PADEMIQ

Shape · Build · Deliver · Grow

AI GRCA PRACTITIONER PROGRAMME

Governance, Risk, Compliance & Assurance Basics

Module 3

60-minute live session · Case study: BankFraud AI

Facilitated by Collins Obasuyi

PADEMIQ / AI GRCA Lab

From spotting harm to managing it



Module 1

What AI is, how it works, who uses it and who it affects



Module 2

Responsible AI and harm: good technical performance can still hurt people



Module 3 · Today

How organisations assess, control, own, evidence and govern that risk



This is where we stop simply identifying what could go wrong, and start understanding how organisations should manage it.

Our case study: Northbridge Bank

The bank
Northbridge Bank, a fictional UK retail bank

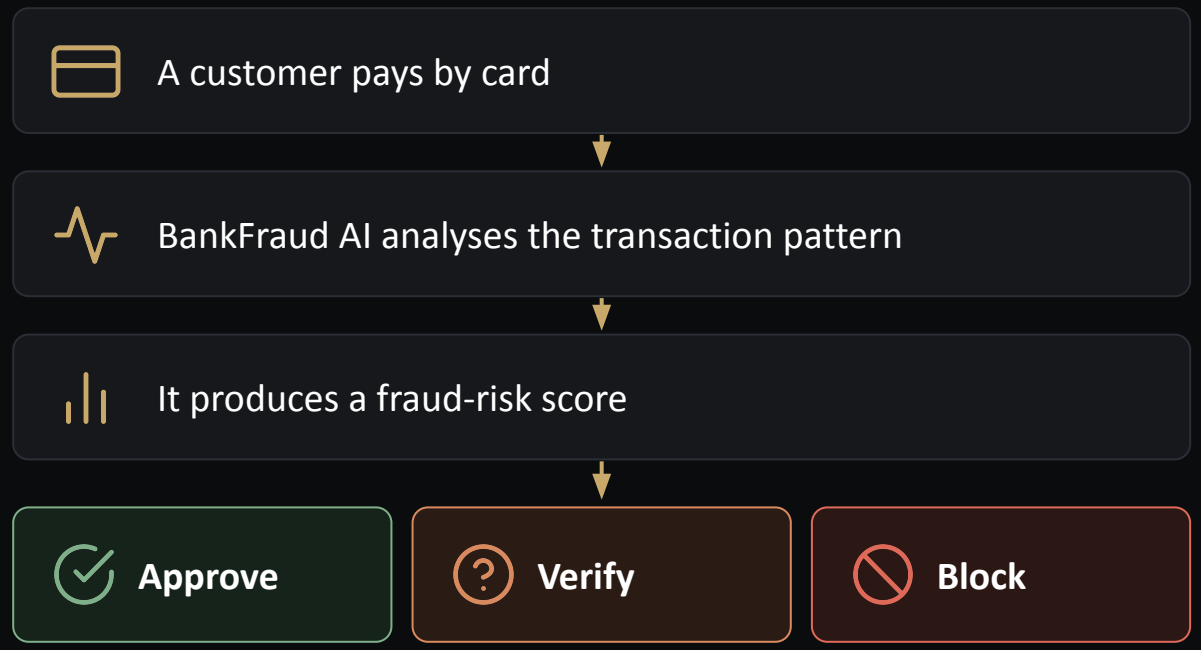
The problem
Card fraud is costing the bank and its customers money

The answer
BankFraud AI, a machine-learning system that scores every card payment for fraud risk

The people
Customers, fraud analysts, customer operations and the AI team

Training simulation: no real bank, customers or live financial decisions.

HOW BANKFRAUD AI WORKS



The score helps decide what happens to the payment

What it promises, and what could go wrong

WHAT THE BANK WANTS

- Reduce financial fraud
- Protect customers
- Detect fraud better

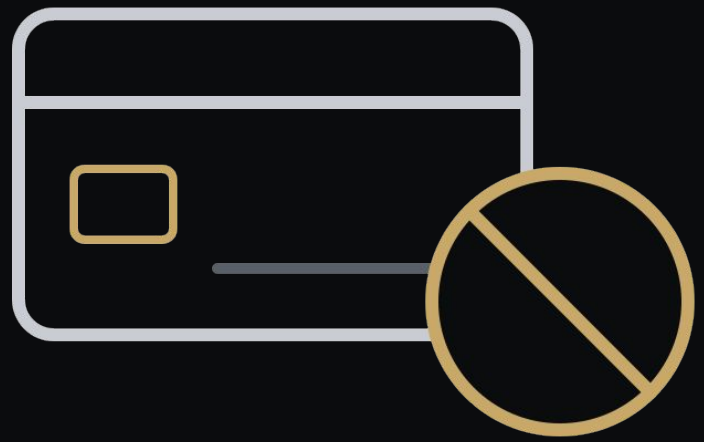
WHAT COULD GO WRONG

- Blocks genuine payments
- Misses real fraud
- Disadvantages some customers
- Staff over-rely on the score
- Delays urgent payments
- Exposes financial data
- Decisions staff can't explain

The same system that protects customers can also harm them. That tension is what today is about.

PICTURE THIS

It stops thousands of frauds. Then it blocks you



A Northbridge customer. A genuine payment. Declined.

BankFraud AI has stopped thousands of fraudulent payments. Then it blocks a genuine customer who needs money for food, travel or a place to stay tonight.

- Who is responsible for that decision?
- Who assesses the risk?
- What controls should exist?
- How does the bank prove those controls work?
- Who has the authority to decide if it keeps running?

Spotting the risk is only the beginning



Identifying AI risk is only the beginning. A practitioner must understand how risk is assessed, controlled, owned, monitored, evidenced and governed.

0:06 · 8 MIN

What is governance?

DEFINITION

The system through which an organisation sets direction, assigns accountability, establishes rules, makes decisions and oversees its activities.

IN PLAIN TERMS

Making sure the right people make the right decisions, follow the right rules and take responsibility for the outcomes.

GOVERNANCE ANSWERS

- Who is responsible?
- Who makes decisions?
- Who approves the system?
- What policies apply?
- What level of risk is acceptable?
- What happens when it goes wrong?
- Who can challenge a decision?
- Who can stop or restrict it?

Think of a school



IN SCHOOL	IN AI TERMS
Teachers Deliver the lessons	The team that builds and runs the AI
Headteacher Manages the school day to day	Management: plans, resources, operations
Governors Oversight, challenge and accountability	Governance: direction, approval, oversight

A school can have excellent teachers and still be poorly governed. The same is true of AI.

Working well is not the same as well governed

THE TECH TEAM SAYS

“The model detects suspicious transactions accurately and runs without technical errors.”

Good governance doesn't block innovation. It helps you introduce it responsibly.

THE GOVERNANCE TEAM ASKS

1. Who approved AI for blocking payments?
2. Has the customer impact been assessed?
3. Who owns the risk of wrongly blocked payments?
4. Can customers challenge decisions?
5. What if someone needs emergency access to money?
6. Are there limits for acceptable false positives?
7. Who reviews system performance?
8. Who can suspend automated blocking?

What is management?

DEFINITION

Responsible for planning, implementing and operating activities.

IN PLAIN TERMS

Getting the work done: planning it, doing it and running it day to day.

In our school: the headteacher.
In BankFraud AI: the rollout plan, daily review of flagged payments, keeping the model running.

MANAGEMENT ASKS

- How will we build or deploy the system?
- Who will do the work?
- What resources are needed?
- How will we implement the controls?
- How will we monitor day-to-day operations?

Governance versus management

Governance asks
Should we? What's acceptable? Who's accountable? What evidence do we need?

Management asks
How do we build it? Who does the work? What resources? How do we run it?

GOVERNANCE	MANAGEMENT
Sets direction and expectations	Executes organisational plans
Defines accountability	Assigns operational responsibilities
Establishes risk appetite	Implements risk treatments
Oversees important decisions	Manages daily decisions
Requires evidence	Produces operational evidence
Approves or delegates authority	Delivers and operates the system

Same upgrade, two different questions

The situation: Northbridge's AI team has an upgraded fraud model ready. In testing it catches more suspicious transactions. Two people look at the same upgrade.

MANAGEMENT: THE DELIVERY MANAGER

"We can deploy the fraud detection upgrade on Friday."

Focus: delivering the system

GOVERNANCE: THE REVIEWER

"Have we shown the upgrade won't cause unacceptable customer harm, and that monitoring, escalation and rollback controls are in place?"

Focus: is its use authorised and controlled?

QUESTION FOR THE ROOM

If the upgrade cuts fraud by 30%, should it go live on Friday? What else would you need to see?

0:14 · 8 MIN

What is risk?

RISK, SIMPLY PUT

The chance that something goes wrong, and how much harm it would do if it did.

Likelihood *How likely is it to happen?*

Impact *How bad would it be if it did?*

TO WRITE A RISK DOWN, NAME THREE THINGS

Cause
Why it might happen

→

Event
What might happen

→

Impact
What it would lead to

Because of **[cause]**, **[event]** may occur, resulting in **[impact]**.

Weak and strong risk statements

WEAK	STRONG: CAUSE → EVENT → IMPACT
“My car might break down.”	Because my car has not been serviced regularly, it may break down during a long journey, resulting in missed appointments and unexpected repair costs.
“BankFraud AI may make mistakes.”	Because BankFraud AI may incorrectly classify legitimate transactions as fraudulent, genuine customer payments could be blocked, resulting in financial inconvenience, customer distress, complaints and potential regulatory concerns.

The strong version gives you enough to start identifying controls.

Two different risks, both matter



FALSE POSITIVE

Genuine payment flagged

Because BankFraud AI may wrongly classify legitimate transactions as fraudulent, genuine payments could be blocked, resulting in distress, complaints and regulatory concern.

Who's hurt first: the customer



FALSE NEGATIVE

Fraud slips through

Because BankFraud AI may not recognise new fraud patterns, fraudulent transactions may go undetected, resulting in financial losses for customers and the bank.

Who's hurt: customers and the bank



Push one down and the other usually rises. Where to strike that balance is a governance decision, not just a technical one.

Risk, issue, harm, control, evidence



Risk

Something that might happen

Genuine payments may be incorrectly blocked



Issue

A problem that has happened or is happening

Testing finds a high rate of wrongly blocked payments



Harm

The negative consequence

A customer can't pay for essentials



Control

A measure designed to manage the risk

Rapid customer verification and manual review



Evidence

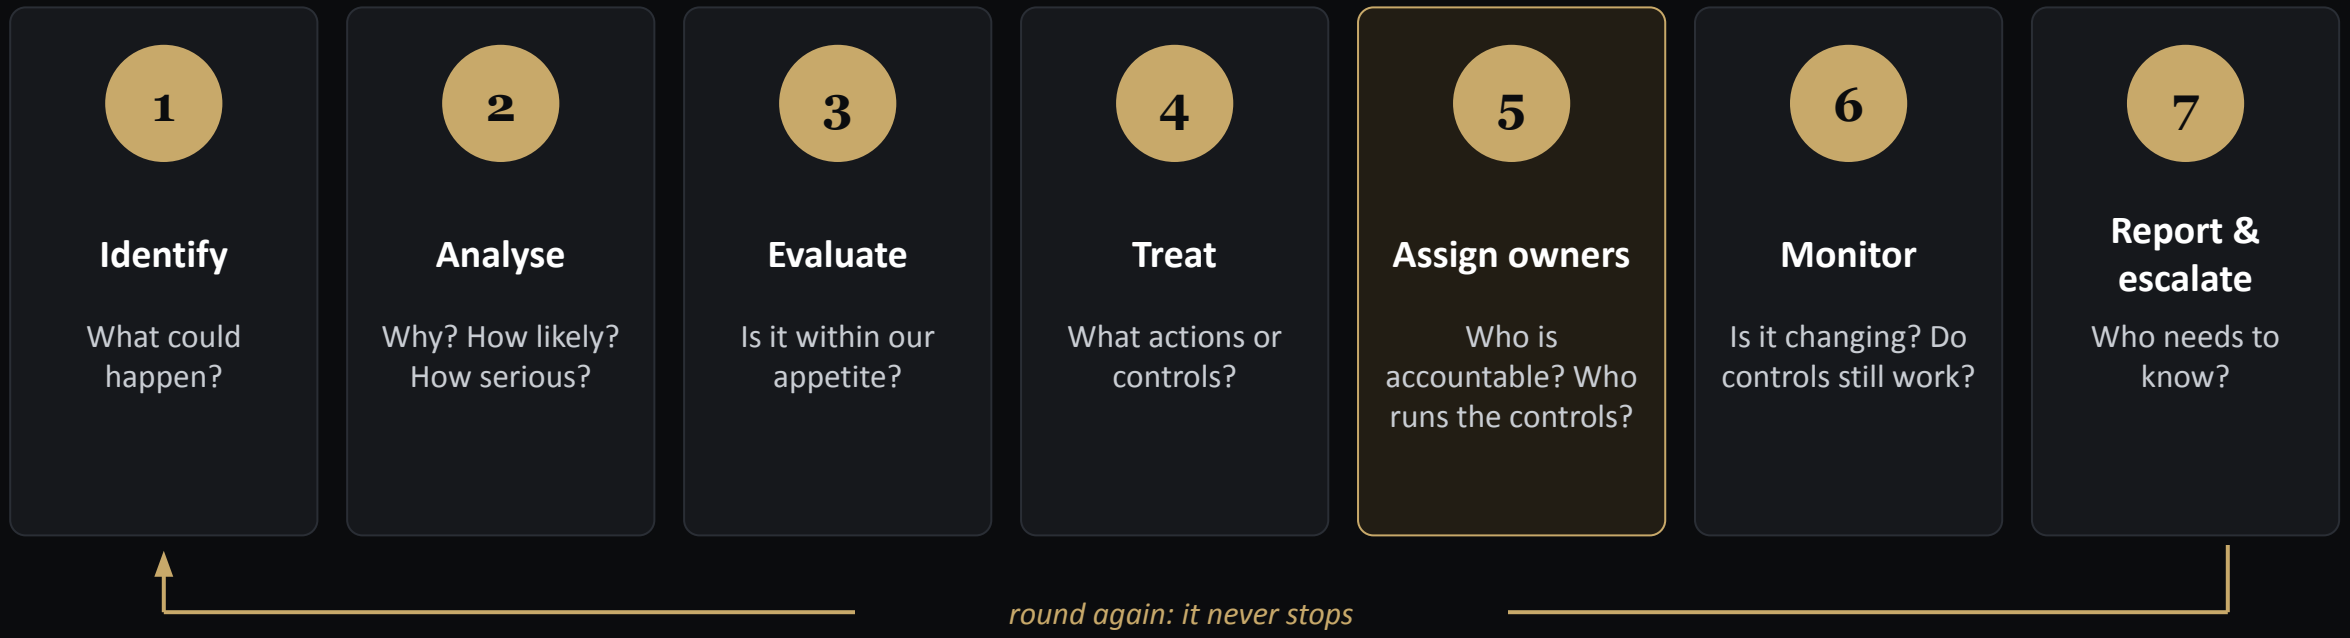
Proof supporting the control or assessment

Review logs, test results, complaint records



A risk might happen. An issue has happened. Harm is the consequence. Controls address the risk. Evidence shows they work.

Risk management is a cycle, not a document



A risk register is not risk management by itself. It's a tool. The real work is ownership, action, monitoring and decisions.

The cycle applied to BankFraud AI

STEP	BANKFRAUD AI: GENUINE PAYMENTS WRONGLY BLOCKED
1 Identify	Genuine customer payments may be blocked incorrectly
2 Analyse	Impact could be serious when customers need urgent access to funds
3 Evaluate	Does it fall within the bank's approved tolerance?
4 Treat	Customer verification, rapid manual review, escalation for vulnerable customers, false-positive monitoring
5 Assign owners	Fraud Operations is accountable for the operational risk
6 Monitor	False-positive rates, complaints, time taken to resolve blocked transactions
7 Escalate	Unacceptable outcomes go to authorised management, who may restrict or suspend automatic blocking

0:22 · 9 MIN

What is compliance?

DEFINITION

Meeting applicable legal, regulatory, contractual and internal organisational requirements.

IN PLAIN TERMS

Knowing which rules apply, and making sure the organisation follows them.



Five-star reviews don't remove food safety rules. Strong AI performance doesn't remove financial regulation.

What may apply to a UK bank using AI for fraud

- UK GDPR & data protection
- Relevant FCA requirements
- Fair treatment & Consumer Duty
- Operational resilience & incidents
- Information security
- Outsourcing & third-party risk
- Internal model governance & risk policy
- Contractual obligations

The exact obligations depend on the bank, the system and how it is used.

Compliance is essential. But compliance alone doesn't guarantee an AI system is safe, fair or effective.

Four lenses, four questions



Governance

Who decides and oversees?

BANKFRAUD AI

A designated authority approves the AI use and its risk limits



Risk

What could go wrong, and what are we doing about it?

BANKFRAUD AI

Genuine payments may be wrongly blocked



Compliance

Are we meeting the requirements that apply?

BANKFRAUD AI

Data protection and financial services obligations



Assurance

What evidence shows our controls work?

BANKFRAUD AI

Testing, customer review records, monitoring evidence

What is assurance?



“My car is safe because I look after it.”

DEFINITION

Evaluating evidence to develop confidence that risks, controls and governance arrangements are working as intended.

IN PLAIN TERMS

Checking whether the organisation is actually doing what it says it's doing.

Like an MOT: an independent check against defined requirements at a point in time. Evidence, not a guarantee it can never fail.

“Show me it works”

THE BANK SAYS

“Customers can request a manual review of blocked transactions.”

THE ASSURANCE REVIEWER ASKS

- Is there a documented review process?
- Can customers actually access it?
- Are staff trained?
- How quickly are reviews completed?
- Are decisions recorded?
- Are urgent cases escalated?
- Do case records show it working?
- What happens when it fails?

A policy doesn't prove a control works. A procedure doesn't prove it was followed. Assurance needs evidence.

Designed well? Working in practice?

Control: customers can request urgent manual review of a blocked transaction.

Design effectiveness

Is it built to manage the risk?

- Is there a documented procedure?
- Are urgent cases defined?
- Are responsibilities assigned?
- Can staff release payments where appropriate?
- Are target review times set?
- Is there an escalation route?

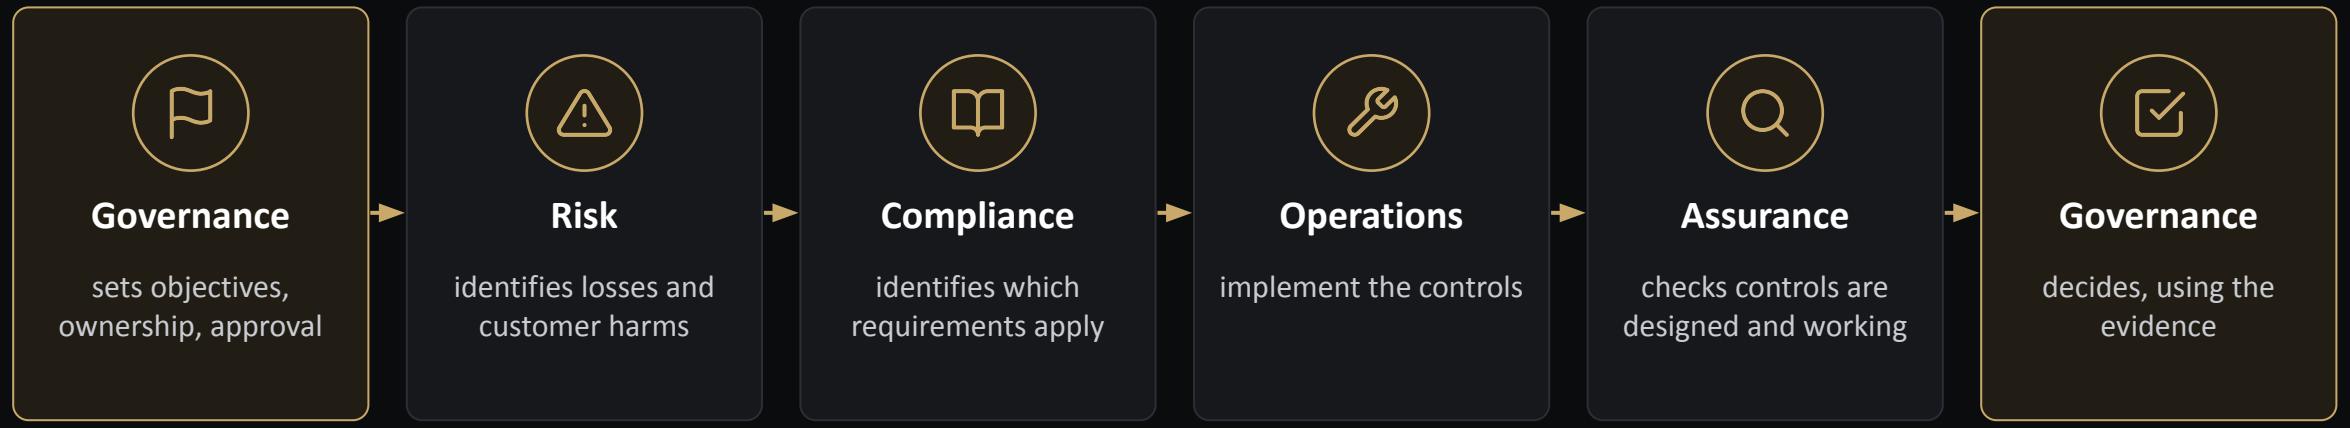
Operating effectiveness

Does it work in practice?

- Are requests handled through the process?
- Do records show appropriate decisions?
- Are review times monitored?
- Are failures investigated?
- Are overdue cases escalated?

A well-designed control can fail if people don't follow it. A control performed every day can still be badly designed.

The complete picture



THE MAIN LESSON

Governance without risk management **is incomplete**

Risk management without controls **is ineffective**

Controls without evidence **are unverified**

Evidence without accountability **isn't a governance decision**

0:31 · 8 MIN

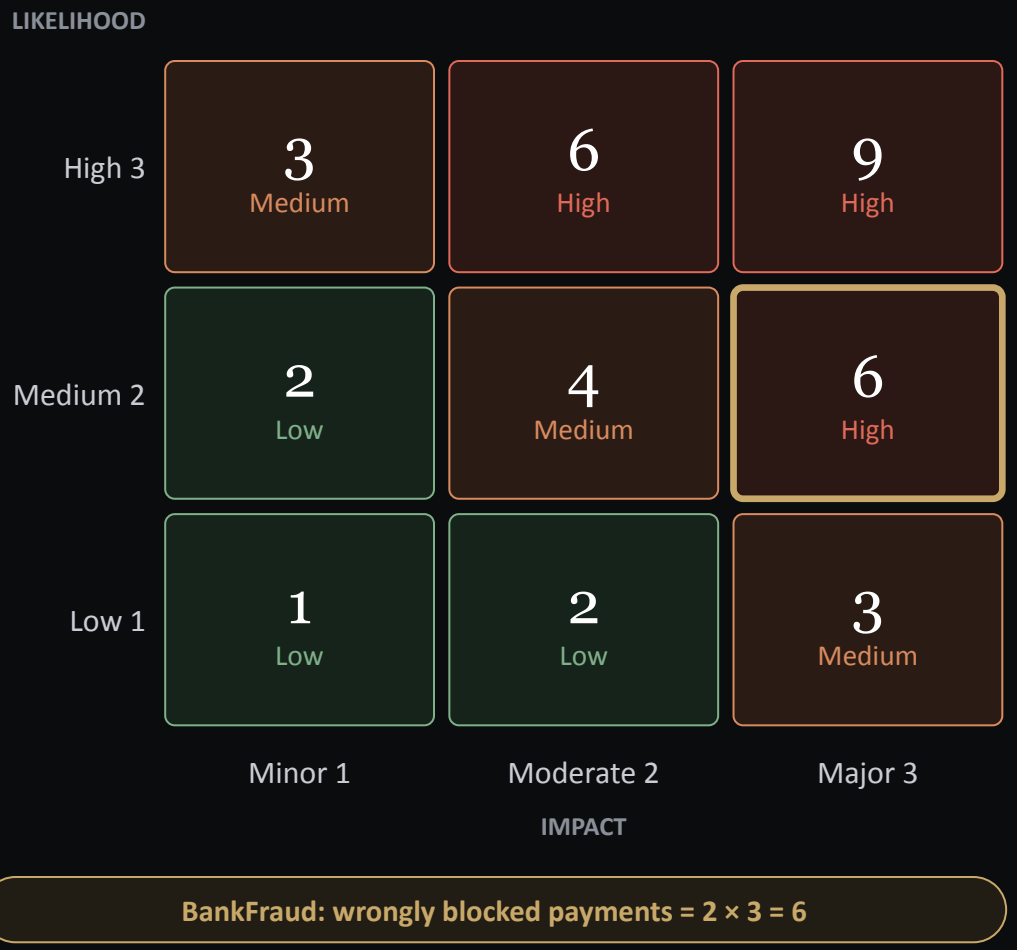
Likelihood × impact

Likelihood
What are the chances of this happening?

Impact
How bad would it be if it happened?

Risk score = **Likelihood × Impact**

A score supports judgement; it doesn't replace it. A rare event can still be unacceptable if the consequences are severe.



BankFraud: wrongly blocked payments = 2 × 3 = 6

What is risk appetite?

RISK APPETITE, SIMPLY PUT

How much risk are we willing to take on to get what we want?

FORMAL DEFINITION

The type and amount of risk an organisation is willing to accept in pursuit of its objectives. Set by senior leadership.

APPETITE IS A SCALE, AND IT DIFFERS BY TYPE OF RISK

Averse	Minimal	Cautious	Open	Eager
Avoid it	Only very low risk	Some, if well controlled	Willing, for a clear benefit	Actively seek it

Northbridge: harming vulnerable customers

Northbridge: some extra checks on genuine payments, to stop fraud

Everyday: one person keeps savings in cash, another buys shares. Different appetites. Neither is wrong, as long as it's a deliberate choice.

What is risk tolerance?

RISK TOLERANCE, SIMPLY PUT

The line we won't cross. The point where we say: this has gone too far, act now.

FORMAL DEFINITION

The limits of acceptable variation around an objective or risk exposure.

TOLERANCE LINE

Within appetite	Approaching the limit	Tolerance breached
Carry on, keep monitoring	Watch closely, escalate	Act now

Everyday: your phone battery. You're relaxed until about 20%, then you look for a charger. 20% is your tolerance.

NORTHBRIDGE: ILLUSTRATIVE LIMITS

- Urgent blocked payments resolved within an agreed time
- Genuine payments flagged: above an agreed rate, escalate
- Complaints about blocks: above an agreed weekly level, senior review

Side by side: appetite and tolerance



Planning a journey

Risk appetite

How much risk are we prepared to take?

The type and amount of risk we'll accept while pursuing our objectives.

Journey: You're comfortable allowing 20 minutes extra for traffic.

Risk tolerance

At what point has it gone too far?

The limits of acceptable variation around an objective or exposure.

Journey: At 45 minutes late, you change your plans.

Appetite has to become practical limits

WE CAN ACCEPT

Some extra checks on genuine transactions, in order to stop fraud

WE WON'T ACCEPT

- Prolonged blocking of essential payments
- Poor treatment of vulnerable customers

SO THE POLICY COULD DEFINE

- Acceptable performance ranges
- Escalation triggers
- Maximum review times for urgent cases
- When manual review is required
- Thresholds for senior management intervention
- When automated blocking must be restricted

If everyone says a risk is unacceptable but nobody defines how to detect or escalate it, governance is weak.

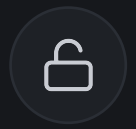
What is inherent risk?

INHERENT RISK, SIMPLY PUT

The raw risk: how big it is before we do anything about it.

FORMAL DEFINITION

The level of risk assessed before taking account of mitigating controls.



EVERYDAY

A house on a dark street with no locks, no alarm and no lights. That's the inherent risk of burglary.



BANKFRAUD AI

Blocking payments automatically, with no human review, no customer verification and no monitoring.

Likelihood: Medium · Impact: High = Inherent: HIGH



It shows how much work your controls are doing. If they fail, the inherent risk is what you're left with.

What is residual risk?

RESIDUAL RISK, SIMPLY PUT

The risk that's left once our controls are in place and working.

FORMAL DEFINITION

The risk remaining after controls have been implemented and considered.

Inherent risk
the raw risk

–

Controls that work
proven by evidence

=

Residual risk
what's left

Everyday: lock the doors, fit an alarm, turn the lights on. Burglary is less likely, but still possible.

BankFraud AI: with verification, rapid review and monitoring working, residual might drop to Medium.

Residual risk is never zero. Leadership decides whether what's left sits within appetite.

Side by side: inherent and residual

INHERENT

Before controls

High

Likelihood: Medium
Impact: High

→

PROPOSED CONTROLS

- Better transaction classification testing
- Customer verification
- Rapid manual review
- Escalation for urgent and vulnerable cases
- False-positive monitoring
- Incident response procedures

→

RESIDUAL

After controls

Lower?

Only if controls are implemented and shown to work



Like driving: seatbelts, speed limits and safe driving reduce the risk. Some always remains.

Listed is not the same as working

Proposed

It's on paper

Implemented

It's in place

Effective

Evidence shows it works

Only this step can justify a lower residual rating

QUESTION FOR THE ROOM

A control has been proposed but not yet implemented. Can you lower the residual risk rating on the strength of the proposal alone?

Four ways to treat a risk

Avoid

Stop, or don't do, the risky activity

No automatic blocking of certain critical payments until safeguards are shown

Reduce

Controls that lower likelihood or consequences

Human review, model testing, rapid customer verification

Share

Allocate some consequences via contracts or insurance

Insurance or indemnities. Doesn't remove accountability to customers

Accept

Keep an understood risk within approved limits

A documented, monitored level of routine false positives

Always: monitor and review. A risk acceptable yesterday can change with new fraud patterns, model updates or control failures.

0:39 · 9 MIN

One risk, one owner, many controls

Risk owner
Accountable that the risk is understood and managed

RISK Genuine transactions may be incorrectly blocked
Risk owner: Head of Fraud Operations

Control owner
Operates and maintains one specific control



Customer verification

Control owner

Fraud Operations Manager



Model performance testing

Control owner

AI Model Validation Lead



Complaints escalation

Control owner

Customer Support Ops Lead



Access restrictions

Control owner

Information Security Lead



"IT" or "the business" is not an owner. Significant AI risks need a clearly named, accountable role.

What is RACI?

RACI, SIMPLY PUT

A simple chart that shows who does what: for each task, who does it, who owns it, who's asked and who's told.

WHY IT MATTERS

Without it, everyone assumes someone else is doing it, or too many people think they're in charge.

EVERYDAY: ORGANISING A TEAM LUNCH

R	**Responsible** does the work	Sam books the restaurant
A	**Accountable** owns the result, makes the final call	The team manager approves the budget and the choice
C	**Consulted** asked before it's done	Colleagues with dietary needs are asked first
I	**Informed** kept in the loop	The whole team is told the time and place



Golden rule: one A per task. If two people are accountable, nobody is.

RACI in practice: BankFraud AI

R — Responsible
Performs the work

A — Accountable
Owns the result or decision

C — Consulted
Gives specialist input

I — Informed
Receives relevant information

EXAMPLE: BANKFRAUD AI MODEL REVIEW

ROLE	RACI	WHAT THAT MEANS HERE
AI Model Validation Lead	R	Responsible for testing
Designated Model Risk Owner	A	Accountable for the review outcome
Compliance and Customer Operations	C	Consulted
Governance Committee	I / approve	Informed or approving, depending on delegated authority

Policy, standard, procedure, guidance

Policy

What must happen?

All material AI systems must have a documented risk assessment before deployment.

Standard

What requirements must be met?

Risk assessments must record affected people, ratings, controls, evidence and accountable owners.

Procedure

How do we do it?

Register the use case → assess risk → identify controls → test → approve → record the decision.

Guidance

How do we understand and apply it?

How to spot false-positive risk in AI fraud detection systems.

QUESTION FOR THE ROOM

In one sentence: what policy statement would you write for Northbridge Bank's use of AI in fraud detection?

Stop it, spot it, fix it

A control is something that helps stop a problem, discover it, or correct it.



Preventative

Stop it happening

- Model validation before deployment
- Restrict automated blocking for defined transaction types
- Access controls
- Change approval requirements
- Extra customer verification rules



Detective

Spot it when it does

- Monitor false-positive rates
- Monitor false-negative rates
- Analyse customer complaints
- Detect unusual block patterns
- Periodic model performance review



Corrective

Put it right afterwards

- Manually release wrongly blocked payments
- Customer remediation
- Investigate model failures
- Roll back to an approved model
- Retrain or change rules, with approval

A control is not a good intention

WEAK

“Fraud staff should check AI decisions.”

STRONG

“Fraud analysts must review transactions meeting defined escalation criteria, record the evidence considered, document their decision and escalate urgent customer-impact cases within the approved service standard.”

FIVE QUESTIONS FOR EVERY CONTROL

1. What risk does it address?
2. Who owns it?
3. When does it operate?
4. How does it work?
5. What evidence proves it operates?

Evidence: “show me proof”

THE CLAIM	THE EVIDENCE
The model was tested before deployment	Model validation report
Fraud analysts review flagged cases	Completed review logs
Customers have a challenge route	Procedure and actual case records
False positives are monitored	Monitoring dashboard and review minutes
Model changes are controlled	Change approval records
Staff understand the procedure	Training and competency records
Senior management approved the risk	Dated approval and risk acceptance record



Information supporting a claim, activity, finding or decision.

What good evidence looks like

WEAK

"The project manager confirmed everything was checked."

STRONG

Dated test results, documented review findings, signed approvals and operational records showing the controls were followed.

GOOD EVIDENCE IS…

- ✓ Relevant
- ✓ Accurate
- ✓ Complete enough for its purpose
- ✓ Traceable to the system
- ✓ Dated
- ✓ Reliable
- ✓ Appropriately retained
- ✓ Protected from alteration

A procedure proves the procedure exists. It doesn't prove anyone followed it.

The Three Lines Model

Governing body: oversight of the whole organisation, receives reporting from all three lines

1 First line
Operational management

- Fraud Operations
- Payments Operations
- AI product & engineering
- Customer Operations

2 Second line
Specialist oversight & challenge

- Risk Management
- Compliance
- Data Protection
- Information Security
- Model Risk Management*

3 Third line
Independent assurance

- Internal Audit: are AI governance requirements followed, and do selected controls work?

**depending on the bank's structure*

Those who build and run the AI shouldn't be the only ones judging their own controls.

0:48 · 4 MIN

Three frameworks you'll keep meeting



NIST AI RMF

US National Institute of Standards and Technology · 2023

A free, voluntary framework for managing AI risk

THINK OF IT AS

A playbook for spotting and managing AI risk

Voluntary, no certification



ISO/IEC 42001

International standards bodies ISO and IEC · 2023

Requirements for an AI Management System (AIMS)

THINK OF IT AS

The rulebook for running AI responsibly across a whole organisation

Organisations can be certified



COSO

Committee of Sponsoring Organizations of the Treadway Commission · 2013

A framework for internal control, used by auditors

THINK OF IT AS

The blueprint for checking internal controls hold together

Not AI-specific; AI controls sit inside it



They don't compete. NIST helps you manage AI risk, ISO/IEC 42001 runs it as a system, COSO checks the controls hold together.

NIST AI RMF: four functions

GOVERN

Accountability, policies and responsibilities

BankFraud AI: Define AI risk owners and approval authority

MAP

Understand the context; identify impacts and risks

BankFraud AI: Identify affected customers and possible harms

MEASURE

Assess and evaluate AI risks and characteristics

BankFraud AI: Evaluate false positives, false negatives, reliability, fairness

MANAGE

Prioritise, respond to and monitor risks

BankFraud AI: Implement controls, make risk decisions, monitor outcomes

Deep dive from Module 5 onwards.

ISO/IEC 42001 and COSO

ISO/IEC 42001

AI Management System (AIMS)

How an organisation sets up, runs, maintains and continually improves the way it manages AI, across every AI system it uses.

Responsibilities · Policies · Risk assessment & treatment · Documented information · Operational controls · Monitoring · Internal audit · Management review · Continual improvement

A risk register for one AI system is not a complete AI Management System.

COSO Internal Control

Five components that work together

1. **Control environment**
   Leadership defines accountability
2. **Risk assessment**
   Fraud, customer harm, operational risk
3. **Control activities**
   Validation, manual review, access limits
4. **Information & communication**
   Reporting, escalation, procedures
5. **Monitoring activities**
   Performance, incidents, control effectiveness

Governance starts before the first line of code

01 Idea
Is AI suitable for this problem?
Business case

02 Discovery
Who is affected?
Stakeholder & harm assessment

03 Requirements
What risks and controls apply?
Risk register

04 Design
Are controls built into the design?
Control requirements

05 Development
Have controls been implemented?
Technical records

06 Testing
Does it perform as required?
Test & validation reports

07 Release
Is the remaining risk acceptable?
Governance approval

08 Operations
Is it behaving as expected?
Monitoring & incident records

09 Changes
Have new risks been introduced?
Change assessment

10 Retirement
Can it be safely withdrawn?
Decommissioning records

Governance is continuous. Passing a pre-release review doesn't guarantee future performance.

0:52 · 8 MIN

Then the complaints arrive

Fraud losses appear to be falling. Management is pleased.

Complaints rising: genuine payments blocked, some customers without urgent access to money.

Fraud team: stricter controls are needed to stop criminals.

Customer operations: the disruption is unacceptable.

The AI Governance Committee requests a formal risk review.

WHAT SHOULD HAPPEN NEXT?

Governance Establish who has authority to decide what happens to the system

Risk Identify both the fraud risk and the harm from incorrect blocking

Compliance Confirm the financial services, data protection and internal requirements

Assurance Get and review evidence on whether the controls are effective

The practitioner's questions

01 What is BankFraud AI meant to achieve?	02 What evidence supports the reported fraud reduction?
03 What are the false-positive and false-negative rates?	04 Which customers are affected?
05 Are some groups disproportionately affected?	06 What is the impact of incorrect blocking?
07 What controls exist?	08 Have they been tested?
09 What risk remains?	10 Who is responsible?

QUESTION FOR THE ROOM

Should automated blocking continue without changes?

Continue as is · continue with changes · pause automated blocking

Saturday: the Northbridge Bank file

100,000
transactions reviewed by AI

4,000
flagged for additional action

1,000
flagged and confirmed fraudulent

3,000
flagged but confirmed genuine

120
complaints linked to incorrect blocks

25
urgent cases involving blocked payments

Fictional, simplified figures for training. Same observation period; all flagged transactions reviewed and classified. Not a real bank or model.

Read the numbers carefully

4,000 FLAGGED · 1 SQUARE = 100



1,000 confirmed fraud

3,000 confirmed genuine

True: 75% of flagged transactions were false positives (3,000 of 4,000).

Not true: 75% of all transactions were wrongly blocked. It's 3,000 of 100,000: 3%.

Unknown: The false-negative rate. We don't know how much fraud was missed.

QUESTION FOR THE ROOM

The supplier says the model is “95% accurate”. What does that actually measure, and is it enough to judge fraud detection and customer harm?

Saturday: you are the AI GRCA team

0–10 min

Understand the system

Objective, purpose, users, affected customers, decisions, evidence, gaps

Output: System & stakeholder overview

10–25 min

Identify & assess risk

At least 5 risks: cause → event → impact; rate on the course matrix

Output: 5 risk statements + ratings

25–40 min

Controls & owners

Preventative, detective, corrective; risk owner; control owner; evidence

Output: Risk-control entries

40–50 min

Review the evidence

Proposed, implemented, verified or not yet assessed? Which gaps matter?

Output: Evidence needs + actions

50–60 min

Make the call

Each group presents a recommendation

Output: Governance Decision Record

POSSIBLE DECISIONS

- Continue with existing controls
- Continue with extra restrictions
- Mandatory human review
- Suspend specified automated blocking
- Suspend wider use while investigated

Your evidence pack and homework

EVIDENCE PACK

Initial AI Risk Register for BankFraud AI

- Short system description
- 5+ clear risk statements
- Affected people
- Likelihood & impact
- Inherent ratings
- Existing & proposed controls
- Named risk & control owners
- Evidence requirements
- Implementation & verification status
- Residual risk, where supported
- Outstanding actions
- Governance recommendation

WRITTEN ASSIGNMENT

150–250 words

“Should Northbridge Bank continue using BankFraud AI in its current form?”

Use the risks, the data (correctly), the controls, who's accountable and what's still uncertain. Don't claim untested controls have reduced risk.

Optional: compare with a HireAI risk (unfair candidate rejection)

Quick-fire check

01

What's the difference between a risk and an issue?

02

Inherent risk vs residual risk?

03

Risk owner vs control owner?

04

Design effectiveness vs operating effectiveness?

05

What is a false positive in fraud detection?

06

Why isn't model accuracy alone enough to approve deployment?

Full 32-question knowledge check and glossary: learner pack.

CLOSING MESSAGE

An AI system doesn't become trustworthy simply because it's technically advanced.

It becomes more governable when organisations understand the risks, define accountability, implement controls and use evidence to make decisions.

On Saturday, we become the AI GRCA team reviewing BankFraud AI.

PADEMIQ

Shape · Build · Deliver · Grow

Thank you

Facilitated by Collins Obasuyi

PADEMIQ / AI GRCA Lab

NEXT

Saturday: BankFraud AI risk investigation (60 min)

Then Module 4: AI and the Wider Legal Landscape