AI Governance & Risk for Financial Services

Regulator-grade AI governance using your real workflows and failure modes — not stock demos. Built around the controls and risk your team owns.

WHO IT'S FOR	FORMAT	LENGTH
Compliance, risk, model-risk and internal-audit teams inside regulated financial firms.	In-house workshop	Half-day or full-day

What your team walks away with

Assess and govern an AI system against your regulator's expectations — model risk, explainability and the failure modes that draw attention.

What you'll learn

- ✓ Describe what 'AI governance' means as concrete controls and records, not an abstract policy
- ✓ Find AI use across your business lines and grade each system by the risk it carries
- ✓ Set clear demands for third-party and embedded AI and verify vendor claims instead of trusting them
- ✓ Explain explainability, bias and model drift in terms your whole team can act on
- ✓ Document AI decisions so the record survives a regulator's questions later
- ✓ Stand up an AI risk register your team will actually keep up to date

Curriculum

01 What "AI governance" really means inside a regulated firm

- The people, controls and records a supervisor expects to see, not a glossary
- Governing AI across its whole life cycle: design, deployment, monitoring, retirement
- How global expectations (EU AI Act, model-risk principles, financial-regulator guidance) line up
- Roles and responsibilities across the three lines of defence
- Turning a governance policy into controls people actually run

02 Spotting and grading AI risk across your business lines

- A simple method to discover where AI is being used, including 'shadow' AI no one logged
- Building an inventory of AI systems and what each one is used for
- Risk-tiering each system by impact, autonomy and who it affects
- Flagging the high-risk use cases that need the most oversight
- Keeping the inventory current as new tools appear

03 Third-party and vendor AI: what to demand and how to check it

- The questions and evidence to require from an AI vendor before you rely on them
- Auditing embedded AI inside tools you bought rather than built
- Testing vendor claims about accuracy, bias and security instead of taking them on trust
- Contract terms, data handling and the right to review or audit
- Where responsibility stays with your firm even when the model is someone else's

04 Explainability, bias and model drift in plain terms

- What explainability really means and how much you can reasonably demand
- How bias enters a model through data and design, and how to test for it
- Model drift: how a model quietly gets worse over time as the world changes
- Monitoring signals that tell you a model is degrading or behaving unexpectedly
- Communicating these concepts to non-technical colleagues and committees

05 Documenting AI decisions so they survive a regulator's questions

- What a defensible record of an AI decision looks like
- Capturing why a system was approved, how it was tested and who signed off
- Logging model versions, changes and the data behind a decision
- Keeping evidence retrievable months or years after the fact
- Aligning documentation with what supervisors actually ask for

06 Building a practical AI risk register your team will actually maintain

- What belongs in an AI risk register and what's just noise
- Linking each AI system to its risks, owners and controls
- Setting review cadences so the register stays alive after the audit
- Connecting the register to escalation and reporting
- Tailoring a register template to your team's real workflows

You keep

> An AI governance framework and model-risk template tailored to your team.
