



TOPICS:

Model Risk

SOURCE:

The Alan Turing Institute

ATI: GenAI Model Risk Management and Governance in Financial Services - From Principles to Practice

- This report by the Alan Turing Institute (ATI) examines how financial institutions (FIs) can adapt existing **model risk management (MRM) frameworks** - namely the U.S. Federal Reserve's SR 11-7 and the UK PRA's SS1/23 - **to govern generative AI (GenAI)** systems. GenAI introduces new dimensions of model risk due to its reliance on large, vendor-hosted foundation models, qualitative outputs and dynamic pipelines, often incorporating Retrieval-Augmented Generation (RAG) architectures. The report outlines a taxonomy of emerging risks - data, vendor, architectural, and human factors - and their associated governance priorities.
- **RAG systems** pose specific challenges such as **data quality inconsistencies, legal/compliance burdens** and **lack of stable ground truth**, exacerbated by the evolving nature of document corpora and opaque vendor updates. Vendor dependencies amplify risks tied to availability, cost, artefact versioning, and contractual ambiguity. **Architectural risks** include system fragility from modular integrations, behavioural drift, and reproducibility gaps.

Human factors such as automation bias and cognitive offloading further complicate assurance.

- **Two case studies** - the Digital Credit Platform and the Lead Recommendation Engine - illustrate practical adaptations. Both use robust validation, real-time monitoring and interdisciplinary governance involving risk, compliance, and AI Centres of Excellence. Key themes include the importance of monitoring over static validation, integrated oversight across components and treating the GenAI pipeline - not just the LLM - as the unit of risk.
- The report recommends **extending model inventories to capture full GenAI workflows**, refining risk tiering to account for GenAI-specific dimensions, formalising lifecycle monitoring, embedding vendor oversight and developing cross-functional governance structures. It concludes that GenAI does not require a separate regulatory regime but demands extensions to existing MRM practice that reflect its volatility, complexity, and third-party reliance.

