Delta Capita

AI Governance

Controlling AI Risk in Financial Services

Delta Capita | AI Advisory & Consulting Services

2026

Delta Capita

The Governance Deficit | AI Without Guardrails

AI adoption in financial services is accelerating, but most institutions are deploying AI faster than they can govern it. With the EU AI Act enforceable, regulators are issuing model-specific fines, and boards face personal liability for AI outcomes they cannot explain. The risks are clear - yet 46% of firms report only partial understanding of the AI technologies they use. Delta Capita bridges this gap with domain expertise, proven accelerators, and a risk-first approach.

€35M

maximum fine per violation under EU AI Act

46%

of firms report only partial understanding of the AI technologies they use

55%

of AI use cases have some degree of automated decision-making

33%

of AI models in production are third-party controlled

Enforcement Timeline

Feb 2025
AI literacy obligations began

Aug 2025
Governance rules and obligations began

Aug 2026
AI Act fully enforced and applicable

The Blind Spot

Shadow AI
Unregistered models in production

Vendor Black Boxes
Third-party AI with no explainability

Stale Models
Never retrained or reviewed

The cost of ungoverned AI

- Regulatory fines: EU AI Act penalties up to 7% of global turnover
- Reputational damage: biased outputs erode customer trust
- Operational loss: model failures with no rollback capability

Sources: (1) European Commission, AI Act; (2) EU AI Act, Art. 99 (Penalties) & Art. 5 (Prohibited Practices); (3) EUR-Lex Reg. (EU) 2024/1689; (4) Bank of England, AI in UK Financial Services (2024); (5) FCA, AI in UK Financial Services; (6) BIS, Regulating AI in the Financial Sector; (7) U.S. Treasury, AI in Financial Services (2024); (8) Federal Reserve SR 11-7; (9) OCC Sound Practices for Model Risk Management.

Delta Capita

Governance in Practice | Lessons from the Industry

The pattern is now clear: robust AI governance means clear ownership, model inventories, risk classification and monitoring

What went wrong

Air Canada / Customer Care Errors

Air Canada's website chatbot told a customer he could retroactively claim bereavement fares, but the airline later denied that policy in the Moffatt v. Air Canada (2024 BCCRT 149) tribunal

Negligent misrepresentation, airline ordered to pay damages

Source: Amazon News

OpenAI / Data Privacy Failure

In March 2023, Italy's data protection authority (Garante) imposed an immediate temporary limitation on ChatGPT's processing of Italian users' data and opened an inquiry

Garante gave OpenAI 20 days to communicate corrective measures and warned of potential fines of up to €20 million

Source: Italian Data Protection Authority (Garante) 31 March 2023

What went right

DBS Bank / AI Model Governance

Centralised AI model inventory with PURE principles (Purposeful, Unsurprising, Respectful, Explainable), and senior-level committee oversight for 370 AI use cases and 1,500 AI models

Full model visibility, regulator-ready audit trail

Source: DBS - Responsible AI in banking: Gaining a competitive edge

NatWest / Responsible AI Framework

NatWest states that every significant AI use case undergoes an Ethical AI Impact Assessment reviewed by a dedicated AI & Data Ethics team. High-risk "edge cases" are escalated to a cross-bank AI and Data Ethics Panel.

Reduced bias, proactive FCA engagement

Source: NatWest Group - Upholding ethical use of AI and data management

AI Governance | Six Capabilities You Need

A mature governance framework covers the full risk & compliance lifecycle - from model registration to ongoing monitoring.

AI Model Inventory

Critical

Centralised register of all AI/ML models with risk tiering, ownership, and regulatory classification

75% of firms said they already use AI, the median firm expects use cases to rise from 9 to 21 in three years, and 16% of use cases were already rated high materiality

Bank of England / FCA 2024 survey

Bias & Fairness

Critical

Pre-deployment and ongoing fairness testing across protected characteristics with drift monitoring

Only 34% of firms reported AI-specific practices for data ethics, bias, and fairness (with 40% relying on non-AI-specific practices)

Bank of England / FCA 2024 survey

Model Explainability

Critical

Interpretable outputs for regulators, customers, and stakeholders including LIME/SHAP integration

81% of AI-using firms employed at least one explainability method, with 72% using feature importance and 64% using SHAP

Bank of England / FCA 2024 survey

Policy Management

High

Living AI policy framework with clear roles, escalation paths, and version-controlled standards aligned to EU AI Act

84% of firms had an accountable person for AI, 72% assigned accountability to executive leadership, but only 34% said they had a complete understanding of the AI they use

Bank of England / FCA 2024 survey

Audit Trail & Reporting

High

Immutable decision logs, model lineage tracking, and automated regulatory reporting for supervisory inquiries

Under the EU AI Act, high-risk AI systems must maintain logs automatically generated throughout the system lifecycle

EU AI Act, Article 12

Third-Party Risk

High

Governance of vendor-supplied AI including transparency requirements, validation rights, and concentration risk

33% of AI use cases were third-party implementations (up from 17% in 2022) including cloud, data and model providers

Bank of England / FCA 2024 survey

Delta Capita

The Accountability Gap | What Boards Can't Answer

Most boards are asking questions their organisations cannot answer.

The Board Asks...	*The Gap*
Inventory: “How many high-risk AI models are in production?” No centralised AI inventory. Models scattered across business lines with no consistent classification or ownership.	**We deployed 40+ models but can't tell the board which are high-risk**
Fairness: "Can we demonstrate that our AI doesn't discriminate?" No systematic bias testing framework. Fairness checked ad hoc - if at all - never monitored post-deployment.	**Our fairness testing is manual, inconsistent, and doesn't cover production**
Models: “If the regulator asks, can we explain how our AI works?" No explainability standards. Black-box vendor models produce outputs neither operations or compliance can interpret.	**We cannot explain AI decisions to regulators in the time they that expect**
Data Protection: "Is sensitive data protected from AI mis-use?" No DLP controls for AI tools. Employees input PII, client data, and proprietary information into GenAI with no guardrails.	**No DLP guardrails - staff freely paste client data into AI tools**

Delta Capita

Establishing AI Governance | Our 5 Key Phases

Delta Capita offers end-to-end support from readiness assessment through scaled deployment and ongoing governance.

AI Governance Health Check

Approach
- Rapid governance health check across 6 key areas

Our Accelerators
- *AI Governance Health Check* i.e. templated questionnaire and scoring models

Client Case Studies
- Payments: uplifted global risk and compliance policy frameworks

Policy Assessment

Approach
- Rapid policy assessment across 7 compliance areas

Our Accelerators
- *AI Policy Toolkit* i.e. 72 templated questions pre-mapped to 90 use cases

Client Case Studies
- Payments: implemented AI enabled compliance tooling assessing ~ 130 policies

Governance Framework Design

Approach
- Regulatory-aligned AI risk framework and controls

Our Accelerators
- *AI Risk Assessment Toolkit* i.e. 60+ controls surveyed with industry forums

Client Case Studies
- Commercial Bank: established operating model aligned to AI governance

Operational Workflow Controls

Approach
- Governance-aligned AI control workflows

Our Accelerators
- Access to industry experts + *proprietary workflow tools* to rapidly develop PoCs

Client Case Studies
- Commercial Bank: delivered MI compliance dashboards / attestation workflows

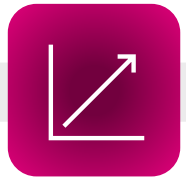

Data Readiness & Compliance

Approach
- Data quality remediation and pipeline architecture

Our Accelerators
- *DCAM certified experts* + AI data templates developed with domain experts

Client Case Studies
- Wealth manager: post-merger data lake buildout for commercial analytics

Delta Capita

Proven Results | Client Case Studies

Delta Capita serves multiple financial services clients for consulting services around their Data & AI maturity

Case Studies	AI-Enabled Third-Party Risk Assessment Global FMI Provider	Attestation Workflows for AI Readiness Challenger Bank	Data Lake & Customer Analytics UK Wealth Manager	AI Policy Assurance and Traceability Global Payments	Agentic AI for Complaints Handling Payments Fintech
Summary	Delta Capita proposed a repeatable, 3-phased onsite Third-Party Risk Assessment to help the Client evidence DORA-aligned oversight of critical ICT suppliers **DORA-aligned oversight**	Delta Capita partnered with a leading Commercial Bank to deliver a cloud-ready MI and BI transformation, establishing a scalable op model aligned to governance standards **10-month cloud delivery**	A leading UK wealth manager required a unified Power BI dashboard to visualise KPIs like customer growth, AUM, profitability, and outflows, with drill-down by region and practitioner **Delivered a 50bps margin uplift**	Delta Capita was engaged to support a Global Payments Provider in uplifting its global risk and compliance policy framework and implementing AI-enabled compliance tooling **130 policies assessed**	Delta Capita implemented an Agentic AI solution for a Payment Services Provider to streamline customer service case handling, reduce response times, and cut the backlog **Processing: 3 days → <1 min**
Approach	• **Onsite Assessment:** Conducted a structured gap analysis and prepared tailored assessment toolkits/protocols, executed onsite assessments for critical third parties to validate evidence, assess operational resilience and security controls, and capture findings consistently. • **Reporting & Recommendations:** Produced standardised third-party assessment reports with issue ratings and remediation actions, delivered a consolidated cross-vendor view and a risk-based remediation roadmap to support prioritisation and ongoing monitoring • **Accelerator (VendorLens):** Accelerated evidence collection and analysis via a controlled vendor submission workflow, contract clause checks and policy/control analysis, producing consolidated risk-based vendor scoring and dashboard outputs.	• **Discovery & Assessment**: Reviewed the existing MI/BI landscape and manual reporting processes (e.g., PowerPoint/Excel outputs), identifying key pain points, control gaps and opportunities to automate and standardise MI production and attestation. • **Solution Design & Prototyping:** Designed and prototyped interactive, evidence-driven MI dashboards and an end-to-end attestation workflow, optimising the user experience and ensuring outputs were aligned to governance and auditability needs. • **Build & Implementation:** Developed and configured Power BI / Power Apps components, integrated dashboards into the target data environment, and implemented Role-Based Access Control (RBAC) to strengthen data security and appropriate user access. • **Deployment & Change Management:** Rolled out the solution with structured change management - stakeholder engagement, training, documentation, and adoption support - to embed sustainable BAU ways of working.	• **Data Integration and Predictive Insight Generation**: Connected live to on-premise SQL databases and built data pipelines to automate refreshes of key commercial and fee data. Used predictive analytics to create dashboards that gave the business real-time insight into commercial performance. • **Risk Modelling and Consumer Duty Governance**: Developed client risk models using behavioural signals such as sub-account closures and reduced activity to identify potential issues early. Also introduced governance workflows around fee data and reporting to improve oversight. • **Reporting Framework and Regulatory Alignment**: Established a scalable reporting framework that improved the quality, consistency, and accessibility of management information. • **Commercial Delivery and Revenue Enhancement**: Supported the launch of the 'Dual Expert' product, contributing to multi-million-pound revenue growth through stronger proposition delivery and business readiness. Delivered a 50bps margin uplift by improving pricing consistency.	• **Policy Review and Advisory:** Assessed ~130 global risk and compliance policies across merchant and consumer sectors. Mapped internal policies to external regulatory requirements. Identified gaps, duplication, and consolidation opportunities to streamline compliance operations. • **Process Mapping and Ownership:** Mapped end-to-end business processes with regulatory tagging. Defined clear ownership structures and documented risk and control flows to enhance transparency and accountability. • **AI Tooling and Self-Serve Portal:** Designed and implemented an AI-powered compliance portal to enable sales and underwriting teams to quickly access relevant policy information, accelerating decision-making. • **Strategic Enablement and Scalability:** Created a centralised, searchable policy and regulatory repository. Delivered a scalable framework to maintain policy alignment post-regulatory changes and support long-term compliance agility.	• **Agentic AI Triage & Categorisation:** Implemented an AI agent to analyse incoming customer requests, detect sentiment, and categorise cases using regulatory policies and customer case history to determine the appropriate handling route. • **Confidence Scoring & Routing:** Introduced confidence scoring to automate routing of cases to the correct team, with a controlled manual review path for low-confidence cases to maintain quality and reduce operational risk. • **Reasoning Summary & Escalation:** Generated concise reasoning summaries to support faster decision-making and escalated cases to relevant teams with clear context, improving consistency and speed of resolution. • **Outcomes:** Reduced case processing time from ~3 days to <1 minute, enabled instant acknowledgement and immediate routing on first contact, achieved a 25–30% reduction in overall case resolution time, and improved customer experience.

Recommended Next Steps

1

AI Governance Health Check

Assess governance maturity across model inventory, policy coverage, validation, monitoring, third parties and board reporting - benchmarked against EU AI Act requirements

2

Model Inventory & Risk Classification

Build a complete, categorised register of all AI/ML models in production and development with risk tiering, ownership, and regulatory classification

3

Governance Framework Design

Define your target-state AI governance operating model including policies, roles (SMCR-aligned), validation standards, and escalation procedures

4

Embed & Operationalise

Deploy governance tooling, monitoring dashboards, and attestation workflows - turning policy into embedded, evidenceable controls

Let's discuss how Delta Capita can help you govern and control AI risk. Get in touch: deltacapita.com/get-in-touch