore a backup, compensate
Directive	Policies, procedures, mandates, training

IN THE LESSON

- Here are the standard examples, worth learning.
- Preventive controls include segregation of duties, meaning no one person can both start and complete a sensitive task, restrictions on system access, approval limits, and system hard stops that simply refuse an invalid entry.
- Detective controls include reconciliations, which compare two records that should agree, exception reports, internal audit, and surveillance of trading and communications.
- Corrective controls include reversing an erroneous entry, restoring data from a backup, and compensating a client who was harmed.
- Directive controls include policies, procedures, mandates and training.
- And any of these can be manual, done by a person, or automated, done by a system.

9 Example 3: classify Ashcombe's controls

Ask: what does the control do, and at what moment?

Second approval over $1m	stops it beforehand	**Preventive**
Daily cash reconciliation	finds breaks afterwards	**Detective**
Recall, reverse wrong payment	repairs the damage	**Corrective**
Payments policy and training	guides behaviour	**Directive**
System blocks unknown payee	automated hard stop	**Preventive**

A detective control does not stop a loss: its value depends on speed

IN THE LESSON

- Let us classify five of Ashcombe's payment controls.
- Ask two questions: what does the control do, and at what moment?
- A rule that any payment over one million dollars needs a second approver stops a bad payment beforehand, so it is preventive.
- A daily cash reconciliation finds breaks after they have happened, so it is detective.
- Recalling and reversing a wrong payment repairs the damage, so it is corrective.
- The payments policy and the training on it guide behaviour, so they are directive.
- A system that blocks any payment to a payee not on the approved list is an automated hard stop, preventive again.
- One trap: a detective control does not stop a loss.
- Its value depends on how quickly the finding is acted upon.
- A reconciliation that nobody reads for a month is worth little.

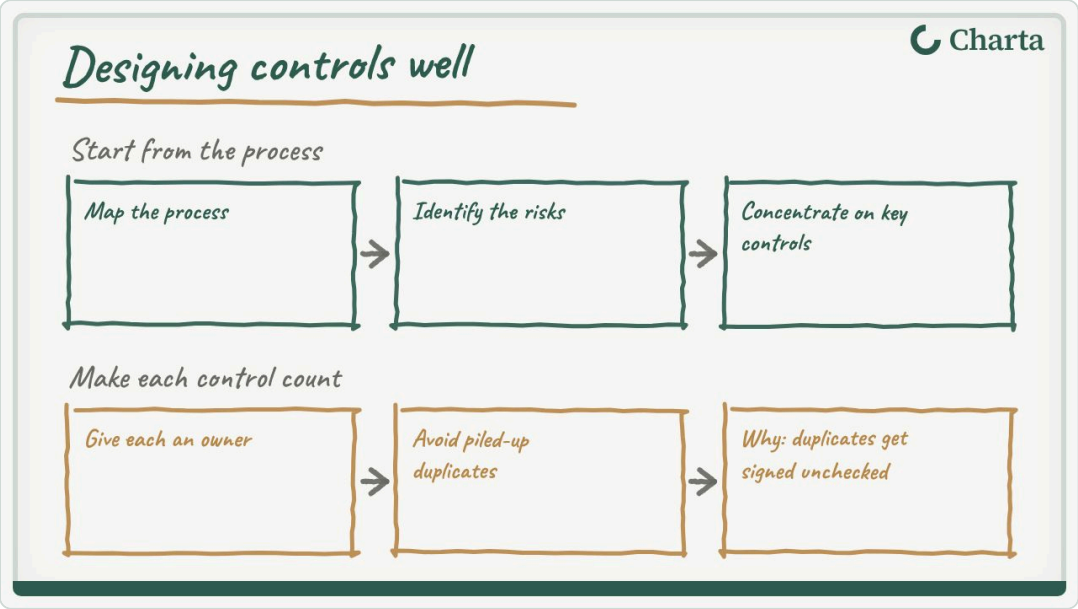
10 Automating controls: gains and dangers

- Gains: less manual error, lower cost, same every time
- Danger: a wrong rule is applied consistently, at speed
- Danger: staff lose the skill to notice failures
- Relies on general IT controls: change, access, backups
- Without them, the control can be altered or disabled

IN THE LESSON

- Many controls can be automated, and automation brings real gains.
- It reduces manual error and cost, and it applies the control the same way every time.
- But there are dangers.
- A wrong rule or a wrong configuration is also applied consistently, and at speed, so one mistake can hit thousands of transactions before anyone notices.
- And staff who no longer do the check by hand can lose the skill to notice when it fails.
- Automated controls also rely on general I T controls.
- These are change management, so nobody alters the system without approval and testing, then access security, and backups.
- Without them, the automated control can be altered or disabled, perhaps without anyone knowing.

11 Designing controls well



IN THE LESSON

- How should controls be designed?
- Good design starts from the process.
- First, map the process, step by step.
- Second, identify the risks at each step.
- Third, concentrate on the key controls, the few that really address the biggest risks.
- Then make each control count.
- Give each control a named owner, who is responsible for performing it.
- And avoid piling on duplicate controls.
- When the same item is checked five times, each checker assumes someone else has looked properly, and duplicates get signed off without checking.
- More controls do not mean more control.
- At Ashcombe, a payment that passed five sign-offs still went to the wrong account, because none of the five reviewers compared it with the invoice.

12 Testing: design versus operating effectiveness

Testing is independent of the control owner

Test	Design effectiveness	Operating effectiveness
Question	Would it work as described?	Did it work over the period?
How	Walkthrough, design review	Sampling and re-performance
If it fails	Redesign the control	Fix how it is performed
Then	Remediate and retest	Remediate and retest

IN THE LESSON

- Testing should be independent of the control owner, who cannot mark their own homework.
- It has two parts, and the exam often asks you to tell them apart.
- Design effectiveness asks whether the control would work if it were performed as described.
- Operating effectiveness asks whether it actually worked, consistently, over the period.
- Design is tested by a walkthrough, following one transaction through, and by reviewing the design against the risk.
- Operating effectiveness is usually tested by sampling transactions and re-performing the control.
- If the design fails, the control must be redesigned.
- If the operation fails, the way it is performed must be fixed.
- In both cases the deficiency is remediated and then retested.

13 Example 4: testing a payment approval

Control: every payment over $1m needs a second approval

Design: would it work as written?	walkthrough of one payment	**Yes**
Operating: sample size	payments over $1m, full year	**25**
Exceptions found	no second approval	**2**
Exception rate	2 ÷ 25	**8%**
Verdict: Ashcombe tolerates zero	remediate, then retest	**Not effective**

Design flaw instead: approval only above $5m when the risk starts at $1m

IN THE LESSON

- Ashcombe's risk function tests the rule that every payment over one million dollars needs a second approval.
- First, design.
- A walkthrough of one payment shows the rule would stop a bad payment if performed as written, so the design is effective.
- Second, operation.
- The tester samples twenty five payments over one million from the whole year and re-performs the check.
- Two had no second approval.
- The exception rate is two divided by twenty five, eight per cent.
- Ashcombe's tolerance for this key control is zero exceptions, so the control is not operating effectively.
- It must be remediated, then retested.
- Contrast a design failure: had the rule required approval only above five million dollars, when the risk starts at one million, it would fail even if performed perfectly.

14 Better processes, fewer errors

- Standardise and simplify processes
- Fewer hand-offs and less re-keying of data
- Move toward straight-through processing
- Each manual step adds a chance of error

IN THE LESSON

- Process quality and human error are tackled together, because most errors happen inside processes.
- The first set of techniques changes the process itself.
- Standardise and simplify it, so every case is handled the same way.
- Reduce hand-offs, where work passes from one person or team to another, and reduce re-keying, where someone types again data that already exists in another system.
- And move toward straight-through processing, in which a transaction flows from start to finish without anyone touching it by hand.
- Why does this matter so much?
- Because each manual step adds its own chance of error, and those chances build up across a process, as the next slide shows.

15 Errors build up across steps

All n steps right = (1 − p)^{n}
At least one error = 1 − (1 − p)^{n}
Small p: roughly n × p

p: error per step	Chance of error at one step
n: steps	How many independent steps
1 − p: step right	Chance one step is done right
Error somewhere	At least one, anywhere

IN THE LESSON

- Let p be the chance of an error at one step, and n the number of independent steps.
- The chance that one step is done right is one minus p.
- If the steps are independent, the chance that every step is done right is one minus p, raised to the power n.
- So the chance of at least one error, somewhere in the process, is one minus that: one minus, open bracket, one minus p, close bracket, to the power n.
- When p is small, this is close to n times p, so doubling the number of steps roughly doubles the chance of an error.
- That is why cutting steps is such a powerful way to raise quality.

16 Example 5: cutting steps cuts errors

Each manual step: 1% chance of error. Ashcombe books 10,000 trades a month

Manual: 8 steps	1 − 0.99^8	**7.73%**
Trades with an error	10,000 × 7.73%	**773**
Straight-through: 3 steps	1 − 0.99^3	**2.97%**
Trades with an error, after	10,000 × 2.97%	**297**

IN THE LESSON

- Ashcombe's trade booking process has eight manual steps, each with a one per cent chance of error, and the bank books ten thousand trades a month.
- The chance of at least one error in a trade is one minus point nine nine to the power eight, which is seven point seven three per cent.
- On ten thousand trades, that is about seven hundred and seventy three trades with an error.
- Now Ashcombe moves toward straight-through processing, leaving only three manual steps.
- The chance falls to one minus point nine nine to the power three, two point nine seven per cent, or about two hundred and ninety seven trades.
- Nobody worked harder or more carefully.
- The errors fell by more than half simply because there were fewer places for them to happen.

17 Error-proofing and checking

- Error-proofing: mandatory fields, drop-downs, formats
- Maker-checker, or four-eyes, review
- Checklists and clear documentation
- A checker helps only if the review is truly independent

IN THE LESSON

- The second set of techniques catches errors at the point of entry.
- Error-proofing makes wrong input hard: mandatory fields that cannot be left blank, drop-down lists instead of free typing, and format checks, for example that an account number has the right length.
- Next, maker-checker, also called four-eyes review.
- One person, the maker, prepares an item, and a second person, the checker, reviews it before it goes out.
- Checklists and clear documentation make sure steps are not skipped.
- But a checker lowers errors only to the extent that the review is genuinely independent.
- A checker who trusts the maker and ticks the box adds a signature, not a control.

18 Example 6: an independent checker or a rubber stamp

5,000 payments a month; the maker's error rate is 2%

Maker errors	2% × 5,000	**100**
Independent checker	catches 90%: 100 × 10%	**10 get through**
Rubber-stamp checker	catches 10%: 100 × 90%	**90 get through**

Same signature on the file, nine times as many errors get through

IN THE LESSON

- Here is the effect in numbers.
- Ashcombe's makers prepare five thousand payments a month with an error rate of two per cent, so one hundred payments contain an error.
- A genuinely independent checker, who compares each payment with its source, catches ninety per cent of them.
- Only ten per cent of the hundred, ten errors, get through.
- A rubber stamp checker, who signs without really looking, catches only ten per cent.
- Ninety per cent of the hundred, ninety errors, get through.
- Both reviews leave the same signature on the file, yet the second lets nine times as many errors through.
- That is why reviewers ask how a check is done, not merely whether it was signed.

19 People, workload and culture

- Manage workload, staffing and fatigue
- Errors rise at peaks and with staff turnover
- Train; report mistakes, find root causes, do not blame
- Track error and rework rates
- Exam: link each technique to the cause of the error

IN THE LESSON

- The third set of techniques is about people.
- Manage workload, staffing and fatigue, because errors rise at peaks, such as month end, and with staff turnover, when new people are still learning.
- Train staff properly.
- And build a culture in which mistakes are reported and their root causes are analysed, rather than individuals blamed.
- If people are punished for owning up, errors are hidden and the same cause strikes again.
- Track error and rework rates, so management can see whether things are getting better.
- For the exam, link each technique to the cause of the error: re-keying calls for automation, fatigue calls for staffing, and a careless checker calls for independence.

20 Why new things are risky

- The firm is doing something unfamiliar
- Untested: processes, systems, models, legal documents, skills
- Volumes start small, then grow fast
- Commercial pressure pushes timelines
- Risk is highest at launch

IN THE LESSON

- Now to change.
- New products, new business initiatives and mergers or acquisitions all raise operational risk for the same basic reason: the firm is doing something unfamiliar.
- The processes, the systems, the valuation models, the legal documentation and the staff skills are all untested.
- Volumes start small, so problems stay hidden, and then they grow fast, and manual workarounds that coped with ten deals a week break at a hundred.
- Meanwhile commercial pressure pushes timelines, and corners get cut to hit a launch date.
- Put together, risk is highest at launch, exactly when the firm knows least about the new activity.

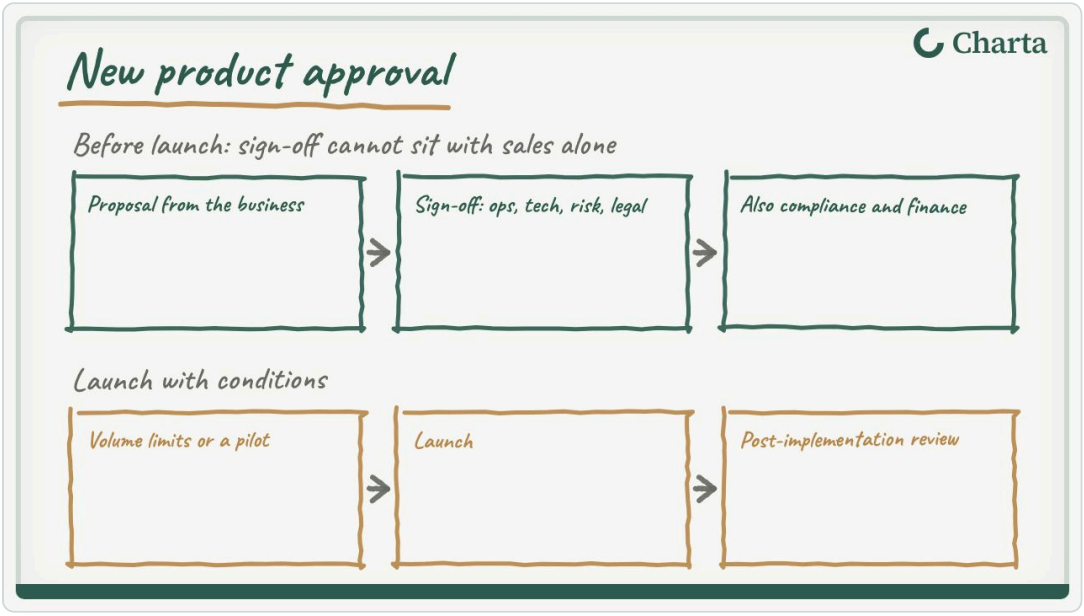
21 Acquisitions add integration risk

- Incompatible systems
- Data migration errors
- Loss of key staff
- Clashing cultures and control standards
- Legacy problems missed in due diligence

IN THE LESSON

- Acquisitions add a further layer, called integration risk.
- Suppose Ashcombe buys Kestrel Finance, a smaller lender.
- The two banks' systems may be incompatible.
- Moving Kestrel's customer records onto Ashcombe's systems can introduce data migration errors, such as balances that are lost or duplicated.
- Key staff, perhaps the only people who understand Kestrel's systems, may leave.
- The two cultures and control standards may clash, so it is unclear whose rules apply.
- And there may be legacy problems, such as products sold to the wrong customers or poor records, that were missed in due diligence, the investigation done before the deal.
- Ashcombe would then inherit those problems along with the customers.

22 New product approval



IN THE LESSON

- The main defence for a new product is a new product approval process.
- The exam point is that before launch, sign-off cannot sit with sales alone, because sales most wants the product launched.
- So the business writes a proposal describing the product, the customers and the expected volumes.
- Then the functions that will have to live with it sign off: operations, technology, risk and legal, and also compliance and finance.
- Approval often comes with conditions.
- Ashcombe might allow its new product only up to a volume limit, or only as a pilot with a few customers.
- Then comes the launch.
- Some months later, a post-implementation review checks whether the product behaves as expected and whether the conditions can be lifted.
- Risk is highest at launch, so that is where the checks bite.

23 Defences for an acquisition

- Operational due diligence before the deal
- A detailed integration plan with owners and dates
- Phased migration, with reconciliation and rollback
- Retain the key people who know the systems
- Defences aimed at launch and integration, when risk peaks

IN THE LESSON

- For an acquisition, the defences follow the same idea.
- First, operational due diligence before the deal: not just checking the accounts, but examining the target's systems, controls, data quality and open regulatory issues.
- Second, a detailed integration plan with named owners and dates.
- Third, phased migration with reconciliation and rollback.
- Ashcombe would move Kestrel's customers in batches, reconcile each batch, meaning check that every balance on the new system matches the old one, and keep a rollback plan to return to the old system if a batch goes wrong.
- Fourth, retain the key people who know the systems until integration is complete.
- In both cases, the defences are aimed at the launch or integration period, when risk peaks.

24 When an event happens: be ready

- Aim: limit the damage and recover quickly
- Incident and crisis management: roles, escalation
- Business continuity plan: alternative sites, people, systems
- IT disaster recovery: servers, data, networks

IN THE LESSON

- However good the controls, some events will happen.
- When one does, the aim is to limit the damage and recover quickly, and that depends on preparation done beforehand.
- First, an incident response and crisis management structure, with clear roles and escalation, so everyone knows who decides, who is told, and when a problem becomes a crisis.
- Second, a business continuity plan, which keeps the business running with alternative sites, alternative people and alternative systems if the usual ones are lost.
- Third, an I T disaster recovery capability, which restores the technology itself: the servers, the data and the networks.

25 Two recovery targets: RTO and RPO

Recovery time objective and recovery point objective

Target	Meaning	Must fit it
RTO	Longest time to restore a service	Failover speed
RPO	Most data that can be lost, in time	Backup frequency

IN THE LESSON

- Two targets drive the design of recovery.
- The first is the recovery time objective, R T O.
- It is the longest acceptable time to restore a service, so it answers the question: how long can the service be down?
- The failover speed, meaning how fast the firm can switch to its backup systems, must fit the R T O.
- The second is the recovery point objective, R P O.
- It is the greatest amount of data that can be lost, measured in time, so it answers: how many minutes or hours of transactions can we afford to lose?
- The backup frequency must fit the R P O, because anything entered since the last backup is lost.

26 Example 7: does the set-up meet the targets?

Payments service: RTO 2 hours, RPO 15 minutes

Backup every 60 minutes	worst loss 60 vs RPO 15 min	**Fails**
Failover takes 4 hours	4 vs RTO 2 hours	**Fails**
Replicate every 10 minutes	10 within 15 minutes	**Meets RPO**
Standby: 90-minute switch	90 within 120 minutes	**Meets RTO**

Tighter targets cost more: set them by how critical the service is

IN THE LESSON

- Ashcombe sets targets for its payments service: an R T O of two hours and an R P O of fifteen minutes.
- Check the current set-up.
- Data is backed up every sixty minutes, so in the worst case an hour of payments is lost, four times the fifteen minute target.
- That fails.
- Switching to the backup site takes four hours, twice the two hour target.
- That fails too.
- The fix is to replicate the data every ten minutes, which is within fifteen, so it meets the R P O.
- And keep a standby site that can take over in ninety minutes, which is within one hundred and twenty, so it meets the R T O.
- Tighter targets cost more, so they should be set by how critical the service is.

27 Communicate, recover, review, test

- Tell customers, regulators, public: timely, factual
- Recover losses: insurance claims, legal action
- Review the causes afterwards
- Test plans through exercises
- Weakness: a backup that shares the primary's failure

IN THE LESSON

- During and after the event, communication to customers, regulators and the public must be timely and factual.
- Silence or spin turns an operational failure into a reputational one.
- Afterwards, the firm recovers what losses it can, for example through insurance claims or legal action against a supplier at fault.
- It also reviews the causes, so the event does not recur.
- Plans only work if they are tested through exercises, such as a simulated cyber attack or a day run from the backup site.
- And beware a common weakness: a backup that shares the same failure as the primary.
- If Ashcombe's backup data centre sat on the same flood plain, or on the same power supply, it would fail at the same moment as the main one.

28 Ways to transfer operational risk

- Insurance: premium paid, defined losses covered
- Outsourcing and contracts, such as indemnities
- Captive insurers: owned by the firm, can reinsure onward
- Capital-market instruments: loss-linked bonds

IN THE LESSON

- Now to transfer, in more detail.
- Operational risk can be transferred in several ways.
- Insurance is the most common: the firm pays a premium and the insurer pays defined losses.
- Outsourcing and contracts move risk to a supplier, for example through an indemnity in which the supplier agrees to cover losses it causes.
- A captive insurer is an insurance company owned by the firm itself.
- It lets a group pool and fund its own risks, and it can reinsure some of them onward, passing them to the wider insurance market.
- And some capital-market instruments, such as bonds whose repayment depends on a defined loss event, pass risk to investors.
- Insurance is the one the exam asks you to compute.

29 How an insurance policy pays

Excess = max(Loss − Deductible , 0)
Payout = min(Excess , Limit)
Retained = Loss − Payout

Actual loss	The firm's actual loss
Deductible	Kept by the firm first
Limit	Most the insurer will pay
Retained loss	Kept at both ends

IN THE LESSON

- For an insurance policy, the firm keeps losses up to the deductible.
- The insurer pays above the deductible, up to the limit, and the firm keeps anything above the limit.
- Write it in two steps.
- The excess is the loss above the deductible, floored at zero: the larger of loss minus deductible, and zero.
- The payout is the excess, capped at the limit: the smaller of the excess and the limit.
- Whatever the insurer does not pay, the firm retains, so the retained loss is the loss minus the payout.
- Notice that the firm is exposed at both ends: the small losses below the deductible, and the huge ones above the limit.

30 Example 8: three losses, one policy

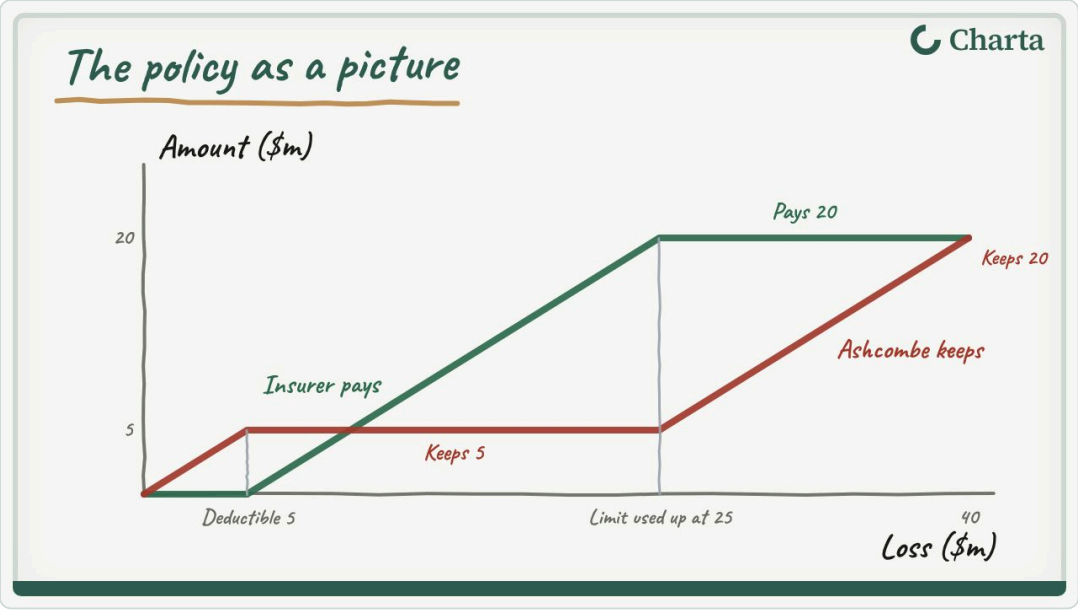
Cyber policy: deductible $5m, limit $20m. Three possible losses

Loss $3m	max(3 − 5, 0) = 0	**Pays $0**
Loss $15m	excess 10; min(10, 20)	**Pays $10m**
Loss $40m	excess 35; min(35, 20)	**Pays $20m**
Retained by Ashcombe	loss − payout	**$3m, $5m, $20m**

IN THE LESSON

- Ashcombe's cyber policy has a deductible of five million dollars and a limit of twenty million.
- Take three possible losses.
- A loss of three million is below the deductible: the excess is the larger of minus two and zero, which is zero, so the insurer pays nothing.
- A loss of fifteen million has an excess of ten million, which is below the limit, so the insurer pays ten million.
- A loss of forty million has an excess of thirty five million, but the payout is capped at the limit of twenty million.
- What does Ashcombe retain?
- Three million, five million and twenty million.
- The largest loss leaves the largest retained amount, because the cover runs out.

31 The policy as a picture



IN THE LESSON

- Here is the same policy as a picture.
- Along the bottom is the size of the loss, from zero to forty million dollars, and up the side is the amount in millions.
- The green line is what the insurer pays.
- It stays at zero until the loss reaches the deductible of five million, then rises dollar for dollar, and goes flat at twenty million once the loss reaches twenty five million, where the limit is used up.
- The red line is what Ashcombe keeps.
- It rises to five million, stays flat while the insurer pays, and then climbs again once the loss passes twenty five million.
- The picture shows the two gaps in the cover: small losses below the deductible, and very large losses above the limit.

32 When insurance works, and when it does not

- Works best: rare, severe, well-defined losses
- Works poorly: frequent small losses inside the deductible
- Limits: exclusions, basis risk, slow payment
- Limits: insurer's credit quality, moral hazard
- Fines and lost franchise value: often uninsurable

IN THE LESSON

- So when does insurance actually work?
- It works best for rare, severe, well-defined losses, like the flood.
- It works poorly for frequent small losses, which fall inside the deductible, as we saw with card fraud.
- Even when cover exists, it has limits.
- Exclusions remove some causes from cover.
- Basis risk is the gap when the cover does not match the actual loss, for example a policy that pays for physical damage when the real loss was lost business.
- Payment can be slow, while the firm needs cash now.
- The insurer's own credit quality matters, since a weak insurer may not pay.
- And moral hazard means the insured firm may take less care because it is covered.
- Finally, fines and lost franchise value, meaning damage to future business, are often uninsurable.

33 Reputational risk

- Arises mainly as a consequence of other failures
- Cannot really be transferred
- Prevent the events that cause it
- Sound conduct and culture
- Fast, honest crisis response

IN THE LESSON

- That last point leads to reputational risk, the risk that the firm's standing with customers, investors and regulators is damaged.
- It arises mainly as a consequence of other failures: a fraud, an outage, a scandal over how products were sold.
- That is why it cannot really be transferred.
- No policy restores customers' trust.
- Instead it is managed at its sources.
- Prevent the events that cause it, with the controls in this lesson.
- Keep sound conduct and culture, so the firm does not create scandals in the first place.
- And when something goes wrong, respond fast and honestly.
- The exam asks whether a transfer actually works in the stated circumstances, and for reputation the answer is almost always no.

34 Try three quick questions

1. Deductible $2m, limit $10m, loss $15m. Insurer pays? Firm keeps?
   Answer: Excess 13, capped: pays $10m, keeps $5m.
2. 5 independent steps, each 2% error chance. Chance of at least one error?
   Answer: 1 − 0.98^5 = 9.61%.
3. Tester samples 40 trades and re-performs the approval. Which test?
   Answer: Operating effectiveness: did it work over the period?

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- A policy has a deductible of two million dollars and a limit of ten million.
- A loss of fifteen million occurs.
- How much does the insurer pay, and how much does the firm keep?
- The excess is thirteen million, capped at the limit, so the insurer pays ten million and the firm keeps five million.
- Question two.
- A process has five independent manual steps, each with a two per cent chance of error.
- What is the chance of at least one error?
- One minus point nine eight to the power five, about nine point six one per cent.
- Question three.
- A tester samples forty trades from the year and re-performs the approval check on each.
- Which kind of effectiveness is being tested?
- Operating effectiveness, because it asks whether the control actually worked over the period.