zango
RESEARCH INITIATIVE

2026

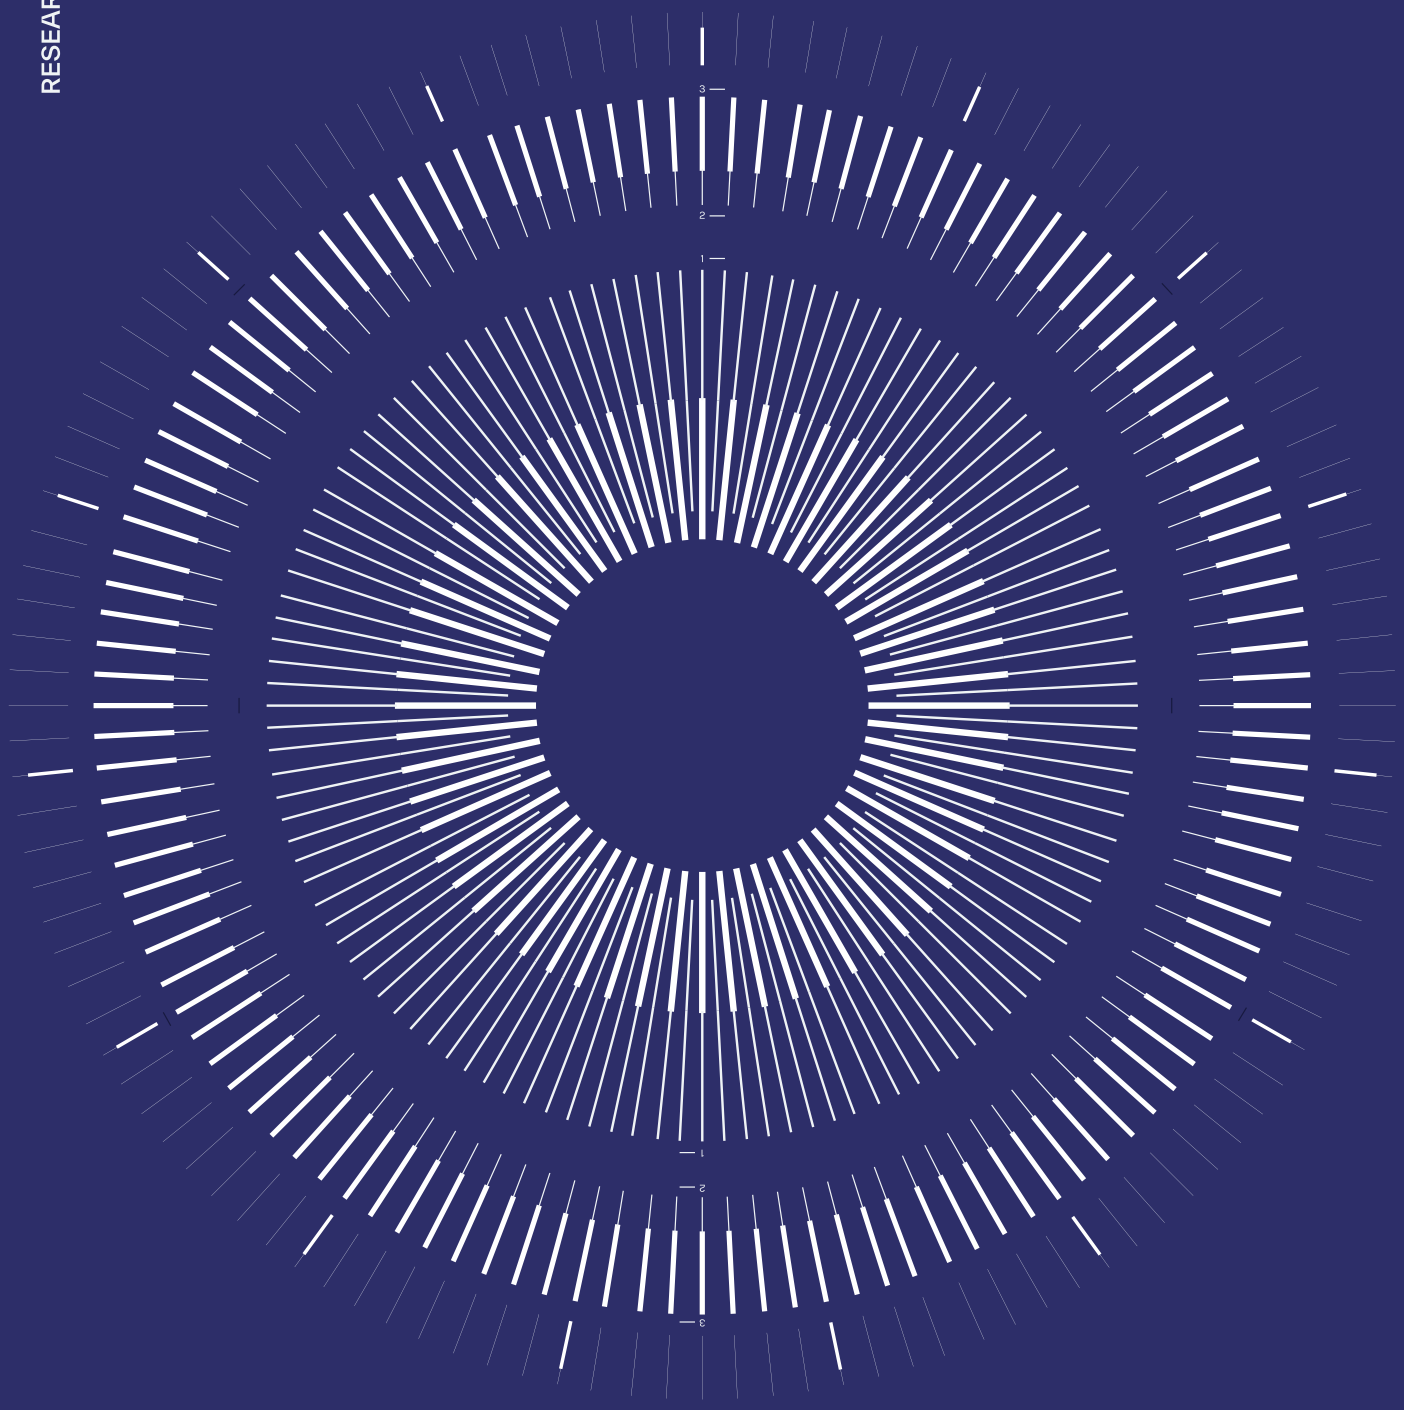

THE FUTURE OF AI GOVERNANCE & COMPLIANCE IN FINANCIAL SERVICES

RESEARCH REPORT

2026

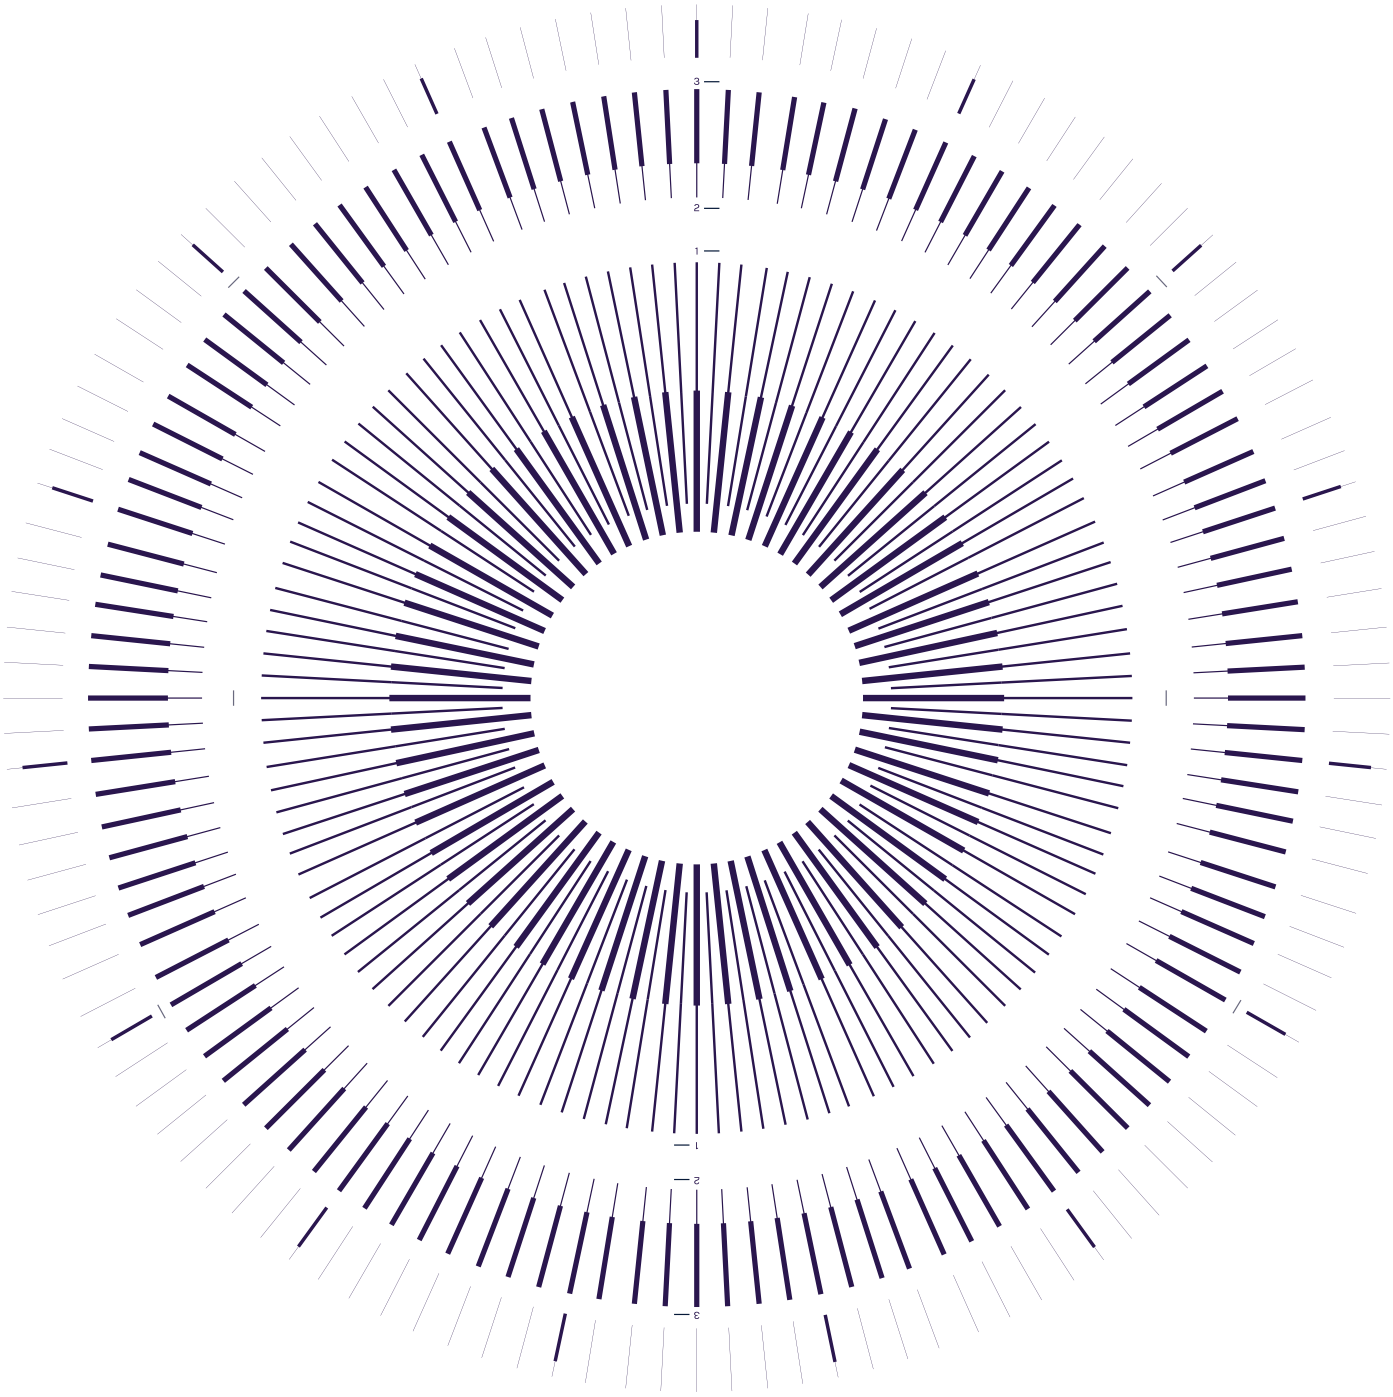

THE FUTURE OF AI GOVERNANCE & COMPLIANCE IN FINANCIAL SERVICES

RESEARCH INITIATIVE
zango

Research team

This research initiative was coordinated by Zango with support from independent research advisers specialising in AI governance and financial regulation.

Research lead

Sam Green
Senior Policy and Partnerships Lead, Zango

Sam previously led the supervision and compliance policy team in the UK Government's AI Regulation Unit.

Research advisers

Dr Alessio Azzutti
Lecturer in Law & Technology (FinTech), University of Glasgow

Andrew Sutton
Visiting Fellow, Oxford Martin School AI Governance Initiative

Industry adviser

Dean Nash
Global Chief Operating Officer (Legal), Santander

Contributors

Lord Tim Clement-Jones
Co-Chair, All-Party Parliamentary Group on AI
House of Lords

Rt Hon John Glen MP
Member of the UK Treasury Committee
House of Commons

Dean Nash
Global Chief Operating Officer (Legal)
Santander

Paul Loftus
General Counsel
St James's Place

Arman Fallah
Chief Risk Officer (UK)
stripe

Ben Ellis
Chief Compliance Officer (Revolut UK Bank)
Revolut

Iain Laing
Chief Risk Officer
monzo

Cosette Reczek
Global Head, Model Risk, Markets
standard chartered

Suzanne Brink
Head of Responsible AI
LLOYDS BANKING GROUP

Ratul Ahmed
Global Head of Model Risk Management & Validation
COMMERZBANK

Mitch Trehan
Chief Compliance Officer
Allica Bank

Willem Wellinghoff
UK Chair and Chief Compliance Officer
ecommpay

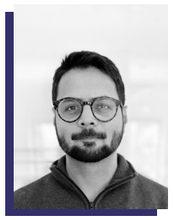

Archit Chamaria
Chief Data and Analytics Officer
novobanco

Rob Phillipson
Managing Director
ICA International Compliance Association

Andrew Sutton
Visiting Fellow
Oxford Martin School, University of Oxford

Dr Alessio Azzutti
Lecturer in Law & Technology (FinTech), University of Glasgow
University of Glasgow

1
Foreword by Lord Clement-Jones

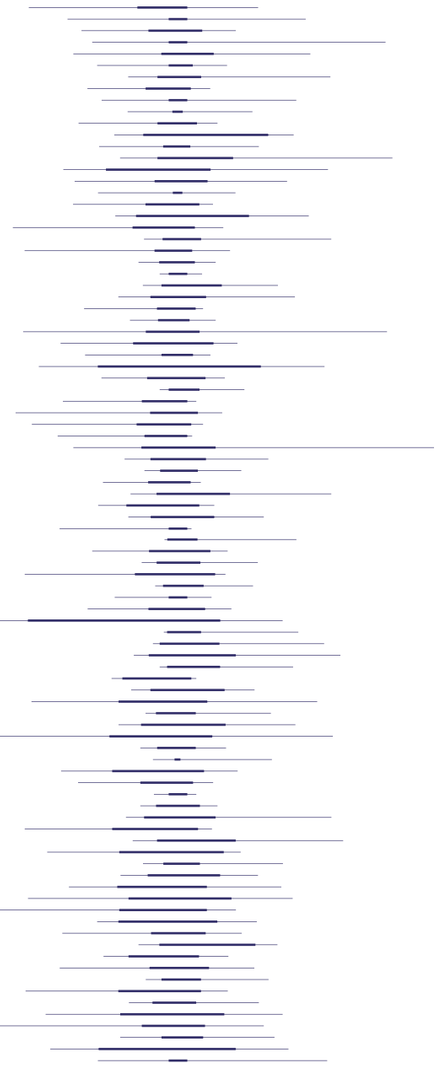

We are currently witnessing a disruptive period of technological change. Artificial intelligence is fundamentally reshaping the financial services landscape, moving far beyond experimentation and piloting into real-world deployment at scale.

The potential benefits of this transformation are undeniable. AI is optimising routine tasks, bolstering fraud detection, enhancing anti-money laundering capabilities, and enabling personalised financial guidance. In the UK alone, three-quarters of financial firms are already utilising AI in some capacity.

However, we need to confront a persistent misconception. The allegation that regulation stifles innovation needs to be turned on its head. In fact, well-designed, proactive regulation is the bedrock of sustainable innovation. Without robust governance and public trust, long-term technological adoption will fail.

As this timely report from Zango powerfully illustrates, the financial services industry is currently deploying AI faster than it can govern it. There is a widening capability gap between the teams deploying these advanced tools and the oversight functions tasked with managing their risks. This is a real source of vulnerability, whether in terms of security or ethics. Generative AI and emerging agentic systems, which produce context-dependent, probabilistic outputs rather than fixed, easily validated scores, complicate oversight immensely. We are no longer just validating static mathematical models; we are expected to govern variable, autonomous behaviour.

In the UK, frameworks like the Senior Managers and Certification Regime (SM&CR) and the Consumer Duty are designed to ensure that the ultimate responsibility for outcomes, whether AI driven or not, rests firmly with human leaders. You cannot outsource your regulatory obligations or liability to an algorithm. Yet senior leaders and oversight functions frequently lack the essential AI literacy needed to be able to challenge these opaque “black box” systems. Independent challenge risks becoming impossible if risk and compliance professionals cannot interrogate how these models truly behave.

Lord Tim Clement-Jones
Liberal Democrat Spokesperson
for Science, Innovation
and Technology, House of Lords

Co-Chair, All-Party Parliamentary
Group on Artificial Intelligence

Transparency and explainability by vendors and deployers are crucial. If consumers are denied credit or offered a higher insurance premium due to an algorithmic decision, they deserve a clear, understandable explanation. Furthermore, we must combat the data biases that can inadvertently lead to unlawful discrimination and financial exclusion for minority or vulnerable groups. We must also be careful that automated systems don’t all react to unusual patterns in exactly the same way, creating a ripple effect that threatens the whole financial system.

What is immediately missing is the translation of high-level regulatory principles into day-to-day operational practice. We cannot simply wait for the aftermath of the first major AI-fuelled financial scandal to force us into action. The financial services industry must take the initiative to build shared, sector-specific implementation guidance to operationalise AI governance effectively.

We have a distinct choice before us: we can deliberately build towards a future of transparent, ethical, and highly effective AI, or we can arrive there only after a crisis we could have easily avoided. I commend Zango for shining a much-needed light on these critical implementation gaps. It is now up to all of us - policymakers, regulators, and industry leaders - to ensure that AI in financial services is governed not by chance, but by conscious, ethical design.

CONTENTS

2
Executive summary

Key findings from interviews with senior leaders across financial services

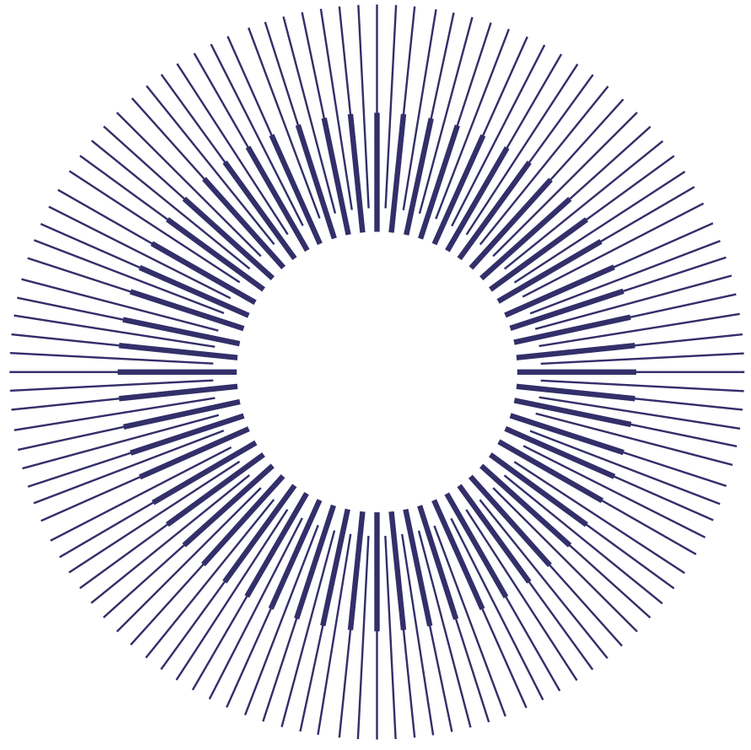

The financial services industry is deploying AI faster than it can govern it. Without action, this gap will widen.

This report draws on qualitative interviews with 27 C-suite and senior leaders responsible for risk, compliance, legal, and AI governance across UK and European regulated financial institutions, including banks, fintechs, payments firms, wealth managers, and digital asset companies. Insights were further developed through four focus groups with interview participants and four industry roundtables involving 60 additional senior practitioners.

1) **The AI now being adopted in financial services is qualitatively different from what came before, and current deployments are just the beginning of that shift.** With many earlier models the same input produced the same output, whereas generative AI systems produce context-dependent outputs with no single correct answer to test against. While agentic deployment remains limited, firms expect to deploy systems that execute actions autonomously and at scale in the near future. Adoption varies across institutions - shaped by data architecture, risk appetite, business needs and strategy, and organisational maturity - but no institution interviewed regards deeper AI deployment as optional.

2) **The functions deploying AI are moving faster than those responsible for overseeing it, creating an oversight gap that will deepen without intervention.** Adoption is most mature in the first line of defence, where data is structured and the commercial case is clear. The second line is following from a significantly lower base, where deployments tend to be more complex and returns less visible - and may struggle to scale as first-line systems become more autonomous.

3) **AI governance maturity varies widely across institutions, and oversight functions often lack visibility over what is being deployed.** In several institutions interviewed, compliance and risk functions had limited insight into which AI tools were being used across the organisation. Some institutions have addressed this by establishing formal AI governance coordination frameworks across functions including model risk, product governance, and data protection. Many have not, and where deployment is decentralised, governance functions can lose sight of what is being used and how.

4) **AI is transforming how the Three Lines of Defence operate in practice, while raising new governance questions.** First-line teams can increasingly run AI-driven validation of their own systems, shifting the second line's role toward setting standards for automated validation and exercising judgement when escalation is required. Oversight functions that rely on periodic sampling risk falling behind as AI enables continuous monitoring across entire datasets. As oversight functions deploy AI themselves, a further challenge emerges, as the overseer systems must themselves be governed.

5) **AI offers genuine potential to strengthen oversight, but many institutions face a skills deficit in the very functions responsible for governing it.** Continuous monitoring, automated validation and AI-assisted audit could materially improve oversight. But compliance and risk professionals who cannot interrogate how AI systems behave cannot realise that potential. Without the technical literacy to examine model behaviour, independent challenge risks becoming independent observation.

6) **Insufficient AI governance creates risks that extend beyond individual firms, both to consumers and to the stability of the financial system.** Weak governance creates vulnerabilities, as institutions with poor oversight of their own AI systems are less able to detect manipulation, adversarial inputs or model failures - and criminal organisations are already deploying AI to exploit these gaps at scale. Conduct failures that once took years to accumulate could now occur in weeks. Across this research, practitioners expect the regulatory framework to be shaped less by proactive standard-setting than by something going wrong first.

7) **A key gap is the AI governance implementation layer that translates regulatory requirements into practice - and industry should build it.** AI governance will always be institution- and deployment-specific, but a common baseline would reduce duplication and stop firms solving the same problems independently. The US and Singapore have already moved to fill this gap through public-private collaboration. The UK and EU have not. Practitioner-built, sector-specific guidance, developed with regulator engagement and modelled on what the Joint Money Laundering Steering Group (JMLSG) has delivered in financial crime, is both achievable and necessary. The future of AI governance will be shaped either by the industry's deliberate choices, or by the aftermath of its first major failure.

Taken together, the findings point to two structural governance gaps emerging as financial institutions deploy AI - and a clear industry response.

CAPABILITY GAP

Oversight functions often lack the skills, tools and infrastructure needed to govern AI at the pace and scale of business deployment.

IMPLEMENTATION GAP

No authoritative sector-specific standard exists in the UK or EU translating regulatory requirements into operational AI governance practice.

INDUSTRY RESPONSE

Shared AI governance implementation guidance

Practitioner-built, sector-specific guidance for the UK and EU - developed with regulator engagement - could provide the common operational baseline firms currently lack, including shared standards for capability assessment, risk classification, and control mapping.

3
Research design & methodology

How the evidence base was built

Methodology overview

Regional coverage
Primarily UK institutions with broader European representation

Institution types represented
Banking • Fintech • Payments
Wealth management • Digital assets

Fieldwork period
December 2025 — April 2026

Participants were selected to reflect diversity of institutional type and role rather than to achieve statistical representativeness. Most interviewees were Chief Compliance Officers or Chief Risk Officers, alongside senior leaders responsible for AI governance, model risk, or legal functions.

The research design and interview framework was developed with input from independent academic advisors with expertise in AI governance, including researchers affiliated with the **University of Glasgow** and the **Oxford Martin School AI Governance Initiative**.

Additional practitioner roundtables were conducted with participants who generally had not taken part in the interviews. These were held under the Chatham House Rule and not recorded, and were used to test and validate themes emerging from the interviews.

Interviews were semi-structured, lasted 30-90 minutes, and were conducted on a non-attributable basis. Quotes are anonymised and attributed by role and institution type. Where the report draws conclusions in its own voice, these reflect the research team's analysis of the evidence.

RESEARCH LEAD PERSPECTIVE

Sam Green
Senior Policy & Partnerships Lead

zango

“This research set out to understand how financial institutions are integrating AI into their businesses, and how they are approaching the governance and compliance challenges that accompany it.

From the interviews, it was telling that even among firms with the most mature AI governance frameworks, practitioners were left asking whether their approach is sufficient - and curious about how their peers are handling the same challenges.

That shared uncertainty is understandable. The governance challenge posed by this technology is new for everyone - businesses, regulators, and governments alike. Definitive answers to difficult questions are hard to find as the landscape continues to evolve.

Practitioners are engaging with these questions seriously and in good faith. But they are largely doing so alone. There is a real opportunity and need to bring that collective expertise together into shared standards before circumstances force the issue.”

4
Framing the governance landscape

Defining AI governance and its regulatory context

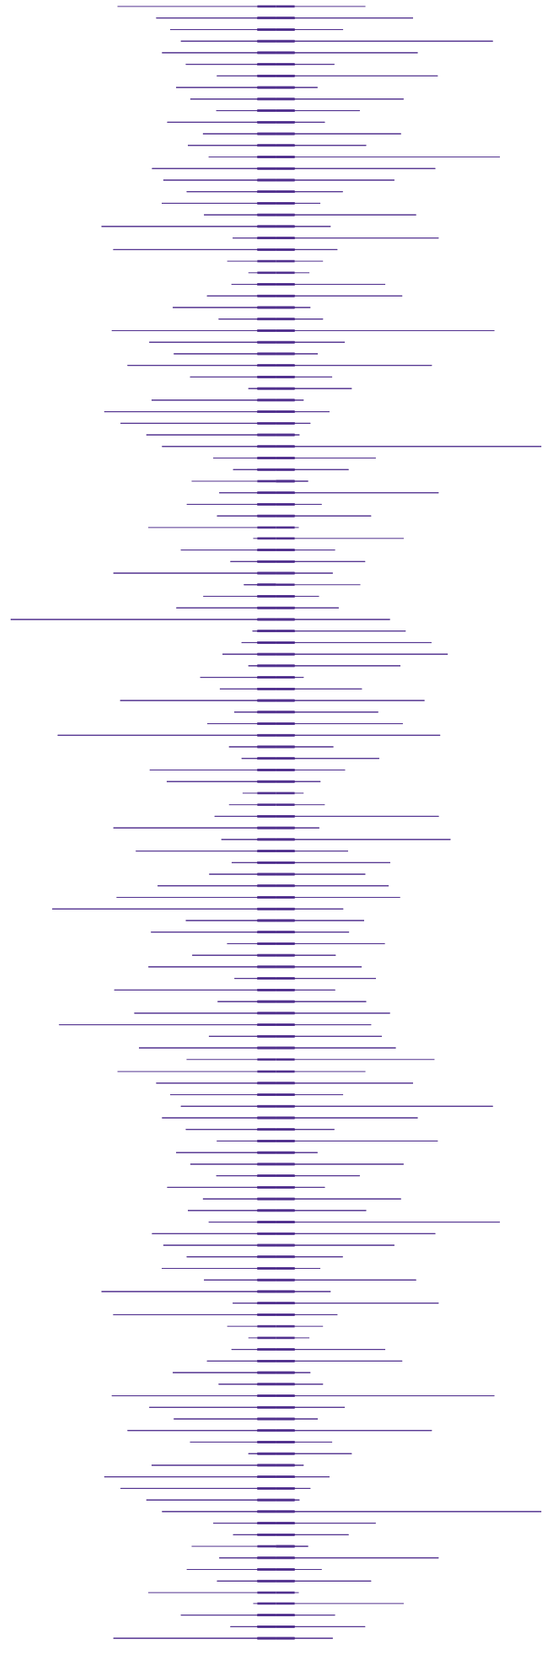

This section sets out how this report uses the term AI governance, and provides brief context on the regulatory landscape within which the institutions interviewed are operating.

Definition: AI Governance

The overarching framework of policies, roles, processes, and controls designed to ensure that AI systems are designed, developed, deployed, and operated safely, responsibly, and in compliance with regulatory requirements and internal standards.

In financial services, this means:

- ensuring clear accountability across the Three Lines of Defence
- managing operational, conduct, regulatory, and ethical risks arising from AI systems
- protecting consumers and market integrity
- ensuring AI-supported processes remain controlled, reviewable, and defensible under supervisory scrutiny

It is useful to distinguish three overlapping components of what AI governance actually comprises in practice:

Component	What it covers
GOVERNANCE STRUCTURES	The organisational arrangements through which accountability is exercised - including senior ownership, and committee oversight.
GOVERNANCE FRAMEWORKS	The policies, processes, and controls governing how AI systems are assessed, validated, and managed - including model risk, product governance, and supplier frameworks.
GOVERNANCE CULTURE AND CAPABILITY	The degree to which the people responsible for oversight are equipped and willing to exercise it effectively.

Divergent regulatory approaches

Approaches to AI regulation are diverging. While the UK has opted for context-based, sector-led regulation, the EU has chosen a horizontal statutory framework.

The Regulatory Landscape: UK

In the UK, AI is primarily governed through existing regulatory frameworks rather than a dedicated rulebook. Financial regulators have adopted a principles-based approach consistent with the UK Government's wider AI regulatory strategy.

Key regimes relevant to AI deployment include:

- **Model risk management** - the PRA's SS1/23 supervisory expectations for firms using internal models.
- **Operational resilience** - PRA and FCA requirements apply to technology-enabled services.
- **Conduct and governance** - the FCA's Consumer Duty, Senior Managers and Certification Regime (SM&CR), and existing conduct rules apply to AI-enabled decision-making.
- **Data protection** - UK GDPR governs personal data processing, including automated decision-making and profiling, which can constrain certain AI uses.

Additional sector-specific guidance, including joint FCA/PRA work on AI and machine learning (DP5/22) and rules governing algorithmic trading, (SS5/18), may apply depending on business model and activity.

The Regulatory Landscape: EU

In the European Union, financial institutions operate under multiple existing regulatory regimes relevant to AI deployment, including:

- **Data protection** - GDPR governs automated decision-making and data processing.
- **Digital operational resilience** - DORA establishes ICT risk management and resilience requirements.
- **Sectoral regulation** - governance expectations are embedded across existing financial services frameworks.

Overlaying these frameworks, the EU AI Act introduces a dedicated cross-sector regulatory regime for AI systems. The Act establishes a risk-based classification framework and imposes enhanced obligations on high-risk systems, including certain financial services applications such as credit scoring.

For financial institutions, the AI Act supplements rather than replaces existing regulatory expectations, adding explicit AI governance requirements alongside prudential, operational and conduct regulation.

The multi-jurisdiction challenge

Firms operating across both jurisdictions must therefore meet specific EU requirements while demonstrating how broader UK regulatory principles are satisfied. Understanding why those requirements take the form they do requires closer examination of what makes AI systems distinct from the technologies financial institutions have historically governed.

EXPERT COMMENTARY

Rt Hon John Glen MP

Member of the UK Treasury Committee
House of Commons

“The future leadership of the financial services sector, perennially dubbed the ‘crown jewel’ of the British economy, depends on our ability to solve the persistent productivity puzzle. As I observed during my tenure at the Treasury, technology is an essential lever for doing more with less, provided we move towards swift implementation.

AI offers extraordinary capabilities to drive the efficiency and competition necessary for growth. However, this potential can only be realised through a truly joined up regulatory approach across HM Treasury, the PRA, and the FCA. The regulatory framework I helped to develop post-Brexit through the Financial Services and Markets Act 2023 was designed to provide agility, replacing rigid, inherited rules with a nimble, UK-specific system.

Just as with the post-Brexit reforms, AI presents an opportunity to apply gold standard accountability frameworks to manage emerging risks without stifling innovation. With the right regulation and attitude, we can transform the challenges of AI into huge opportunities for global competitiveness.”

5
Why this technology is different

Governance assumptions under pressure

Financial institutions have long governed advanced analytics and machine learning systems through established risk management frameworks. Those frameworks assume that system performance can be tested and validated to provide sufficient confidence for deployment, even where underlying behaviour is not fully interpretable. This report examines the deployment of generative AI and agentic systems - technologies that challenge these assumptions.

The nature of the shift

Earlier AI deployments were typically bounded in practice. Models produced outputs within defined categories or numerical ranges, and their behaviour could be tested against historical data through established validation techniques. Where automated decisioning was used - in credit scoring, fraud detection, or risk modelling - the underlying logic could be reasonably documented and audited.

Generative AI changes these assumptions. These systems produce new content - such as text, code, or images - rather than classifying inputs or generating scores. Built on foundation models trained on vast datasets and adapted through prompting or fine-tuning, their behaviour cannot be fully specified in advance, shifts as they ingest new data, and in many cases produces no single correct answer against which to validate performance.

This shifts the governance challenge from validating fixed outputs to governing variable behaviour.

FEATURE	TRADITIONAL ML (DISCRIMINATIVE)	GENERATIVE AI	AGENTIC AI
OUTPUT	Scores, categories, labels	Text, code, media	Actions and tool use (API calls, system tasks)
VALIDATION	Backtesting against ground truth	Scenario-based evaluation, human review	Runtime controls, guardrails, monitoring
HUMAN ROLE	Output quality oversight and review	Reviewer / validator (where required)	Supervisor of system actions (where required)

Foundation model adoption remains early – but is accelerating

Foundation models remain a small share of deployed AI models. The 2024 Bank of England and FCA survey found they account for 17% of AI use cases overall, and just 7% in risk and compliance.[1]

Adoption figures should be read with some caution. Survey data in this area reflects definitional choices, as what counts as "AI" varies between studies and between firms, and self-reported figures may not capture the full picture of deployed systems.[2]

Foundation models as a share of all AI models, by business area (%)
Source: Bank of England / FCA, AI in UK financial services survey, 2024

These figures reflect adoption at the time the survey was conducted. Given the pace of development in generative AI since then, the share of deployments is likely to have increased. Qualitative evidence from interviews conducted for this research supports that trajectory, with institutions reporting growing experimentation with foundation models and expanding generative AI use cases.

[1] Bank of England and FCA (2024), Artificial Intelligence in UK Financial Services 2024.

[2] The 2024 survey was the first to report foundation models as a distinct category. These, which encompass large language models and generative AI systems, are the category most relevant to the systems this report examines.

Agentic AI escalates the stakes further

The sector's governance frameworks are already being tested by current deployments - and the next wave of adoption will be more complex to govern, particularly with the emergence of agentic systems.

Definition: Agentic AI

Agentic AI refers to AI-enabled systems capable of autonomously planning and executing multi-step tasks - accessing data, using tools, and taking actions - to achieve defined objectives with limited human intervention.

The governance implications are significant. While many current deployments retain human review, agentic systems can execute actions directly, and as autonomy increases, the challenge shifts from assessing outputs to governing actions taken on those outputs.

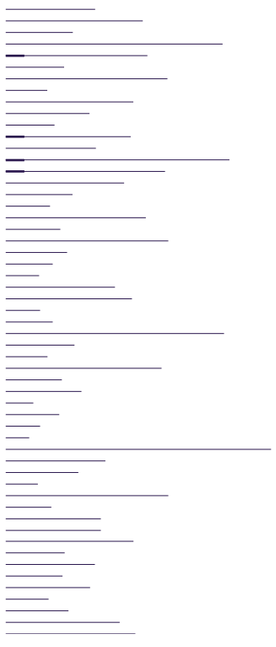

EXPERT COMMENTARY

Dean Nash
Global Chief Operating Officer (Legal)

Santander

"Senior managers in regulated finance are legally accountable for ensuring their areas of responsibility are controlled effectively. Traditional governance assumes that robust policies and deterministic controls ensure manageable outcomes; the manager signs off because the process is bounded and predictable.

Generative AI challenges this approach. Because Gen AI models are probabilistic, context-sensitive, and inherently non-reproducible, they create a "responsibility gap" where legal liability outpaces a manager's technical visibility. When a system's output cannot be guaranteed, tracing data lineage and logic through a traditional audit trail becomes much more difficult.

Closing this gap requires a fundamental rethink of governance architecture. We must shift from controlling and auditing a system's internal logic to governing the dynamic environment in which it operates. In a non-deterministic world, risk management is no longer about predicting every result, but about orchestrating an ecosystem of continuous, real-time guardrails. In other words, this involves managing the integrity of data inputs, prompt structures, and human-in-the-loop interventions to bridge the accountability gap."

6
AI adoption and oversight

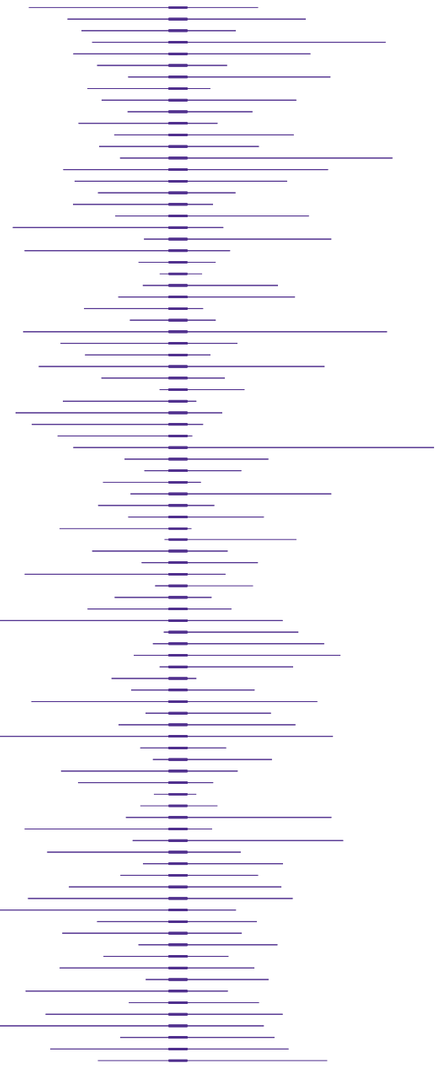

Where AI is being adopted - and where governance is not keeping up

AI adoption in UK financial services is accelerating. The 2024 Bank of England and FCA survey found that 75% of firms are already using AI - up from 58% in 2022.[3] The European Banking Authority's 2025 survey showed 92% of EU banks are using AI, and 55% of banks are already using general-purpose AI or agentic AI in consumer-facing processes.[4] But these headline figures obscure significant variation in what adoption actually looks like in practice.

Adoption is real, uneven, and accelerating

Across the institutions interviewed for this research, AI adoption varies widely - from digitally native firms deploying AI at production scale to senior risk leaders with little or no direct contact with it in their day-to-day work. Most sit somewhere in between, with a growing pipeline of use cases and a shared sense that adoption is no longer optional.

How firms are approaching AI adoption differs significantly - from top-down mandates to conscious caution.

"The CEO has mandated that every single department adopts AI."
— Chief Compliance Officer, major UK fintech

"We never intended to be pioneers on AI. We let the others be pioneers. So it's step by step."
— Senior risk, legal and compliance leader, international investment bank

At the leading edge, adoption is already substantial. At one financial institution, around a fifth of engineers no longer write code, instead supervising AI assistants that generate it.

Several firms described an intermediate position - no longer experimenting, but not yet deeply embedded:

"We're much more in that implementation phase rather than the kind of deeply embedded phase. But I can see just how that will move quite quickly from one to the next."
— Senior risk and compliance leader, global payments firm

Across institutions, interviewees consistently described economic pressure as a key driver: firms are deploying AI to capture efficiency gains even where governance frameworks are still evolving.

"Even though they may not have perfected their model to do the task, I'm pretty sure they will want to capture those cost savings upfront."
— Chief Compliance and Risk Officer, major payments firm

[3] Bank of England & Financial Conduct Authority (2024). Artificial intelligence in UK financial services - 2024.
[4] European Banking Authority (2025). Rising application of AI in EU banking and payments sector.

The Three Lines of defence Model

The Three Lines of Defence Model is a widely adopted governance framework, originating in internal audit and now embedded in financial services risk management.[5]

- **First line** - Business and technology teams that design, deploy, and operate systems, and own day-to-day risk management.
- **Second line** - Risk, compliance, and control functions that set policies and standards, provide oversight and challenge, and monitor adherence.
- **Third line** - Internal audit, which provides independent assurance that controls and governance processes are effective.

The first line pulls ahead

AI deployment is most mature in the first line of defence, where data is structured, workflows are repeatable, and the return on investment is immediate. Fraud detection, transaction monitoring, KYC screening, customer service, and code generation are the most mature use cases across the sample.

Second-line adoption - including AI-driven policy gap analysis and controls testing - is gaining ground but from a significantly lower base. In several firms, compliance and risk teams acknowledged that AI had not yet significantly changed how they work.

"I still think we're at the early stage of using it as a second line team. I wouldn't say it's a core part of people's ways of working yet."
— Chief Risk Officer, UK fintech

Part of the issue is that second-line processes have been designed around human workflows, making straightforward AI adoption more complex than in the first line.

"When you take AI and you try to apply it to a very human-centric process, it actually doesn't work very well, because the process has been designed by humans. When you then just kind of retrofit AI onto the top of it — it doesn't quite work."
- Senior compliance leader, major payments firm

The third line is at an earlier stage still. Internal audit functions have historically been the most removed from day-to-day AI deployment, and in most institutions interviewed their adoption remains relatively limited. The gap between where AI is being deployed and where oversight of it sits is widest here, and is likely to widen further as first-line systems become more complex.

Consequently, the functions deploying AI are moving faster than the functions responsible for overseeing it.

[5] Institute of Internal Auditors (2020). The IIA's Three Lines Model: An update of the Three Lines of Defense.

The maturity of AI use cases in compliance [non-exhaustive]

FIRST LINE OF DEFENCE	SECOND LINE OF DEFENCE	THIRD LINE OF DEFENCE
● Fraud detection	● Horizon scanning	● Audit planning & risk scoping
● Sanctions screening	● Obligation extraction	● Audit testing
● Transaction monitoring	● Policy & controls gap analysis	● Agentic audit assistants
● Customer complaint triage	● KYC QA/QC	● Real-time audit monitoring

● Embedded/Scaled ● Piloting/Partial ● Immature

The infrastructure divide

The depth of adoption also varies by institution type, and data architecture is a core dividing line. Cloud-native firms describe a fundamentally different starting position, with unified systems, clean data, and AI-ready infrastructure. For legacy institutions, fragmented architecture makes meaningful deployment difficult.

"We have a clean, simple data fabric - so we're in a position to just apply AI. In other organisations, you'd have to fundamentally redesign the tech stack first."

— Chief Compliance Officer, UK digital bank

The consequence is that for most legacy institutions, meaningful AI transformation will remain bounded for some time. While data may exist in sufficient quantity, its fragmentation across dozens of interdependent systems limits what AI can actually do with it.

"These companies are neither data-first nor AI-first - it's not going to be a meaningful transformation across the board. It will be very specific verticalisations."

- Group Chief Compliance and Risk Officer, major payments firm"

EXPERT COMMENTARY

Arman Fallah
UK Chief Risk Officer

stripe

"The financial ecosystem is transitioning from "digitised" to "digital-native," driven in large part by agentic and generative AI. This creates both opportunity and obligation for risk and compliance functions.

Data architecture is a decisive factor in AI readiness. Cloud-native institutions can more readily build scalable model-driven detection that rapidly embeds intelligence directly into transaction flows. Legacy institutions with fragmented data architectures face a structural disadvantage: siloed systems impede the real-time, unified data access on which effective AI-driven risk management depends.

However, speed of adoption must not outpace oversight. Responsible AI usage operates with human-in-the-loop validation and accountability. Like all good governance it should be risk-based; higher-risk uses of AI require stronger controls, review, and approval.

As the industry moves toward agentic, high-velocity workflows, the challenge is not whether to adopt AI, but whether governance frameworks can evolve at the same pace - calibrating controls to risk while preserving the agility that makes AI transformative."

KEY INSIGHT

How legacy institutions are approaching adoption despite infrastructure constraints

Legacy financial institutions with fragmented data architecture cannot always pursue the same AI deployment strategies as cloud-native firms. Among those institutions with relatively more mature AI deployments, a common pattern is emerging.

- **Narrow initial use cases:** Early deployments are concentrated in processes where data is already structured or accessible (such as transaction monitoring) rather than attempting organisation-wide transformation.
- **Use cases that rely primarily on external data:** Some early deployments focus on tasks where external rather than internal data is the primary input.

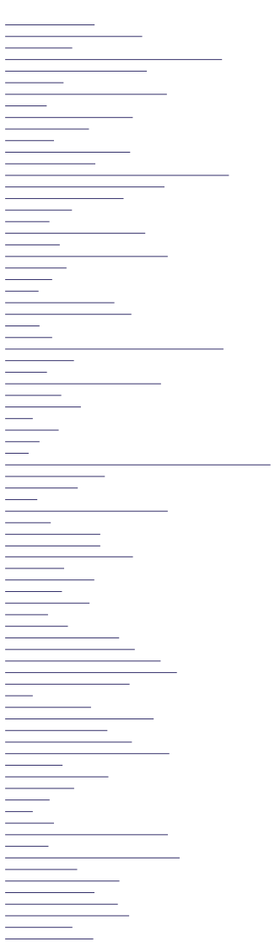

The agentic shift

Agentic AI is generating significant interest, but production deployment remains limited. Most institutions are still experimenting internally, where consequences of errors are contained and human review is straightforward to maintain.

Yet a growing number of institutions are actively scoping customer-facing use cases - in areas including customer service, onboarding, and personalised financial guidance. Some banks have already committed publicly to a customer-facing agentic deployment, and others are not far behind.

As agentic AI increasingly moves from experimentation to deployment, the question of who owns it - and who is accountable when it fails - becomes more urgent.

EXPERT COMMENTARY

Archit Chamaria
Chief Data and Analytics Officer

novobanco

“Financial services has historically not been a leader in AI adoption - and for good reason. The role our industry plays in society means that adverse outcomes can carry serious consequences for society. But I think that calculus has fundamentally shifted. Institutions that continue to move slowly risk becoming targets.

At Novobanco, we are working to ensure AI is adopted by colleagues and enhances enterprise journeys and processes. At the enterprise level we see applications across 3 main areas. The first is customer-facing: moving towards genuine hyperpersonalisation, where AI allows us to deliver a higher level of service with a laser focus on security and applying appropriate guardrails. The second is the middle office - the human-machine interface - where machines are now working alongside humans to manage entire processes that were previously highly human-dependent. The third, and perhaps most consequential, is proactive risk and compliance. We can now surface regulatory and risk information in context, moving the compliance burden away from data gathering. Across all three areas the role of the humans is not diminished but augmented.

What concerns me most is the cyber dimension. If we are not using AI to counter threats that are themselves AI-enabled, traditional tools simply won't be sufficient. As an industry, we have developed shared standards and ways of working to guard against credit risk, market risk, operational risk - we need to do the same for AI-enabled threats. Without a baseline standard, institutions become individual targets for malicious actors, and the consequences of that ultimately fall on the people and the systems we exist to serve.”

7
Organisational structures for AI governance

Who owns AI?

In regulated financial institutions, accountability for AI does not typically sit in a standalone function. Existing regulatory frameworks allocate responsibility across established governance roles, meaning accountability follows business outcomes and organisational structure rather than the technology itself.

How firms then organise AI governance in practice, and the challenges that distributed accountability creates, is the focus of this section.

AI and SM&CR

In the UK, the Senior Managers and Certification Regime provides the regulatory framework within which AI accountability sits. Under SM&CR, regulatory accountability is assigned to named Senior Management Function holders for specific prescribed responsibilities - SMF4 for risk oversight, and SMF16 for compliance.

No Senior Management Function exists specifically for AI, meaning responsibility is absorbed into existing roles by default. As with other cross-cutting systems, a single AI deployment cuts across risk, compliance, data, and technology simultaneously, meaning accountability is fragmented across multiple SMF holders by the structure of the regime itself.

Accountability follows outcomes, not technology

The dominant approach across the institutions interviewed is to assign accountability based on the outcome the AI system supports, rather than on the technology itself.

"We don't necessarily assign accountability for AI. What we do assign accountability for is the outcomes of AI."

— Chief Compliance Officer, major UK fintech

This approach works within existing SM&CR structures and keeps ownership close to the business decisions being made. For regulated firms already operating structured accountability frameworks, existing governance machinery provides a workable foundation. As one practitioner observed, SM&CR means institutions are already accustomed to assigning clear responsibility, and AI governance can, to a degree, be built on top of that.

But the approach relies on function owners having sufficient AI literacy to discharge that accountability meaningfully - a condition that, as Section 10 sets out, is not yet consistently met. Further, it addresses only the formal assignment of accountability, not the question of whether governance structures are configured to identify and respond to AI risk as it arises across the organisation.

EXPERT COMMENTARY

Ben Ellis
Chief Compliance Officer
(Revolut UK Bank)

Revolut

"Accountability structures are a key element of any framework aiming to ensure AI is safely adopted across firms in the UK. As seen with accountability structures such as SM&CR, creating a system that holds people accountable for their actions and decisions helps ensure decision-makers act in the best interests of consumers and market integrity.

As AI is sometimes seen as a pure technology capability, it is important that accountability structures do not fall into the trap of assigning accountable individuals solely in Tech teams. Where a firm is deploying AI, there needs to be accountability across the firm, from strategy setting to product design to underlying technical maintenance."

Cosette Reczek
Global Head, Model Risk, Markets

standard chartered

“In UK financial institutions accountability for use of AI outputs within business activities ultimately resides within the Senior Manager Functions as prescribed within the UK’s SM&CR regulation. In a broader sense, all employees have accountability, as they must evidence they act with “due skill, care and diligence”, meaning that everyone must understand the impact of using AI outputs within their daily work.

As AI proliferates within the businesses and functions of financial institutions, AI review and approval functions, particularly those created on a group-wide, centralised basis, should evolve to enable identification of risks emerging from AI use across the SMFs as an embedded risk, just like any other risk in the execution of daily business and risk management activities.

To aid this, oversight activities and governance forums can constructively review and challenge that the use of AI outputs in business decisions is fitting and based upon safe and controlled execution. Examples could include consideration of AI outputs in the assessment of new product customer suitability and operational readiness, and also model governance whereby second-line independent review considers the performance of models with AI components to ensure drift, bias and other performance health considerations are actively managed.”

Coordination and visibility challenges

Effective governance depends on visibility. Where AI deployment is decentralised - with individual teams adopting tools independently, sometimes with limited central oversight - governance functions can quickly lose sight of what is being used and how.

“Without visibility you miss things. We've seen a huge range of quality in use cases. Some excellent. Some real shockers. If I were at a bank without central oversight I would be worried.”

- Head of Model Risk, major international bank

A single deployment can cut across technology, data privacy, model risk, compliance, and customer outcomes simultaneously, with accountability distributed across functions and no single owner maintaining a complete view.

"You need a coordinating role around AI governance, because it isn't one thing: it's tech, it's non-tech, it's cyber, it's legal."

- Head of Responsible AI, major UK bank

At institutions with more established governance frameworks, use cases tend to be routed through approval committees, with model risk or responsible AI functions maintaining oversight of the portfolio. At others, the picture is more fragmented - individual teams have adopted tools on their own initiative with limited central guidance, and no single function holds a complete view of what is being used across the organisation.

"If I ask the question - show me everywhere AI is being used across this organisation - I wouldn't be able to get an answer."

- Head of Compliance, regulated financial institution

The emergence of dedicated AI leadership

Dedicated AI leadership roles, such as Chief AI and Data Officers or Heads of Responsible AI, are emerging in response to coordination challenges. These roles rarely hold direct accountability, but instead act as coordination layers across functions. Their emergence reflects recognition that, at least at this early stage of mass adoption, the coordination burden of AI governance exceeds what any single existing function can carry.

“Institutions that federate too quickly struggle to capture value and manage risk.”

- Head of Responsible AI, major global bank

However, the appetite for centralisation varies. At some institutions, AI governance is deliberately disaggregated with principles set centrally, and deployment decisions handled at business line level within defined parameters. The appropriate balance between central coordination and local ownership is one of the more actively contested questions in how firms are structuring their approach.

What is consistent across the sample is that these coordinating roles remain relatively new, and their authority and remit is still being defined. In most institutions, they operate alongside existing functional accountabilities, with limited formal power to compel alignment across the business lines they are meant to coordinate.

Suzanne Brink
Head of Responsible AI

LLOYDS BANKING GROUP

“At Lloyds Banking Group, accountability for AI is structured around the Three Lines of Defence. Within the first line, the Responsible AI team, based in the AI Centre of Excellence, enables safe and scalable adoption through practical guidance, tooling and accelerators that embed risk controls into delivery. The team also acts as a focal point for thought leadership on emerging AI risks, convening cross-functional forums, including on agentic AI, to continually assess whether governance and controls remain fit for purpose.

Clear ownership sits with model and use-case owners, who are accountable for managing AI risks across the lifecycle - from design and approval through to ongoing monitoring - using established decision and accountability frameworks.

From a second-line perspective, the Model Risk Office provides coordinated oversight of AI use cases, bringing together specialist risk functions such as data, cyber, legal and compliance to support holistic risk assessment and independent assurance.”

KEY INSIGHT

How AI governance structures are organised

Formal accountability for AI sits with function owners across every institution in the sample. What varies is which existing function acts as the organisational anchor for AI oversight:

- **Model risk as second-line anchor:** At firms with established model risk functions, model risk has emerged as the primary oversight mechanism for AI - with some actively building out that capability specifically for AI systems.
- **Operational risk as second-line anchor:** At some institutions, AI governance sits under operational risk as a risk class.
- **Technology or COO function as AI strategy owner:** At several firms, AI strategy sits within a technology or operations function.
- **Dedicated AI leadership roles:** Larger institutions - particularly banking groups - are appointing Chief AI and Data Officers at C-suite level alongside existing risk and compliance leadership. Some firms have also established Heads of Responsible AI or equivalent as first-line coordination functions below that.

Emerging pressure points

The organisational response to AI governance is still taking shape. Coordinating roles are being established while remits are being defined, and the balance between central oversight and local deployment is being worked out in practice. Assigning accountability and defining organisational structures is a necessary starting point, but whether the frameworks through which AI is governed day to day are keeping pace is a separate question. That is the focus of the next section.

8
AI governance frameworks in practice

How institutions are building and adapting their governance frameworks

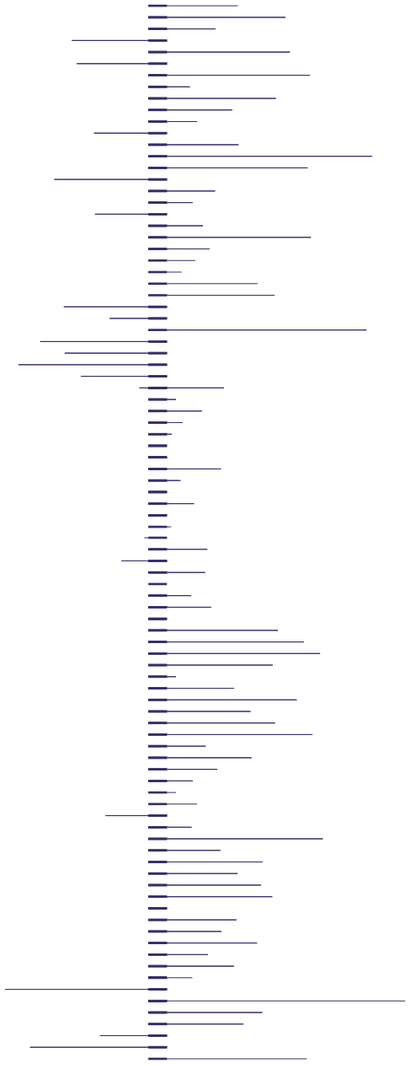

While dedicated AI leadership roles are emerging to address coordination challenges, governance itself is exercised primarily through existing risk and compliance frameworks rather than through structures built specifically for AI. Some institutions have layered a formal AI governance framework on top of these, but many have not. Whether those frameworks are ultimately sufficient for the characteristics of modern AI systems remains an open question.

AI governance frameworks: a fragmented picture

The maturity of formal AI governance frameworks varies widely. Some firms have established dedicated programmes with defined processes for assessing, escalating, and monitoring AI use cases. Others have no overarching framework at all - governing AI instead through fragmented extensions of existing processes, with no consistent standard applied across deployments.

"We have an AI risk framework... that explains how the different risks related to AI are managed in their natural home. It also describes some more AI-focused ways of working and some requirements for AI-specific governance and oversight."

— Head of Responsible AI, major global bank

"There are pieces of how to use it — but there's no holistic AI risk management framework yet."

— Group Chief Compliance and Risk Officer, major payments firm

At one financial institution, the absence of any overarching framework had produced a fragmented picture: individual teams deploying AI within their own parameters, without consistent metrics, escalation paths, or testing standards.

"I would feel much more comfortable if we had an overarching governance... this is the strategy, this is what governance looks like, this is how you show good governance at the use case level, these are the metrics required."

— Chief Compliance Officer, financial institution

How firms are operationalising AI governance

Across the institutions interviewed, several broad governance patterns emerged. These reflect different ways firms are integrating AI into existing governance frameworks - including model risk, product governance, vendor management, information security, and data governance. Many institutions combine elements of more than one model.

MODEL	CHARACTERISTICS
Principles-based with committee escalation	Firm-wide AI principles set centrally; business areas deploy within those principles without requiring central approval; deployments above defined thresholds escalate to relevant governance committees
Model risk-anchored	AI deployments enter governance through established model risk processes; independent validation assesses robustness and performance; escalation determined by materiality; the model risk function coordinates input from data, legal, security, and compliance
Product governance as primary vehicle	AI-enabled products and processes route through standard product governance committees; AI-specific checkpoint items - covering test results, output boundaries, and human-in-the-loop requirements - are embedded within existing product rollout processes
Operational risk-anchored	AI governance sits under operational risk as a risk class; first line completes a structured scorecard assessment for each use case; second line scrutinises the scorecard and owns oversight
Data governance-modelled	AI governance mirrors the firm's existing data privacy governance structure; AI deployments route through a responsible AI committee that follows the same logic as GDPR governance; AI and data governance formally connected
Decentralised / early-stage	No formal overarching AI governance framework in place; individual teams deploy within high-level principles if any exist; no consistent escalation paths, metrics, or testing standards across the organisation

How existing frameworks are being adapted

Where AI governance exists as a formal structure, it typically sits above and coordinates existing processes. Where it does not, AI is being absorbed into existing frameworks directly. In both cases, it is extension rather than reinvention.

"We're not creating new governance per se - we've already got governance in place, it's just being enhanced to accommodate the AI lens."

— Chief Compliance Officer, major UK fintech

Practitioners broadly rejected the idea that AI requires a standalone governance structure separate from existing risk management:

"I don't think there should be separate AI governance. It should all be part of the same thing."

— Chief Risk Officer, global payments platform

Where existing frameworks strain

While most institutions are extending existing governance frameworks, several practitioners questioned whether this approach will ultimately prove sufficient. Frameworks designed around systems with stable, bounded behaviour were not built for probabilistic outputs and rapid iteration cycles.

"The analogue thinking of traditional risk and compliance oversight does not translate well into AI. I find it hard and not fully compatible."

— Chief Risk Officer, UK financial institution

One place where this strain is particularly visible is in the way in which the Three Lines of Defence model is currently operationalised in practice - the primary structure through which financial institutions distribute oversight responsibility. How AI is reshaping that model is the subject of the next section.

KEY INSIGHT

What "extending existing frameworks" actually looks like

In practice, “extending existing governance frameworks” involves adding new control layers and documentation requirements to existing processes:

- **Within model risk:** Validation processes are supplemented with AI-specific documentation requirements - covering training data provenance and performance monitoring calibrated to probabilistic outputs. As model development cycles accelerate, some institutions are beginning to automate parts of documentation generation to keep pace.

- **Within product governance:** Where AI affects customer-facing processes, standard rollout processes are supplemented with AI-specific checkpoint items - such as test results, output boundaries, and human-in-the-loop requirements - before deployment approval.

What tends to go wrong is having no clear picture of where the gaps between frameworks lie, and no way of knowing when an AI deployment has fallen through them.

EXPERT COMMENTARY

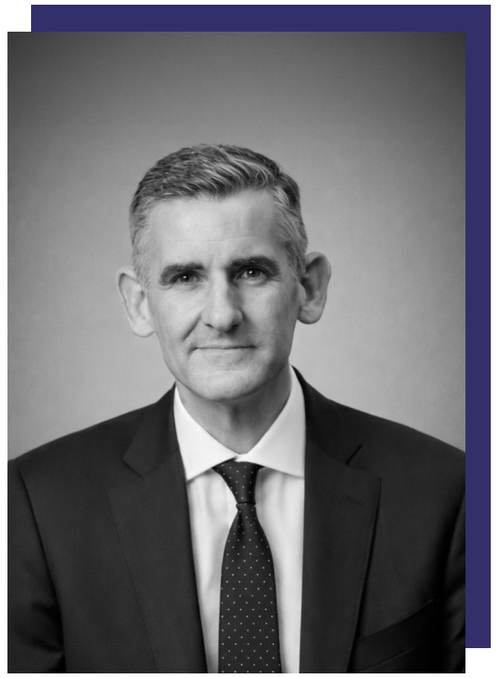

Paul Loftus
General Counsel

St James's Place

“The core role of a General Counsel and a legal team is to identify, manage and mitigate risk. But to do that well, one needs to horizon scan for new and emerging risks and ways to address that novel threat. AI sits at the nexus of these complementary roles.

AI will undoubtedly revolutionise processes and drive efficiencies within firms. At SJP, we believe agentic AI will complement human led process and, if interoperability is achieved carefully - as we are working to do - the benefits for our clients, our firm and the wider economy will be significant.

The role of the second line and the legal teams then becomes one of ensuring that the AI is working as it should and that its ingestion leads to good customer outcomes, and that it is deployed in the right circumstances, operates consistently, has appropriate safeguards and remains a servant of the firm and accretive to risk management.

This requires us to think differently about our client journeys and engagement models. To reimagine how we gather information, assess suitability, compare against the broader market and make decisions.

The output of this reimagination is a renewed governance framework. There is no single source for how we construct these frameworks. It is for lawyers and risk professionals to distil their collective experience into a new model for compliance.

Constant consideration, assessment and modification is the only way risk management suitable for AI can be delivered.”

9
How AI Is reshaping the Three Lines of Defence

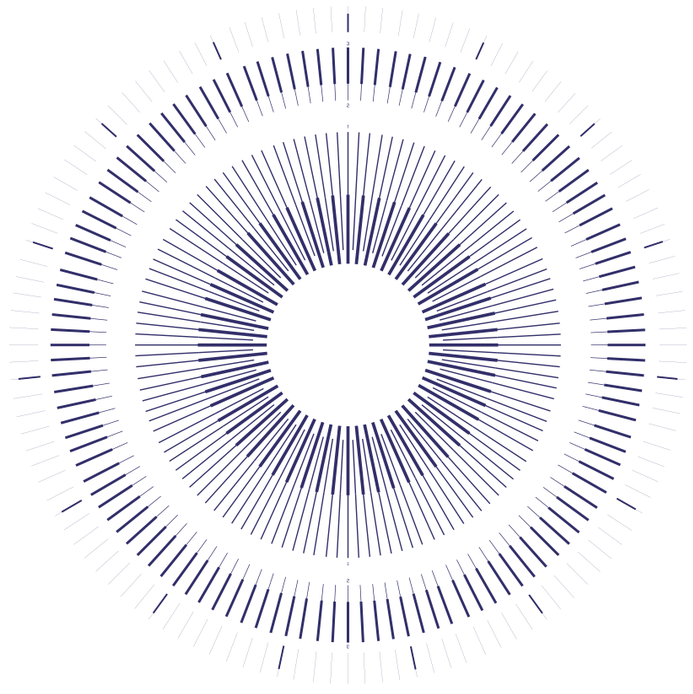

Same framework, evolving practice

The practitioners interviewed were confident that the Three Lines of Defence model remains an appropriate framework for governing risk in regulated financial institutions. What AI changes is what operating that framework effectively requires in practice.

Historically, independent challenge depended on the ability of the second and third lines to review - and where necessary independently validate - what the first line had built. As AI systems become more complex, probabilistic, and embedded in operational processes, the methods through which each line discharges its responsibilities must adapt. At the same time, AI is creating new possibilities for how oversight and assurance functions operate.

When the first line becomes its own judge

One emerging development is the use of AI systems to evaluate the outputs of other AI systems. In this approach - sometimes described as “LLM-as-a-judge” - a second model evaluates the outputs of the first, assessing for errors, bias, or performance degradation. In many cases this validation can now be run directly within the first line.

"Can I build another large language model as a judge in the second line? Why would I do that? I would ask the first line to do it."

— Head of Model Risk, major European bank

Domain expertise, access to the system architecture, and operational context often sit with the teams that built the model. As a result, the first line may be able to run validation processes - including AI-driven evaluation - at a depth and speed the second line cannot replicate.

“We’re seeing more and more QC and QA happening in the first line.”

— Chief Compliance Officer, UK SME bank

What the first line is performing in this scenario is internal quality assurance rather than independent challenge. But for increasingly complex AI systems, it may also be better placed to perform that validation.

Rethinking the second line: from model oversight to control oversight

The second line's role has historically been defined by direct engagement with what the first line builds: inspecting models, testing outputs, and challenging assumptions. As first-line teams increasingly run AI-driven validation of their own systems, the basis of that independent challenge is shifting - away from direct model validation and toward setting the standards by which automated validation is conducted, and maintaining the judgement to escalate when it fails.

"The first step in validating a model could now be asking another model what faults it sees. The first line could do that themselves... That could put the second line into a mode of control oversight rather than model oversight. Instead of validating the model ourselves, we check whether the first line's automated validation process is working properly. That would be quite a difference."

— Senior risk leader, UK financial institution

In practice, this means the second line is increasingly focused on defining standards for automated validation, ensuring monitoring systems are functioning as intended, and escalating issues that require human judgement. This conception of independent challenge requires different skills, and a different relationship with the first line.

EXPERT COMMENTARY

Iain Laing
Chief Risk Officer

monzo

“AI will fundamentally change ways of working in firms, and that will impact the sector and society. Good governance is technology agnostic, but AI will require firms to adapt fast.

The Three Lines model, for example, separates duties and ensures everything meaningful to the risk of the firm has an independent pair of eyes on it. Generative AI doesn't undermine that, but it changes how it works in practice.

In the past, models often produced binary or numerical outputs. Those could be back-tested against real outcomes, but AI models are more complex. Commonly they generate sequences of decisions or strings of text. Overseeing generative AI can be more like governing human judgement and work rather than just testing a prediction against reality.

This technology is also shaping businesses. First line colleagues can validate AI models themselves by using other models. That will change how risk and control functions work, with more emphasis on verifying automated control environments than testing each action or each model. Successful firms will adapt fast to make use of these technologies, while preserving the critical benefits of independent scrutiny.”

From sampling to full coverage

The same technologies reshaping the first line are beginning to transform oversight functions themselves - albeit at a slower pace and from a lower base. Traditional compliance monitoring and internal audit have historically relied on sampling methodologies, which inherently leave gaps in coverage. AI changes this constraint, as entire datasets can be analysed in real time, enabling oversight functions to detect emerging issues far earlier than periodic review cycles allow.

“Our sample sizes are larger... because we’re using AI to do the data analytics over it.”

— Chief Compliance Officer, UK SME bank

“[The third line] have built AI agents internally to analyse that data... to look for anomalies and outliers. They’re picking up on things much more rapidly than a human review would.”

— Chief Compliance Officer, UK SME bank

Continuous monitoring and shared live controls

As AI enables each line to operate with greater coverage and speed, a different model for assurance is beginning to emerge - one in which a single operational control serves all three lines simultaneously: the first operates it in real time; the second monitors its performance continuously; the third independently verifies its design and effectiveness.

"The control should be one control - available to everyone, so each line of defence can perform its duties on the same thing."

— Group CCO/CRO, major payments firm

System integration, data quality, and regulatory acceptance of automated assurance remain significant barriers. But this is the most coherent answer yet to how the Three Lines model adapts to AI at scale - by rebuilding existing structures around shared live data.

Overseeing the overseer agents

The deployment of AI agents within oversight functions introduces a new governance challenge: when agents are used to perform monitoring and assurance, they must themselves be subject to monitoring and assurance.

"We have to oversee the oversight agents and make sure that they're doing it correctly. So in a way it's quite a lot of layers of risk."

— Senior compliance leader, major payments firm

There is a risk that monitoring agents may produce outputs that appear authoritative but are not. The greater the reliance placed on AI-assisted review, the harder failures within the review mechanism become to detect.

Whether the potential of these tools is realised depends on whether the people within each line develop the capability to use them effectively - and to exercise meaningful judgement when they fail.

EXPERT COMMENTARY

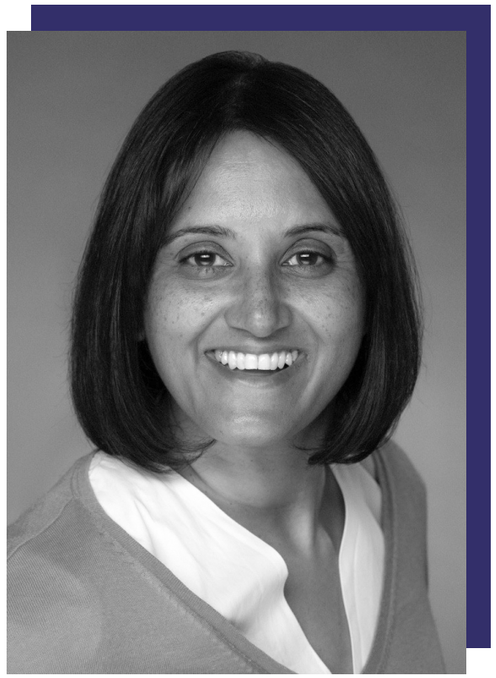

Ratul Ahmed
Global Head of Model Risk Management & Validation

COMMERZBANK

“The frameworks this report describes are largely fit for purpose, the real question is whether the people who need them can reach them.

AI governance has historically been concentrated in specialist second-line functions. Built by the technically fluent, for the technically fluent whilst the teams making consequential deployment decisions, daily, are left to navigate by instinct. As an example, the emergence of LLM-as-judge capability does not simply redistribute a validation task. It exposes something more structural: that governance frameworks predicated on periodic, sampling-based oversight were designed for a world of bounded, deterministic models. Probabilistic, context-dependent systems operating at scale require something different; not just better tools, but a fundamentally different distribution of capability.

The AI adoption gap and the AI governance gap are the same problem, manifesting in the same place, at the same time. Institutions anchoring second-line oversight in direct model validation will find that position increasingly difficult to defend as first-line teams deploy AI-driven evaluation at a depth and cadence the second line cannot replicate.

Those that reposition governance as a democratised capability, very much embedded in model risk standards, automated validation frameworks, and the technical literacy of every line, will no doubt find both gaps close together.

That is the real evolution the Three Lines model is being asked to undergo, not a restructure but a redistribution of capability.”

KEY INSIGHT

AI is changing how the Three Lines of Defence model is operationalised in practice

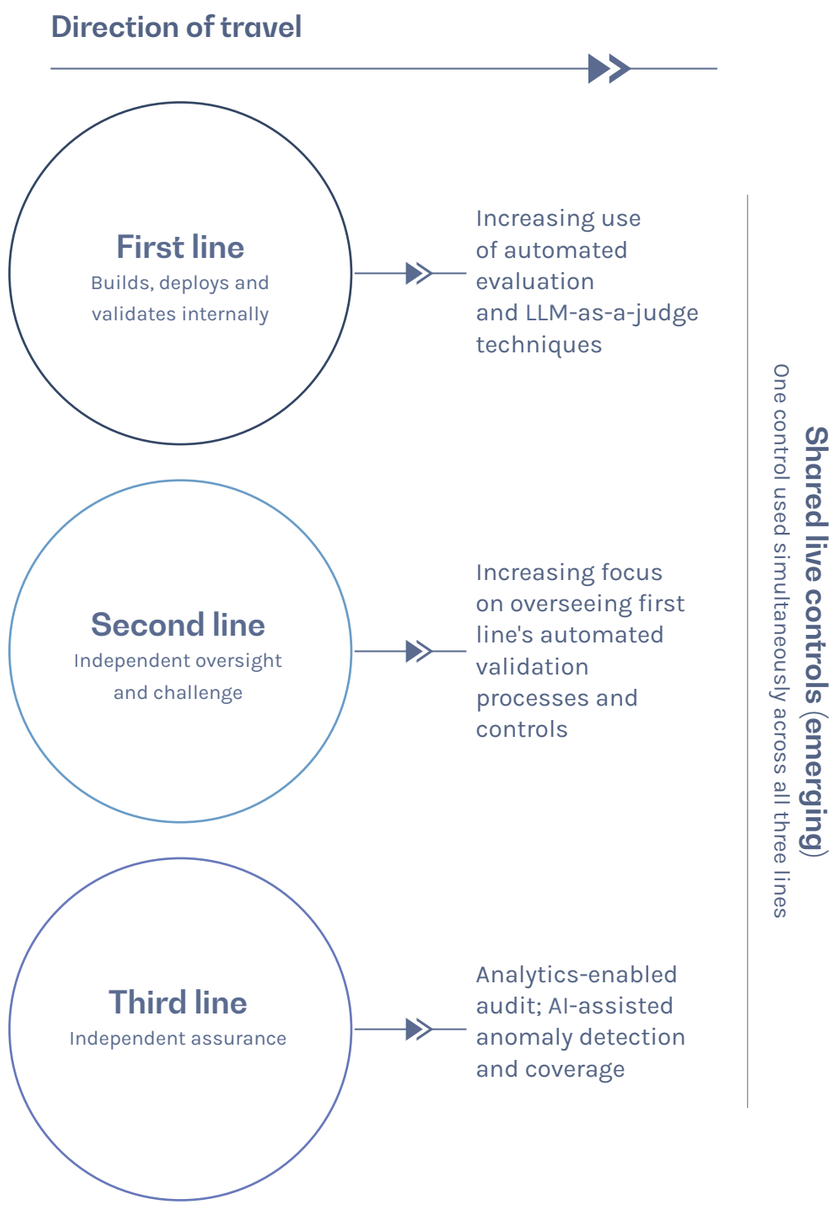

10
The skills gap as a governance problem

Why governance capability is struggling to keep pace

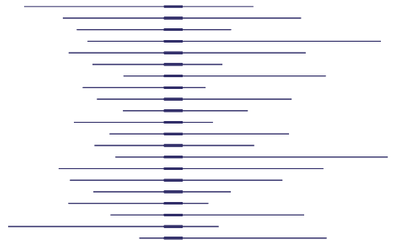

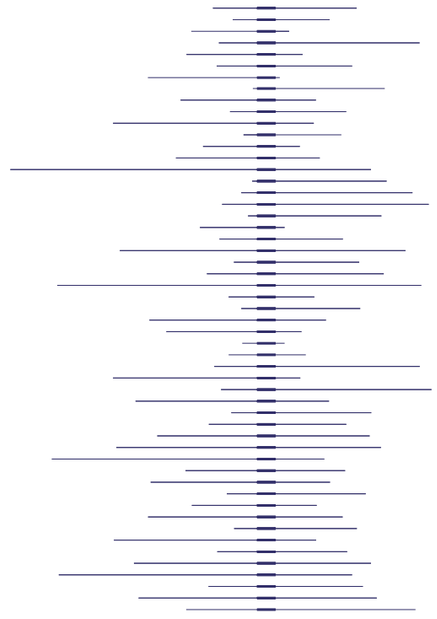

The previous section identified a significant opportunity: AI has the potential to materially strengthen oversight. But that opportunity is only realisable if the people responsible for oversight can identify issues, and exercise sound judgment when they go wrong.

Recent research from the UK government shows that the key AI skills gaps in financial services are in “governance, ethics, and interpreting AI outputs, especially in compliance and legal teams”.[6] The evidence from this research corroborates this finding - for many institutions, that capability is not yet there. Furthermore, institutions are attempting to close a gap they cannot clearly define, as there is no shared standard for what AI governance capability in oversight functions actually requires.

Control functions are not keeping pace

The gap is most visible at the interface between first-line deployment and second-line oversight. Engineers are building increasingly complex systems, often leveraging external models and evolving workflows. Compliance and risk teams, tasked with challenging those systems, are often operating without the technical literacy required to do so effectively.

"The coders and the engineers, they're fantastic. But I don't know that the control functions are there yet in terms of how do you govern all of this."

- Chief Compliance Officer, major payments firm

The skills required are distinct from traditional compliance expertise: understanding how models behave under different conditions, what they exclude as well as include, how thresholds are calibrated, and how outputs should be interpreted in context.

"I need people who understand compliance at a substantive level, but I also need some of them to be tech-able - to be able to challenge the business - because otherwise we're not going to be able to provide sound advice."

- Chief Compliance Officer, major payments firm

Where these capabilities are absent, independent challenge becomes difficult to exercise in a meaningful way.

Many are resistant to engaging with AI

The challenge is not only one of access to training. Several practitioners described significant resistance within control functions - reflecting anxiety about what deeper engagement with AI might mean for existing roles.

“People are still quite fearful of the topic, and therefore fearful to engage."

- Chief Risk Officer, major UK wealth manager

At one financial institution, this translated into a clear internal divide - with some teams actively experimenting while others remained largely disengaged, concerned that deeper engagement with AI would ultimately threaten their roles.

Unless addressed, this cultural resistance will outlast investment in tools or training. Effective AI governance requires control functions to engage with the technology. Where hesitation persists, culture itself may become a governance risk.

[6] Skills England (2025), AI skills for the UK workforce

EXPERT COMMENTARY

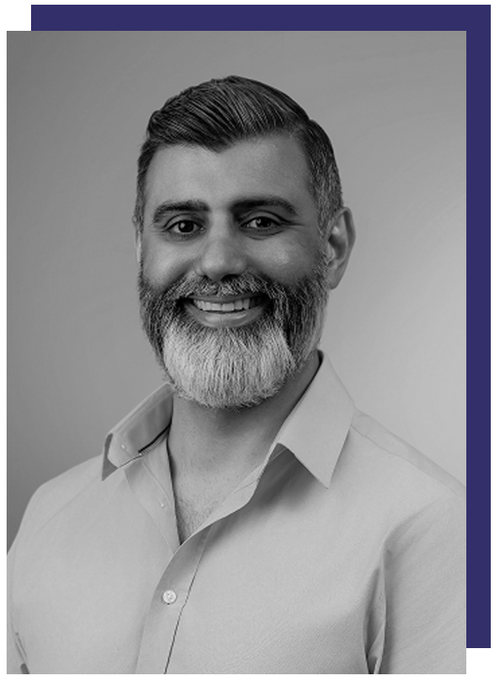

Mitch Trehan
Chief Compliance Officer

Allica Bank

“AI is here to stay. The real question is how to embed it into an organisation - and bring people with you.

Everyone is talking about AI. Far fewer are turning that talk into something that meaningfully changes how a company operates. That shift requires more than tools; it demands cultural adoption at every level.

At Allica, AI is now embedded in our Company Scorecard. Each Executive Committee member is measured on how effectively AI is adopted within their function, making accountability explicit.

We’ve also established an AI Risk Working Group to oversee the tools we deploy, ensuring innovation is balanced with control. And every week, we run a company-wide AI “show and tell,” where colleagues share how they’re using AI in practice - spreading ideas, building confidence, and accelerating adoption across the business.”

Voluntary training fails to move the dial

Firms engaging seriously with this challenge are finding that access to training alone is not sufficient. In several institutions, learning resources are already available, but uptake is limited. Without clear expectations or incentives, capability development remains uneven and often deprioritised over immediate operational demands.

"Sometimes you have to enforce it. If you don't have someone pushing you, you're going to opt for the easy solution, which is carry on as you are."

- Head of Compliance, major international bank

More structured approaches are beginning to emerge, including mandatory training pathways, role-specific capability frameworks, and internal AI champion networks to support adoption within control functions.

No shared definition of what good looks like

The potential for AI to strengthen oversight and free compliance professionals to focus on judgement rather than process is significant. But that potential will not be realised in organisations where the workforce remains undertrained, or where the benefits of the technology are poorly understood. Alongside technical training, institutions will need to articulate clearly how AI can improve the work of risk and compliance professionals rather than replace it.

Addressing this is partly an institutional challenge, but it is also an industry one. There is currently no shared framework defining what AI governance capability in oversight functions requires - what skills are needed, at what level, and in which roles. Without that, firms are left to develop their own definitions in isolation. The skills gap will persist not only because training is insufficient, but because there is no common standard for what closing it requires.

EXPERT COMMENTARY

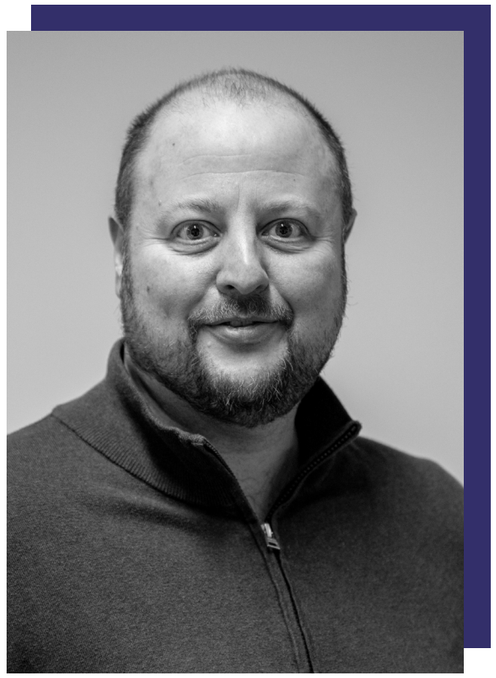

Rob Phillipson
Managing Director

ICA INTERNATIONAL COMPLIANCE ASSOCIATION

"AI literacy within compliance functions is foundational to the safe and effective adoption of the technology across financial services.

Compliance professionals need a functional understanding of how these systems are built, where they may fail and how their risks map onto existing regulatory frameworks covering but not limited to conduct, data protection, financial crime prevention and operational resilience. Without that grounding, compliance is reduced to a reactive gatekeeper, slowing adoption without meaningfully reducing risk. With it, compliance becomes a genuine enabler of innovation.

The best-in-class modern compliance functions are embedded within the business rather than sitting apart from it. They partner with product, engineering and commercial teams from the outset, shaping how AI is designed, tested and deployed so that controls are built in rather than bolted on. Their instinctive response to a complex regulatory scenario is not to block, rather it is to say "here's how". That posture - technically literate, commercially engaged and solution-oriented - is what allows firms to harness AI at pace while supporting sustainable, well-governed growth."

11
Towards shared AI governance implementation guidance

The missing implementation layer

Across the institutions interviewed for this research, firms are solving the same governance problems independently, without shared standards to work from. They are interpreting the same regulatory principles in isolation, building capability frameworks from scratch, and navigating the same implementation questions without reference to how peers are approaching them. The duplication of effort is significant and unnecessary - there is a clear case for collective action.

The implementation gap

UK financial services regulation is principles-based, and practitioners broadly support that approach. But principles require interpretation - and without a shared operational layer translating regulatory expectations into practice, each firm is left to develop that interpretation alone.

The EU AI Act provides codes of practice for general-purpose AI model providers, covering transparency and systemic risk, but not how financial institutions should govern AI deployments in practice.

In both the UK and EU, there is currently no authoritative, sector-specific guidance translating regulatory expectations into operational AI governance practice. While each firm is different and proportionate governance will differ, the absence of shared implementation guidance is creating real uncertainty. Firms interpret existing rules conservatively, avoid use cases they cannot clearly justify, and delay deployment where expectations are unclear.

[7] Cyber Risk Institute (2026). Financial Services AI Risk Management Framework. https://cyberriskinstitute.org/artificial-intelligence-risk-management

Addressing the AI governance implementation gap

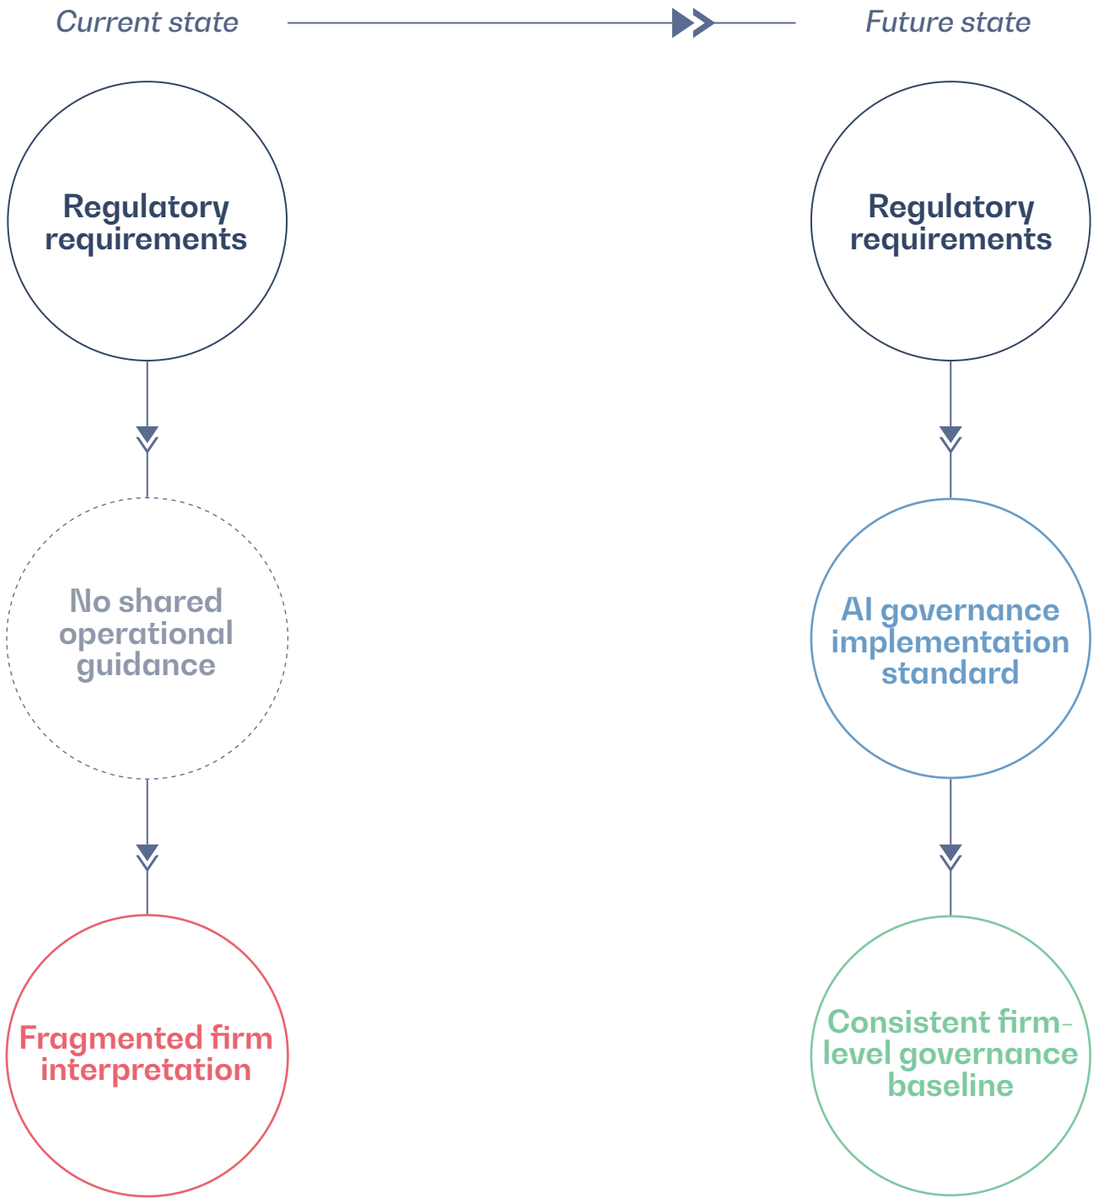

Industry should build the standard - not wait for one

The solution does not necessarily need to come from regulators. In the UK, a useful precedent already exists in financial crime. The Joint Money Laundering Steering Group (JMLSG) provides detailed operational guidance developed by industry and approved by government. It carries authority because it reflects how the sector actually operates.

"You could have that for AI - not prescriptive, but endorsed. Industry-based guidance contributed to by banks and other financial institutions and core players, endorsed by government."

- Chief Compliance Officer, payments firm

Other jurisdictions have already acted

The US has already recognised this gap and moved to fill it. NIST's AI Risk Management Framework (2023) established a voluntary standard for managing AI risk, and in February 2026 a public-private collaboration - including the Cyber Risk Institute (CRI), 108 financial institutions, and the US Treasury - adapted it into a practical, Financial Services AI Risk Management Framework.[7]

Similarly, the Monetary Authority of Singapore (MAS) has led equivalent work through Project MindForge, a public-private initiative. An AI Risk Management Toolkit was published in March 2026.[8] At the technical level, GovTech Singapore has developed the Agentic Risk and Capability (ARC) Framework - a structured approach to risk classification and control mapping specifically designed for agentic AI systems, and among the most technically rigorous work in this space globally.[9]

No comparable public-private initiative or sector-specific guidance for financial services currently exists in the UK or EU.

[8] Monetary Authority of Singapore (2026). Project MindForge: AI Risk Management Toolkit. https://www.mas.gov.sg/news/media-releases/2026/mas-partners-industry-to-develop-ai-risk-management-toolkit-for-the-financial-sector

[9] GovTech Singapore (2025). Agentic Risk & Capability Framework. https://govtech-responsibleai.github.io/agentic-risk-capability-framework/

Willem Wellinghoff
UK Chair and Chief Compliance Officer

ecommpay

“Artificial intelligence is no longer just a technological discussion topic; it is becoming a fundamental part of the engine driving the future financial ecosystem. But to truly harness its potential, we must take responsibility for its governance today.

It is critical that we, as an industry, take the lead in defining AI standards. The technology is evolving at significant pace. If we wait for government and regulators to mandate legislation, we risk being constrained by static frameworks that fundamentally cannot keep pace with innovation.

Instead, we must champion an industry-led approach that is dynamic, nuanced, and continuously adaptable. By establishing robust guidelines that earn the trust and endorsement of governments and regulators, we build a collaborative model that helps us embed ethical, transparent practices across our organisations, protecting consumers while keeping us agile to innovate. Ultimately, we must be the architects of our own compliance.”

The case for shared implementation guidance

The governance challenges described throughout this research point to the need for an implementation layer translating regulatory requirements into operational practice. Rather than prescribing how firms deploy AI, practitioner-developed guidance could help institutions interpret existing regulatory expectations in a consistent way, while preserving flexibility for different organisational structures and risk appetites.

For agentic AI in particular, the need is urgent. As financial institutions increasingly deploy autonomous agents across their operations, the absence of sector-specific shared standards - covering risk classification, control mapping, documentation, monitoring, and oversight - risks real harm to consumers and to financial stability.

KEY INSIGHT

What industry-developed AI governance guidance might address

Institutions differ in size, structure, and AI maturity, and proportionate governance will vary accordingly. But shared operational guidance tailored to the UK and EU financial services context would provide a common reference point. Based on the challenges identified, it might include:

01
AI risk and control mapping
A shared standard for assessing system capabilities, classifying deployment risk, and mapping controls to regulatory requirements.

02
Validation and testing
Common expectations for pre-deployment and ongoing testing — including for probabilistic outputs — so second-line functions can assess whether first-line validation is sufficient.

03
Documentation
Minimum requirements covering model identity, data provenance, testing conducted, and deployment rationale.

04
Monitoring and metrics
Shared expectations for what failure trigger escalation and how customer-facing AI outputs should be reviewed.

05
Third-party AI
Expectations for due diligence, contractual controls, and monitoring where AI is provided by external vendors or embedded within third-party platforms.

06
Lifecycle and change governance
Standards for how to update models, prompts, or training data should be approved, tested, and documented — including when changes require re-validation.

12
The risks of inaction

How insufficient AI governance creates risks to consumers and financial stability

The governance gap identified throughout this research has consequences beyond the firms that experience it. In financial services, failures in AI oversight can propagate quickly across the system. This creates systemic risk: an AI system with ineffective oversight can accumulate harm at a pace and scale that manual processes could never achieve, before anything visibly goes wrong.

This section does not attempt an exhaustive taxonomy of those risks. It focuses on the themes raised consistently by practitioners across the interviews, and on what they collectively suggest about the nature and scale of the problem.

Harm at speed and scale

AI systems optimised around measurable metrics can produce outcomes their operators did not intend - and traditional compliance monitoring, built around periodic sampling and retrospective review, was not designed to catch problems accumulating at the speed that AI facilitates. The consequences of that mismatch can compound significantly before anything visibly goes wrong.

The reference point raised by more than one interviewee was Payment Protection Insurance (PPI) - a mis-selling problem that built up over years and resulted in more than £38 billion in redress paid to consumers.[10] AI significantly compresses that timeline.

"Think how long it took to build up PPI. That could happen in two weeks with AI — billions of pounds of liability quickly because something wasn't right."

— Director of Compliance, major UK wealth manager

An AI system generating, say, personalised customer communications at scale - flagging account features, prompting product upgrades, or summarising terms - could systematically mislead customers. If the AI makes the same error consistently, at volume, then by the point something visibly goes wrong, the damage may already be substantial.

Markets, infrastructure, and the concentration problem

Beyond conduct risk, practitioners raised concerns about AI's potential to amplify market instability. The 2010 "Flash Crash" - in which US equity markets briefly lost around $1 trillion in market value within minutes - demonstrated how algorithmic systems can produce herd-like behaviour that rapidly magnifies market moves.[11] As agentic AI becomes more prevalent in trading and real-time decision-making, the potential for correlated failures increases as similar models respond similarly to the same conditions.

"Undetected weaknesses could cause market shocks — we've seen it before with algo trading models causing mini crashes."

— Chief Compliance Officer UK, major international fintech

The concentration of AI infrastructure compounds this. If the sector becomes reliant on a small number of foundation model providers, a flaw or vulnerability in a shared model becomes a systemic exposure - not a risk local to any individual firm.

"AI could almost be viewed as critical national infrastructure. In the future there'll probably only be a couple of core AIs that everyone's leveraging — it's important to make sure the AI itself doesn't have systemic flaws or weaknesses that nobody picks up on."

— Chief Compliance Officer, major international fintech

[10] Financial Conduct Authority (2020). Monthly PPI refunds and compensation. https://www.fca.org.uk/data/monthly-ppi-refunds-and-compensation

[11] U.S. Securities and Exchange Commission and U.S. Commodity Futures Trading Commission (2010), Findings Regarding the Market Events of May 6, 2010, Report of the Staffs of the CFTC and SEC.

EXPERT COMMENTARY

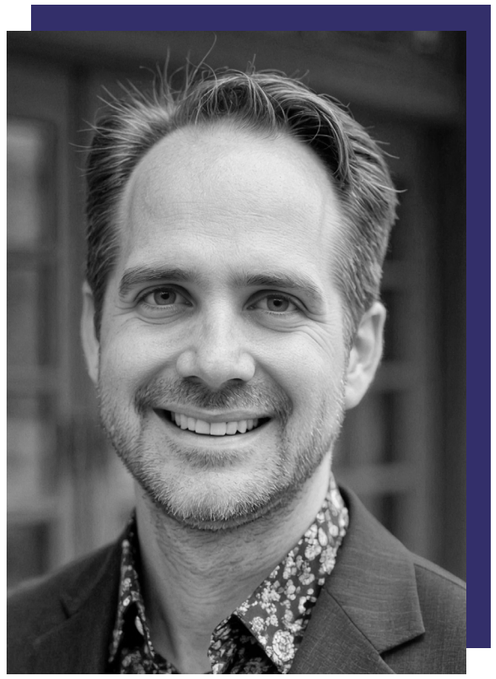

Andrew Sutton
Visiting Fellow

OXFORD MARTIN SCHOOL
UNIVERSITY OF OXFORD

"Financial Institutions have been rightly concerned with managing risks of their own AI deployments. But many risks arise from AI's use by customers, competitors, adversaries, and in wider society. These may emerge where AI makes existing risk management assumptions unreliable or controls less effective. Often such assumptions relate to latent, sometimes unnoticed, protective features within the financial system, such as human inertia and diversity of action.

This may play out in many risk-critical areas, such as liquidity risk management. For example, widespread adoption of AI agents by individuals and companies to manage their finances more efficiently could reduce deposit stickiness in ways that increase the threat of bank runs. An incident might be triggered by something as simple as a set of similar agents moving funds in response to the same change in deposit rates, or the same rumour, or quickly following actions they see in the market. The agents may be acting on publicly-visible information, or on signals that are not easily detectable or interpretable to human overseers.

Such challenges cannot be addressed by financial institutions alone, especially as there is often a public policy trade-off: features that make advanced AI powerful and convenient are often the same ones that create risk. Detecting and predicting where such challenges are likely to arise, adapting controls as needed, and adjudicating new trade-offs is an urgent, exciting, and important priority for the finance and AI sectors, researchers, and government."

AI in the hands of adversaries

As well as internal risks, the governance gap within financial institutions creates an exploitable surface for malicious actors, which is already being tested.

AI introduces attack vectors that existing security frameworks are only beginning to address. Prompt injection, jailbreaking, and adversarial inputs designed to manipulate model outputs require new controls that sit alongside, rather than within, conventional cybersecurity governance.

"AI security is suddenly something different. New threats need to be managed... like prompt injection, guardrails, jailbreaking."

- Group Head of Responsible AI, major UK bank

The vendor ecosystem compounds this, since institutions are dependent not only on their own AI deployments but on the AI embedded within their suppliers' systems - and their suppliers' suppliers. Oversight of this is patchy at best.

In financial crime and fraud detection, institutions and criminals are running competing models. Weak internal governance means offensive actors can outpace the institutions attempting to defend against them. Global fraud losses are estimated at $579.4 billion in 2025, while 90% of financial professionals report an increase in AI-enabled attacks.[10] Interviewees warned that the scale of criminal revenues allows organised fraud networks to rapidly adopt and deploy new AI capabilities.

"The thought of the scams industry - which has already figured out how to convince a well-educated, ostensibly smart British citizen to send life-changing amounts of money to somewhere they can't recover it - operating at industrial scale using AI tools is really quite scary."

- Chief Risk Officer, major UK digital bank

Individual institutions doing creditable work on controls and detection are struggling to keep up with an adversary that is growing exponentially. If institutions fail to govern their own AI - and to use AI effectively to upgrade their defences - the consequences extend beyond the firm to the wider financial system.

A regulatory dynamic shaped by failure

Across the institutions interviewed, there was a consistent expectation that rather than the regulatory framework for AI in financial services being shaped by proactive standard-setting, it will be shaped by something going wrong.

"When is the first time someone somewhere in the industry gets this wrong - and what [is the regulator] going to do about that?"

- Director of Risk and Compliance, UK fintech

"It's because there's not been a massive AI-fuelled incident yet. Once regulators say 'this is what we expect to see, and there will be sanctions if you don't' - I bet that would change the way risk teams engage."

- Senior executive, major UK banking group

This is a familiar pattern. PPI mis-selling and the 2008 financial crisis both illustrate how risks in financial services are often recognised in principle long before they are addressed decisively in practice. In each case, the eventual cost of intervention far exceeded the cost of acting earlier.

AI changes the speed at which that pattern can unfold. Failures that once accumulated gradually can now emerge far more quickly when decisions are automated and systems operate at scale. In that environment, the consequences of weak governance may become visible only once significant harm has already occurred.

KEY INSIGHT

The risks of ungoverned AI extend beyond the institutions that deploy it

The following reflects themes raised consistently by practitioners across the interviews, rather than an exhaustive account of the risks posed by insufficient AI governance in financial services.

- **Harm at speed and scale.** Conduct failures that once accumulated gradually could now emerge far more quickly when automated systems operate at scale. Oversight frameworks were designed for a world where harm accumulates slowly enough to be caught, but AI changes that assumption.
- **Market and infrastructure fragility.** Common model architectures and training data can produce correlated behaviour across institutions. Concentration in a small number of foundation model providers means that a shared vulnerability is a systemic one.
- **AI in the hands of malicious actors.** Criminal organisations are deploying the same AI capabilities as financial institutions - at scale and with growing sophistication. Weak internal governance undermines institutions' ability to defend against these threats.

These risks interact. An external threat exploiting weak internal governance - in a market where AI infrastructure is concentrated and regulatory response is reactive - is a scenario in which failures at individual firms can quickly become systemic.

13
Conclusion

The financial services industry is deploying AI faster than it can govern it. This research identifies two gaps the sector must close - and a future worth building towards.

There is a capability gap: oversight functions lack the skills, tools and infrastructure to govern AI at the pace and scale of business deployment.

They are being asked to interrogate systems they often cannot see, with capabilities that have not kept pace - and risk falling further behind as institutions move toward agentic AI.

There is also an implementation gap: firms are solving the same governance problems independently, without shared standards.

Sector-specific, practitioner-built AI implementation guidance - developed with regulator engagement and modelled on the JMLSG precedent - is both achievable and necessary. The US has moved to fill this gap. The UK and EU have not.

Insufficient AI governance puts consumer protection and broader financial stability at risk.

Institutions are deploying AI faster than their controls can follow - and criminal actors are already exploiting those gaps at scale. Conduct failures that once took years to accumulate could now occur in weeks. Practitioners interviewed for this research expect regulation to be shaped by something going wrong - a pattern the industry has lived through before.

Yet there is also opportunity: the same AI that creates these risks can also strengthen oversight if governed effectively.

Continuous monitoring, automated validation, and compliance functions freed to focus on judgement rather than process are already emerging at the leading edge of institutional practice. Whether the industry builds toward that future deliberately, or arrives there via a crisis it could have avoided, is a choice the industry is making now.

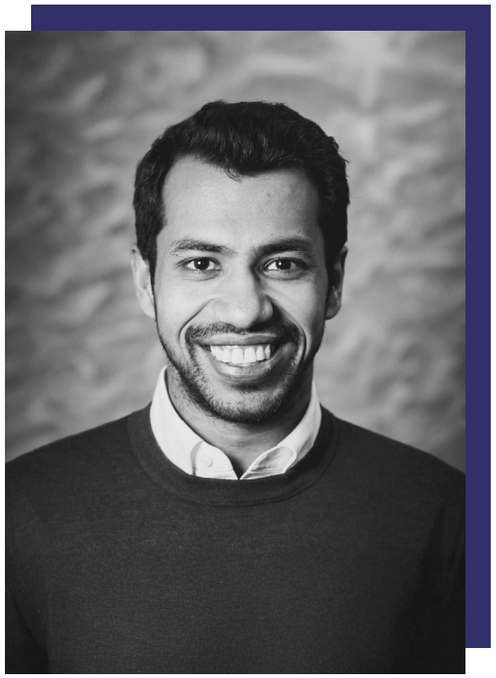

Ritesh Singhania
CEO and Co-Founder
zango

“When we speak to risk and compliance leaders across financial services, it is clear that organisations want to deploy AI more widely, but are still working out how to do so with the right guardrails and oversight in place.

Through our work building and deploying AI agents in risk and compliance teams in financial institutions, we are working hand-in-hand with firms as they reshape regulatory compliance for the AI age.

We coordinated this research to better understand how financial institutions are navigating that transition - what is working, where governance frameworks are under strain, and how oversight models are adapting as AI becomes more deeply embedded across financial services.

This research points to a clear need for practical, industry-developed guidance on what good AI governance looks like in practice. We hope it contributes to that effort.”

About Zango AI

Zango is the AI compliance layer for financial services. The platform uses AI agents trained on financial regulation to help financial institutions monitor regulatory change, assess compliance gaps, and review marketing and product launches against regulatory rules. Founded in 2024 by Ritesh Singhania and Shashank Agarwal, Zango AI serves clients across banking, insurance, payments, and digital assets, with offices in London, Lisbon, and Bengaluru.

Learn more at www.zango.ai

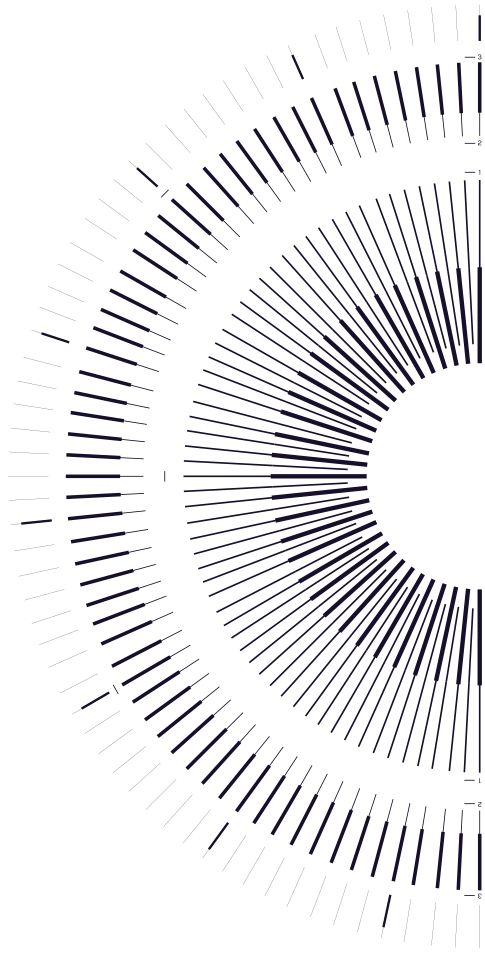

2026

THE FUTURE OF AI GOVERNANCE & COMPLIANCE IN FINANCIAL SERVICES

RESEARCH REPORT

zango
RESEARCH INITIATIVE

zango

www.zango.ai

CONTRIBUTORS

HOUSE OF LORDS	LORD TIM CLEMENT-JONES
HOUSE OF COMMONS	RT HON JOHN GLEN MP
SANTANDER	DEAN NASH
ST. JAMES'S PLACE	PAUL LOFTUS
STRIPE	ARMAN FALLAH
REVOLUT	BEN ELLIS
MONZO	IAIN LAING
STANDARD CHARTERED	COSETTE RECZEK
LLOYDS BANKING GROUP	SUZANNE BRINK
COMMERZBANK	RATUL AHMED
ALLICA BANK	MITCH TREHAN
ECOMMPAY	WILLEM WELLINGHOFF
NOVOBANCO	ARCHIT CHAMARIA
INTERNATIONAL COMPLIANCE ASSOCIATION	ROB PHILIPSON
OXFORD MARTIN SCHOOL	ANDREW SUTTON
UNIVERSITY OF GLASGOW	DR ALESSIO AZZUTTI

DISCOVER WHAT'S NEXT
Closing the implementation gap