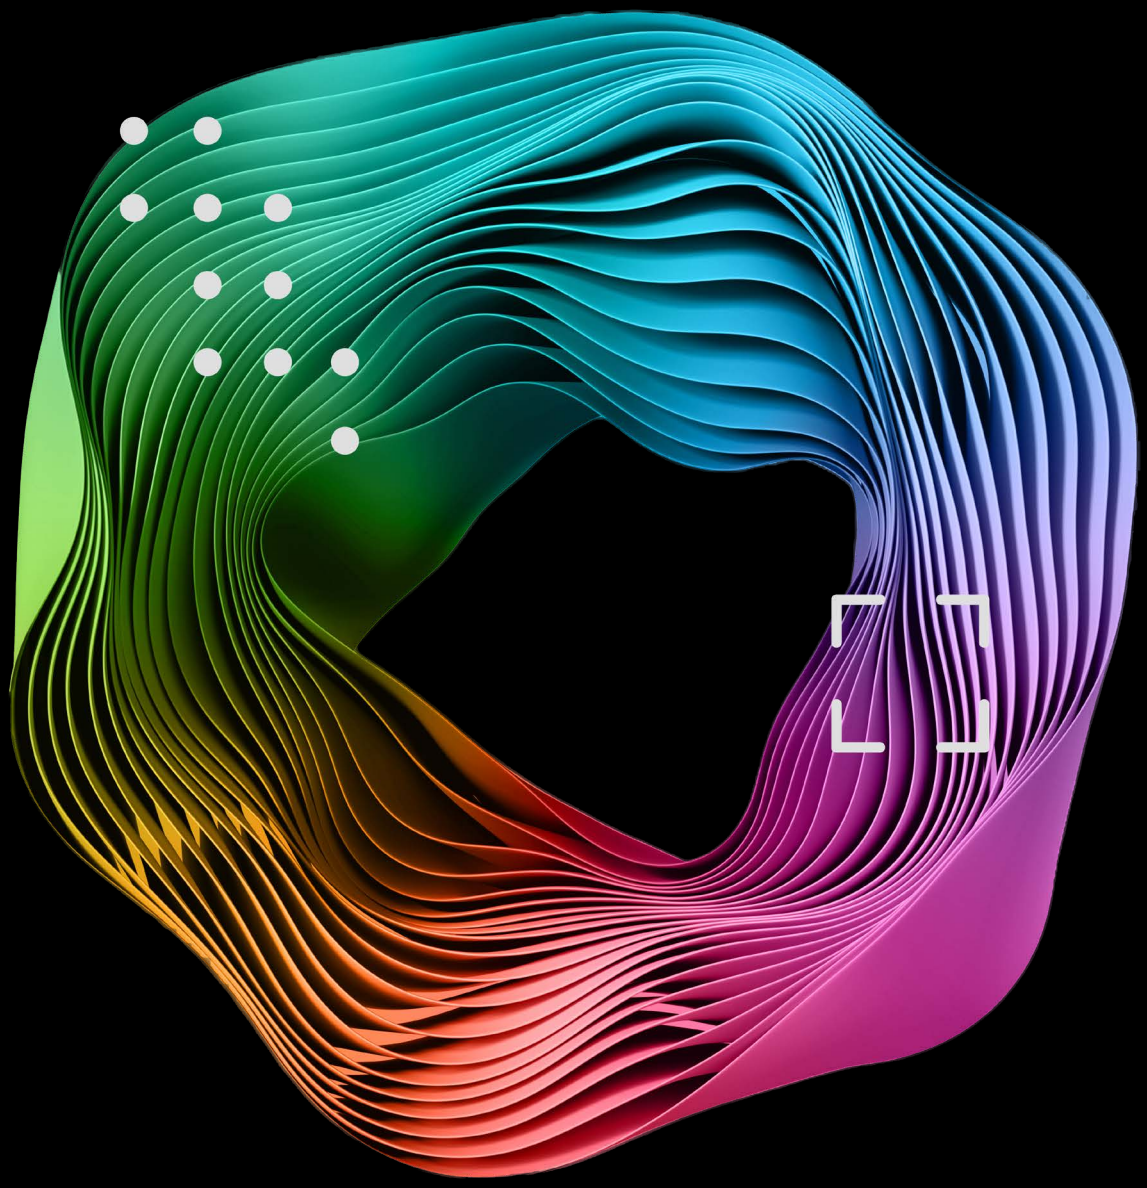
Deloitte.

Together makes progress



Deloitte Access Economics

Banking on trust

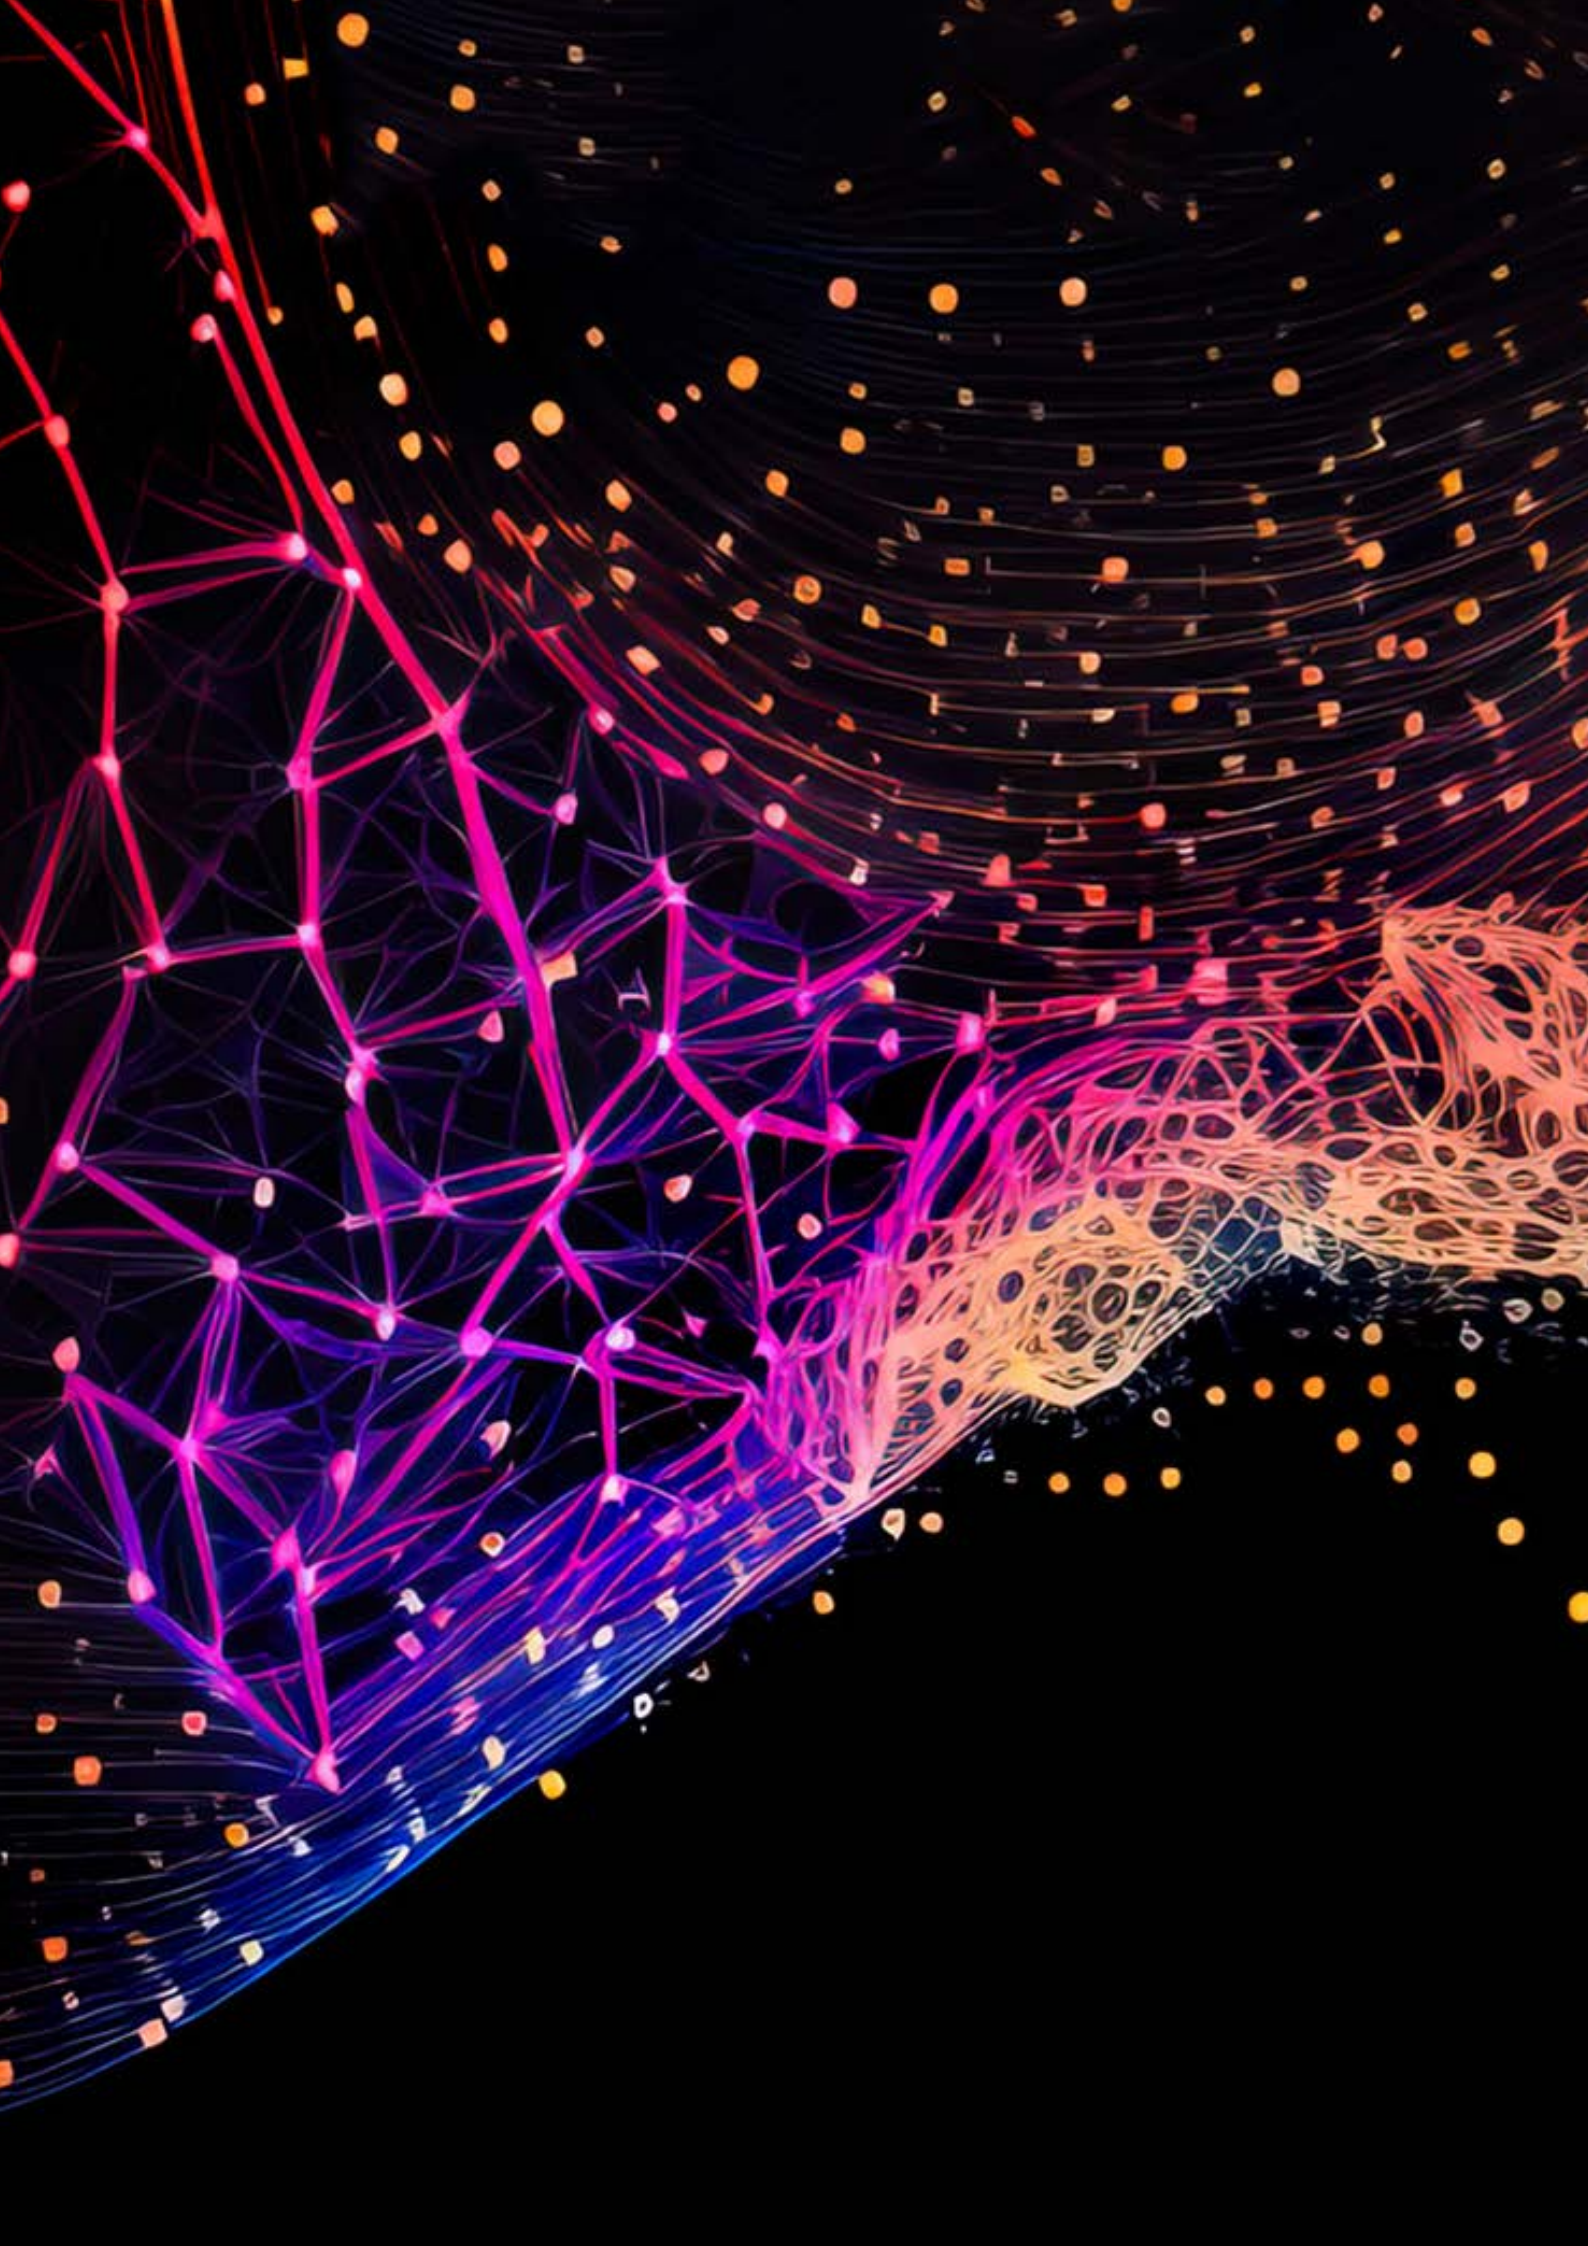
AI governance for growth, resilience and scale

July 2026



Contents

Executive summary	6
1 | Why AI matters for global banks now	8
2 | From experiments to scale: AI in the banking industry	10
3 | What is the state of AI governance in banking?	15
4 | The AI governance dividend	28
5 | How banks can close the AI governance gap	31
Appendix A | Global bank survey	35
Appendix B | Index development and econometric modelling	39
Appendix C | Regulatory developments in select markets	44
Endnotes	46
Contact us	49



Banking on trust: **AI governance for growth, resilience and scale**

AI has become a core strategic capability for global banks, helping drive efficiency, resilience, customer protection and new growth. As adoption accelerates from experimentation to enterprise-wide deployment, banks are also facing sharper data, regulatory, cyber and reputational risks, making effective AI governance essential to scaling safely. This report finds that stronger AI governance is linked to broader AI deployment and stronger revenue growth from better use of the technology, positioning governance as a growth enabler rather than a brake on innovation.

Now is the time: AI adoption is outpacing its guardrails

Growth in Adoption

The share of employees using AI on a weekly basis doubled in the year to 2026 to reach 63% of all employees.

Hazards & Risks

8x

The annual number of AI incidents for financial services increased by nearly 8 times between 2022 and 2025.

Consumer Trust

84%

84% of consumers would switch providers if their data was mishandled by their financial institution.

Nearly all banks could improve their AI governance

Distribution of Trustworthy AI Governance Index



87%

of global banks could improve their AI governance, according to our Trustworthy AI Governance Index.

The prize for improvement: **AI adoption at scale and revenue growth**

31% More weekly users of AI systems in banks with optimised governance level compared to rudimentary levels.

5x More fully implemented AI solutions for banks with optimised governance compared to those with rudimentary levels.

10% The average revenue growth uplift from a 10 point increase in Trustworthy AI Governance Index score.

Still a way to go

72% of banks have less than half of AI use cases in their central register.

44% of banks have risk monitoring across the implementation lifecycle for Agentic AI.

15% of banks provide regularly refreshed AI governance training to all staff.

Executive summary

AI now matters to global banks not merely as a productivity tool, but as a strategic capability shaping competitiveness, resilience and growth. This matters more than ever as banks navigate a more difficult operating environment marked by global energy shocks, weaker growth expectations and tighter margins. In this context, AI offers a powerful means to improve efficiency, strengthen risk modelling, protect customers from fraud and cyber threats, and unlock new sources of revenue.

Global banks are moving quickly from AI experimentation towards full scale enterprise deployment. The share of bank employee use of AI at least weekly has doubled to reach 63% in 2026, and most global banks have fully implemented AI solutions in business areas such as customer service, IT, marketing and sales, operations and finance. Many are also piloting AI solutions in HR, legal and compliance and R&D. The next wave is likely to be even more significant. Agentic AI is creating an automated workforce that is capable of making decisions and executing complex processes faster and more accurately than humans ever could.

As AI becomes more deeply embedded across banking operations, risk exposure is rising just as quickly. Banks face growing data, legal, model, cyber, conduct and reputational risks. Financial Services now records more publicly reported AI-related incidents and hazards than any other industry.* In the first half of 2026 alone, the financial services industry recorded more incidents than in the whole of 2025. At the same time, regulatory expectations are increasing in several jurisdictions, particularly around transparency, accountability and misuse of AI. Trust is also easily lost, with 84% of consumers saying they would switch providers if their financial institution mishandled their data.

The implication is clear: the ability to scale AI safely depends on effective governance.

To assess how well banks are positioned, this report introduces a Trustworthy AI Governance Index tailored to the banking sector. The Index assesses governance maturity across five pillars – principles and policies, organisational structure, procedures and controls, people and skills, and monitoring, reporting and evaluation – using 41 indicators. It was applied to survey responses from 135 senior leaders from Global Systemically Important Banks (G-SIBs), Domestic Systemically Important Banks (D-SIBs) and other large banks. The Index provides management teams with a practical way to assess current governance arrangements, identify gaps and prioritise investment and effort.

The findings show that most banks are not yet where they need to be. Some 87% have scope to materially strengthen AI governance, while only 13% currently achieve an optimised level. Concerningly, 10% of global banks still have governance that is effectively ad hoc.

Although many banks show greater maturity in principles, policies and some procedural controls, the largest weaknesses are in the organisational and human pillars. Organisational structure is the least mature overall, with more than a quarter of banks ranked in the lowest category for the index. People and skills also lag, with many banks yet to build the training and capability needed to apply AI governance consistently in practice.

Banks with stronger governance share several characteristics. They are more likely to treat AI as a transverse risk embedded across existing risk domains rather than as a narrow standalone issue. They are also more likely to have visible executive accountability and leadership from a Chief AI Officer, Head of AI, or senior technology executive. By contrast, banks with less effective governance are more often associated with leadership of AI governance integrated within a broader role in the

* Where possible, this report presents data relevant specifically to global banks or the banking industry. This data has been supplemented by data for the financial service industry where global banking data is not available. To ensure transparency, we used the nomenclature used in the original source.

> The implication is clear: the ability to scale AI safely depends on effective governance.

First Line of Defence (e.g. Chief Data/Analytics Officers) or relying on the Second Line of Defence (i.e. Chief Risk Officers) to drive. This leads to lower strategic coordination and capability gaps across the workforce.

As AI systems becomes more autonomous and agentic, the First Line of Defence will increasingly need to shift from primarily human-led oversight to technology-enabled governance, with embedded controls, continuous monitoring and automated decision oversight operating at the point of use. Despite tech-enabled governance playing an increased role, human First Line of Defence roles will also need greater fluency of what responsible AI use means for their role and application. The growth of both AI adoption and technology-enabled governance across will likely make the impact of a dedicated AI leader or executive even more pronounced over time.

As AI capabilities continue to accelerate, governance is not simply a compliance exercise. Effective AI governance gives senior leaders the confidence to deploy AI at scale, integrate it into higher-value activities and respond effectively when issues arise.

Banks with more mature governance settings are better able to translate AI ambition into operational reality. Optimised banks, on average, have fully implemented AI solutions in five out of 10 business areas – more than double than for banks at the rudimentary level and more than five times the number of banks at the ad hoc stage. Effective AI governance does not appear to suppress adoption; rather, it creates the conditions for broader and faster deployment.

While many senior leaders of banks are looking towards AI to improve commercial outcomes, realising the commercial dividend is far from automatic. A separate study found that 95% of enterprise generative AI investments had produced no measurable impact on profit.

This is where governance can make a difference. The financial value of AI depends not just on the technology itself, but on whether banks can deploy it at scale in a way that aligns with the business strategy, embed it into core processes and manage the associated risks which will depend significantly on effective governance. Analysis undertaken for this report points to a positive relationship between AI governance and revenue growth. Econometric modelling indicates that banks with higher Trustworthy AI Governance Index scores tended to record stronger revenue growth from unlocking potential commercial benefits from AI, even after controlling for the extent of AI use, headquarters location, workforce size and years in operation.

On average, a 10-point increase in the Index is associated with a 10 percentage-point increase in revenue growth. The effect is economically material. For the average bank in our survey, revenue grew by 10% between FY2023–24 and FY2024–25. The modelling suggests that, with AI governance 10 points higher, revenue growth could have been closer to 20% over the same period. For a bank with US$500 million in annual revenue, that implies revenue growth of around US$100 million rather than US$50 million. While this uplift is significant, further analysis would be required to assess how long the uplift in revenue would last.

The strategic implication is straightforward. Stronger AI governance is emerging as a decisive differentiator for global banks. It should be viewed as a growth enabler, not a brake on innovation. In an industry built on trust, control and resilience, governance is what allows banks to move faster with confidence. As AI becomes more deeply embedded in core banking activities, the governance gap will increasingly become a performance gap. Banks that close the gap stand to gain not only in risk management and regulatory readiness, but in adoption, execution and revenue growth.



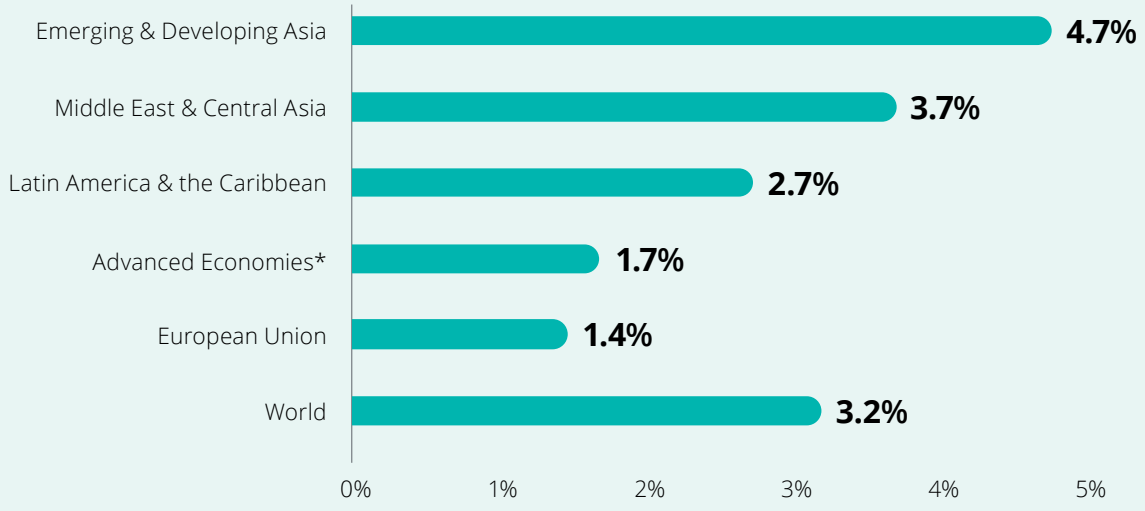
01
Why AI matters for global banks now

Global banks sit at the centre of the international financial system and the broader world economy. Banking is one of the global economy's largest industries by revenue, with total banking assets of US $112 trillion comparable to annual global economic output.[1, 2, 3]

Beyond balance sheet scale, banks, particularly Globally Systematically Important Banks (G-SIBs), Domestically Systematically Important Banks (D-SIBs) and other major banks perform essential economic functions; facilitating payments, intermediating credit, supporting capital markets, financing trade, providing liquidity and enabling risk transfer across borders. Their size, interconnectedness and international reach mean their decisions affect financial markets, as well as business investment, household resilience and wider economic stability.

The operating environment has become more challenging for banks recently, with elevated geopolitical risk and volatility contributing to weaker expected economic growth. The International Monetary Fund (IMF) expects regional growth to remain uneven, with Emerging and Developing Asian economies continuing to drive global growth, while Advanced Economies and the European Union see more subdued expansion.[4]

Chart 1.1 Expected average annual real GDP growth, 2025-2030, by region



Source: International Monetary Fund (2026)

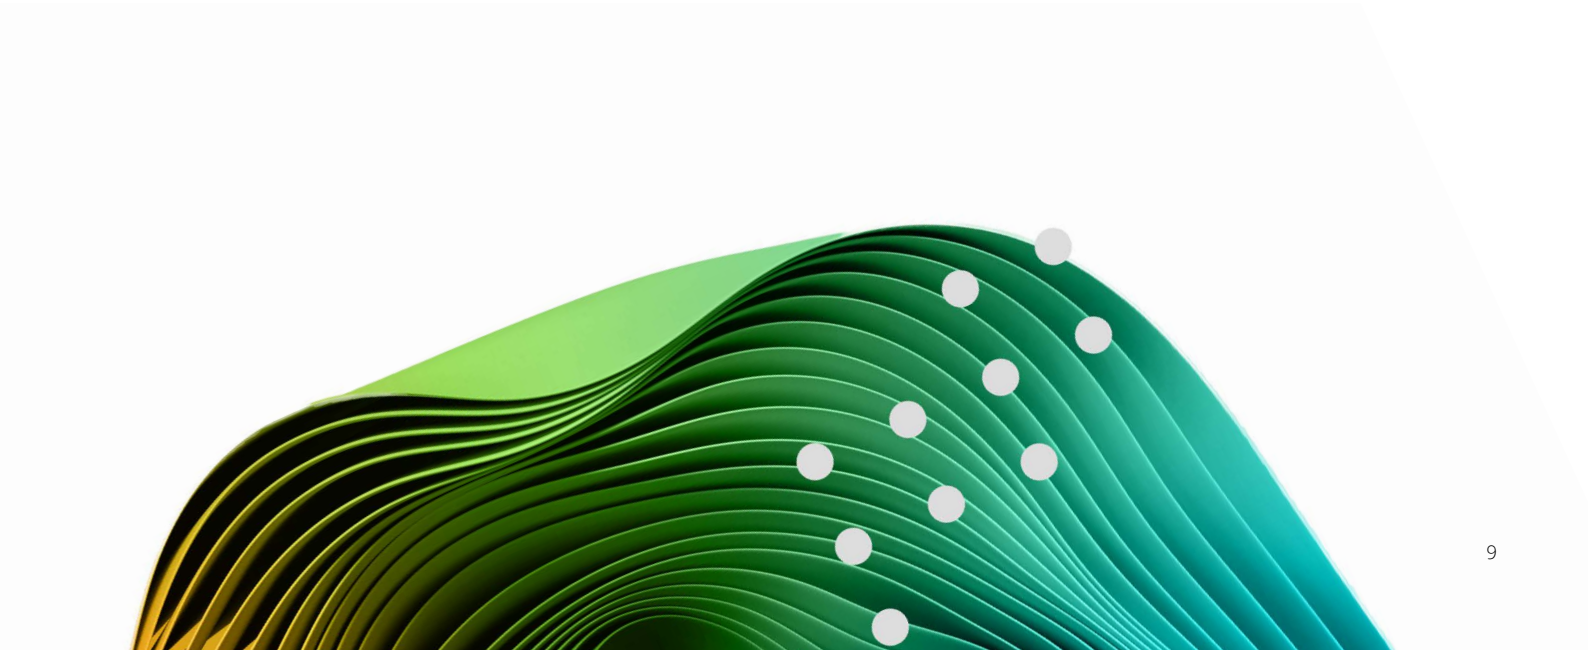
* Advanced economies include the European Union and G7 economies.

These challenges alongside competitive pressure are weighing on bank profitability. Global banks are required to become more efficient, resilient and responsive in an increasingly demanding operating environment. Following a decline in 2025, global net interest margins of banks are expected to remain subdued or compress further in 2026 amid lower interest rates and rising funding costs.[5]

In this context, AI is emerging as a potentially powerful enabler. Across banking, AI offers opportunities to improve productivity, reduce costs and free employees to focus on higher-value tasks. Growing capabilities of AI means that the use cases are expanding rapidly, including real-time fraud detection, risk modelling and management, fraud monitoring, and customer-facing virtual assistants. However, while AI presents significant opportunities if deployed at scale, it also introduces data, legal, model conduct and governance challenges. Integrating AI across a complex and large banking organisation is not straightforward, and the risks it creates must be actively managed.

To provide C-suite executives and senior leaders with insights around the role of effective AI governance in managing these risks and how to assess current governance arrangements, 135 senior leaders were surveyed about their use of AI and governance arrangements. Further details about the survey are contained in Appendix A. The next chapter examines the current state of AI adoption across global banks and key risks that must be navigated to manage the integration. The remainder of the report focuses on the state of AI governance in banking and the benefits of improving AI governance before concluding with some recommended action items for banks.

Integrating AI across a complex and large banking organisation is not straightforward, and the risks it creates must be actively managed.



02

From experiments to scale: AI in the banking industry

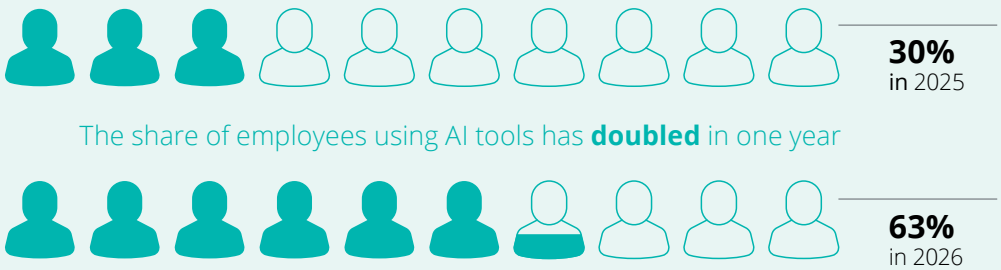
AI adoption in global banks has accelerated from early experimentation to enterprise scaling across core business functions. This creates significant opportunities to improve efficiency and service delivery while also expanding exposure to a variety of risks in real-time. As AI becomes more deeply embedded in day-to-day operations, the quality of governance will increasingly determine whether banks can scale AI systems appropriately and realise their benefits. This chapter explores the state of AI adoption across the banking industry and outlines some of the key risks banks are facing as they undertake full-scale deployment as the capabilities accelerate.

2.1 Accelerating AI adoption across global banks

Global banks are well placed to capture value from AI due to the data intensive and rule-based operations of the industry.[6] As a result, banks are deploying and scaling AI solutions across both operational and client facing activities. Bank of America's AI-powered virtual assistant, Erica, has handled over three billion client interactions since its launch, with more than 98% of users able to access the information they needed within 44 seconds on average, contributing to reduced call centre volumes and improved efficiency and customer experience.[7] More generally, the rapid adoption has seen the share of bank employees using AI doubling in just one year from 30% to 63% in 2026 (Figure 2.1).[8]

Moving from pilot to production is a critical step in capturing value from AI. Our survey of global banks indicates that progress is being made in integrating AI solutions across the business. These AI solutions are more likely to be fully implemented in business areas such as customer service, IT and marketing and sales teams, while business units such as Human Resources and legal and compliance business areas are more likely to be in the trial phase (Chart 2.1). On average, global banks are trialling AI solutions in four out of the 10 business areas presented while a further five business areas have AI solutions that are fully implemented.

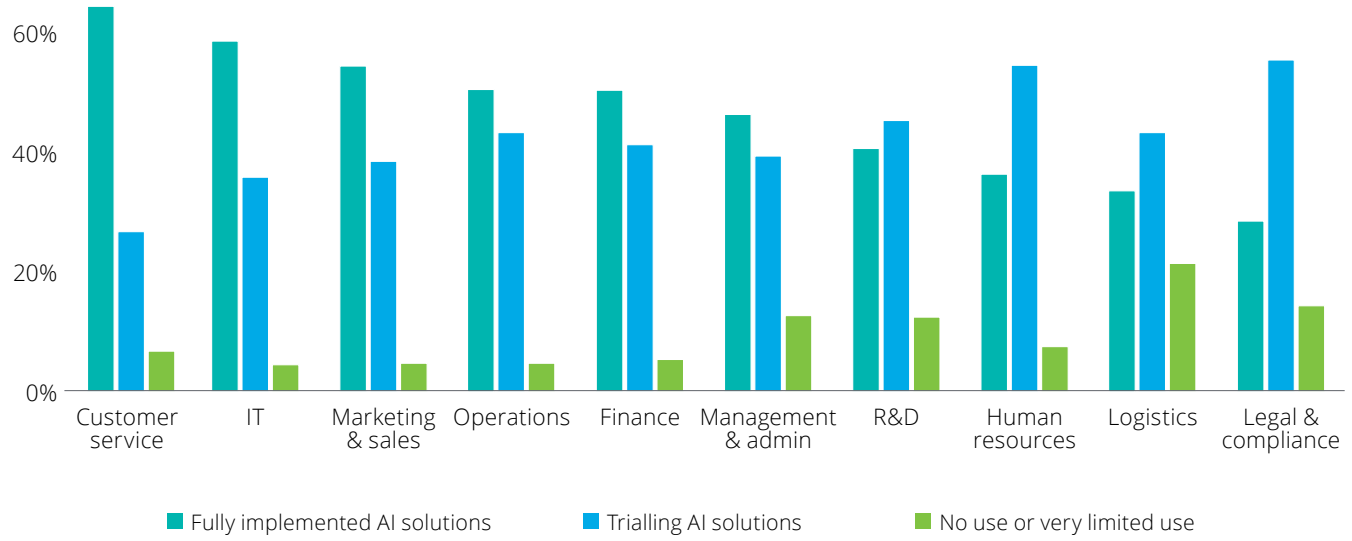
Figure 2.1 Share of employees using AI for work on a weekly basis



Source: Deloitte AI Governance in Banking Survey (2026) and State of AI in the Enterprise (2026)[2]

[2] The 2025 figure represents the numbers of workers in financial services more broadly.

Chart 2.1 Implementation of AI Solutions, by business area



Source: Deloitte AI Governance in Banking Survey (2026)

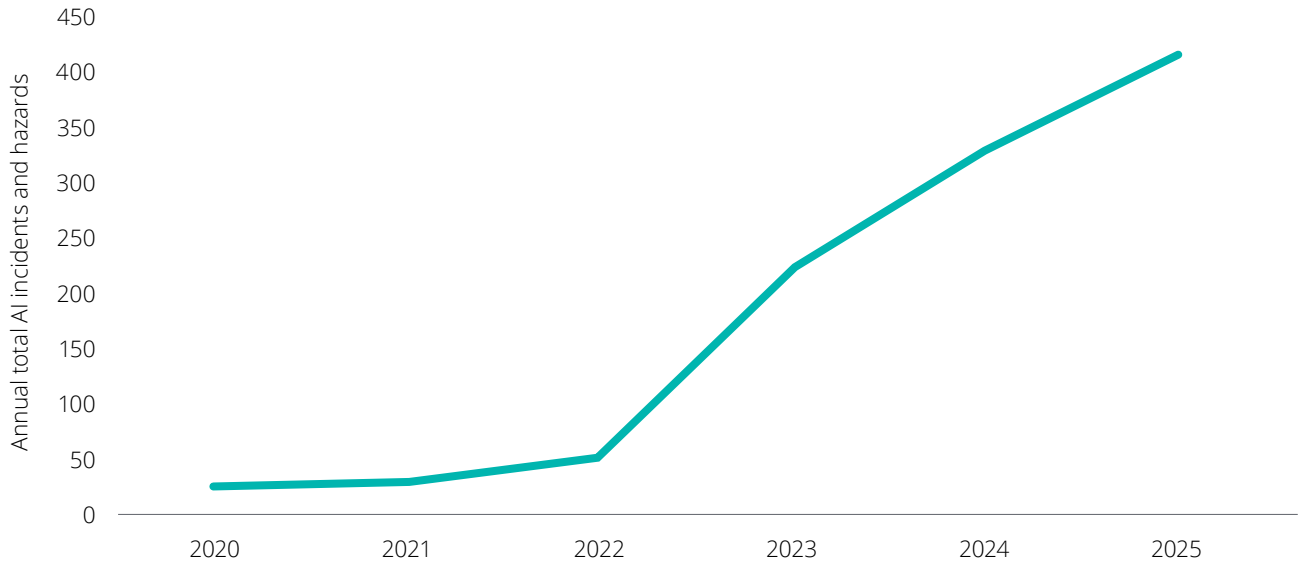
AI adoption is likely to accelerate and deepen as capabilities of AI become more advanced. This is particularly evident in agentic AI, which has the potential to transform the banking industry, acting as an autonomous workforce to make decisions and execute complex processes quicker than humans ever could. Banks such as JP Morgan are already using AI across a vast array of functions, including actively deploying and scaling AI for fraud detection, operational efficiency and client experience among others.[9]

The rapid uptake of AI is increasing banks' exposure to risk at the same time as it creates new sources of value. As AI systems are embedded more deeply within business processes, banks face heightened exposure to data, legal, regulatory, cyber, privacy and model risks. For cyber in particular, the use of AI has simultaneously improved defence as well as threats as the scale, speed and sophistication of AI-enabled cyber threats become more prevalent.

This shift is particularly significant because, historically, discovering and exploiting serious cyber vulnerabilities required rare expertise, deep system knowledge and considerable time. AI is changing that dynamic. Each successive generation of frontier models appears increasingly capable of identifying and exploiting vulnerabilities with greater speed and ease. Some of these AI solutions have discovered long-standing vulnerabilities that have survived decades of human review and automated tests.[10] For banks, this means cyber security capabilities will need to evolve in step with the growing speed, scale and complexity of AI-driven threats.

The Organisation for Economic Co-operation and Development (OECD) AI Incidents and Hazards Monitor indicates that the financial services industry has experienced one of the fastest growth rates in annual AI-related incidents and hazards, compared with other industries over the past four years. Financial services-related incidents have more than doubled as a share of total incidents, increasing from 1.8% in 2021 to 5.2% in 2025. As of May 2026, the sector has recorded 247 publicly reported incidents, with incident volumes already 20% higher than in the same period in 2025, and 72% higher than in 2024 (Chart 2.2).[11]

Chart 2.2 Publicly reported AI-related incidents and hazards for financial services, 2020-2025



Source: OECD (2026)

Cost overruns for AI projects have also become more common as banks transition from pilots to enterprise deployment, leading to usage that is much higher than originally intended. Unlike traditional technology investments, AI introduces ongoing, usage-based consumption costs, which can scale rapidly and exceed spending expectations if not actively governed. Many organisations lack mature processes to monitor usage, attribute costs to business units, and provide transparency at the executive level, increasing the risk of cost overruns and inefficient spending.

2.2 The costs of weak AI governance

These incidents and other events that are not made public have material costs for banks. The lack of governance creates direct and indirect costs for banks, ranging from financial losses and regulatory exposure to reputational damage and slower strategic execution (Figure 2.2).

Figure 2.2 The costs of weak AI governance

Financial losses

Incorrect credit decisions, trading errors, or fraud detection failures could all lead to significant revenue loss and large write-offs for banks. Banks may also experience financial losses from customers switching due to AI or data incidents.

Regulatory fines and legal costs

Non-compliance with regulatory and legal obligations can lead to fines, penalties or legal action.

Reputational damage

Data privacy breaches, discrimination or deployment of unethical AI applications could erode customer trust and reduce brand value.

AI paralysis

Poor governance can lead to an internal lack of trust in the use of AI solutions and may impact planned rollouts of AI pilots or even suggestions to use AI.

Source: Deloitte (2026)

The most apparent cost of weak AI governance is direct financial losses. Where AI is deployed without appropriate controls, banks may face system errors, privacy breaches, misleading outputs, misuse, and operational disruption. In 2025, the average cost of a data breach in financial services was US$5.56 million.[12] The emergence of 'shadow AI' – the use of AI tools without formal approval or oversight – can further increase exposure. Where shadow AI was present in incidents, the average cost increased by US$670,000.[13] The research found nearly all (97%) of AI-related security breaches involved systems lacking proper access controls, and most breached organisations reported having no governance policies in place to manage AI or prevent shadow AI.

Weak governance can create legal and regulatory exposure, which is compounded for global banks by the variation across jurisdictions and the speed of change. While global banks have always faced strict regulatory requirements, regulation related to AI is relatively unique in how requirements are changing in response to quickly growing capabilities and emerging use cases.

International financial regulatory bodies are increasingly recognising the risks associated with AI adoption and are developing guidance to support its safe and responsible use within the banking sector. The Financial Stability Board released its AI governance framework 'Sound Practices for Responsible Adoption of Artificial Intelligence (AI)' for financial services in June 2026, with twelve practices across two key pillars; governance and AI lifecycle management. It focuses on governance outcomes rather than specific models, and outlines qualitatively different risks associated with agentic AI.[14] The International Organization of Securities and Commissions published a Supervisory Toolkit for AI Use in Capital Markets, supporting the supervision and oversight of AI.[15] It covers the full lifecycle of AI systems and applies across system types, identifying potential risks such as implications for investor protection, market integrity and financial stability.

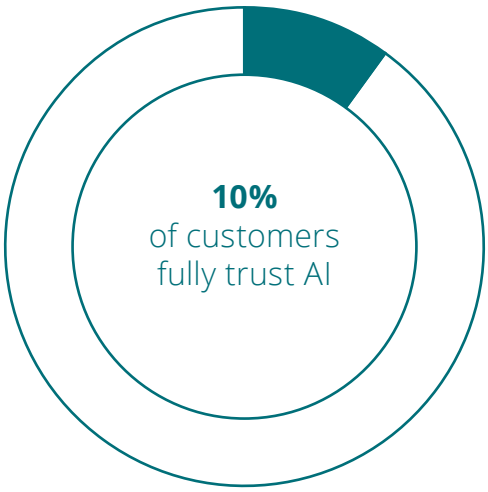
At the market level, approaches to AI regulation vary significantly across jurisdictions, with some markets adopting AI-specific regulation, introducing technology-neutral regulatory requirements while others have introduced policies that encourage adoption of AI in banking. Appendix C provides an overview of selected jurisdictions' AI regulatory frameworks.

Reputational damage is often the most visible consequence of AI failures. Single incidents involving data breaches, unfair outcomes, misleading recommendations, or inappropriate customer-facing interactions can erode trust quickly and attract regulatory, media and public scrutiny.[26] This risk is particularly significant in banking, where confidence, fairness and the secure handling of customer information are central to the value proposition.

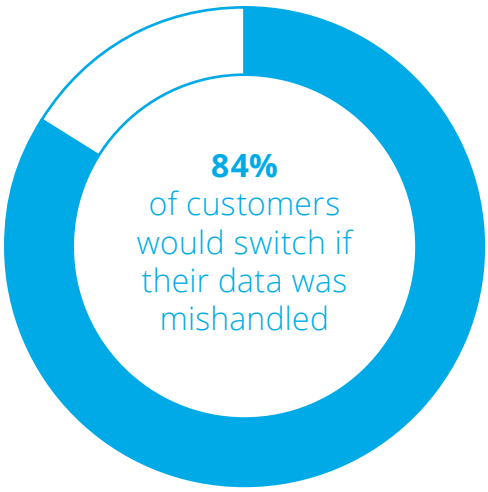
Reputational damage can quickly translate into customer attrition, investor concern, litigation and longer-term brand erosion, making it one of the most consequential costs of weak AI governance. For example, a survey of over 9,500 consumers found 84% would switch providers if their data was mishandled.[27] For global banks, realising the benefits of AI will therefore depend not only on technical capability, but also on the ability to demonstrate that AI is governed responsibly, transparently and securely.

Beyond direct financial and regulatory costs, weak AI governance can lead to AI paralysis. Where governance is unclear or underdeveloped, leaders may hesitate to scale AI because they lack confidence in the organisation's ability to manage the risks. This can slow adoption, reduce returns on investment, and leave banks at risk of falling behind competitors. Conversely, prolonged hesitation may also lead to rushed deployment later under competitive pressure, without the safeguards needed for safe and effective implementation. In this way, poor governance can result in both excessive caution and insufficient control.

Figure 2.4 Consumer attitudes and trust towards AI in financial services



Most customers trust their financial service provider less than a year ago amid concerns about data security and use.



This is a higher share than who would switch 12 months previously and provides a clear imperative for banks when considering their AI governance.

Source: Salesforce (2025)

03

What is the state of AI governance in banking?

Developing trustworthy AI is essential for banks to navigate the risks that come from rapid AI adoption and fully leverage this transformative technology. Trustworthy AI provides a level of certainty that technology is ethical, lawful and technically robust. Developing Trustworthy AI does not happen automatically, banks must have robust AI governance to provide the structure that ensures AI solutions deliver these outcomes.

To assess AI governance in a holistic and standardised way, we have developed a bespoke Trustworthy AI Governance Index tailored for banks. The Index scores each bank across five pillars: principles and policy, organisational structure, procedures and controls, people and skills and monitoring, reporting and evaluation. Together, these five pillars and their underlying 41 indicators are used to generate a single score that categorises AI governance as ad hoc, rudimentary, established and optimised, as shown in the table below. The Index gives senior leaders a single, comparable view of how ready their bank is to use AI safely and at scale.

Further detail on the methodology to develop the Trustworthy AI Governance Index is contained in Appendix B.

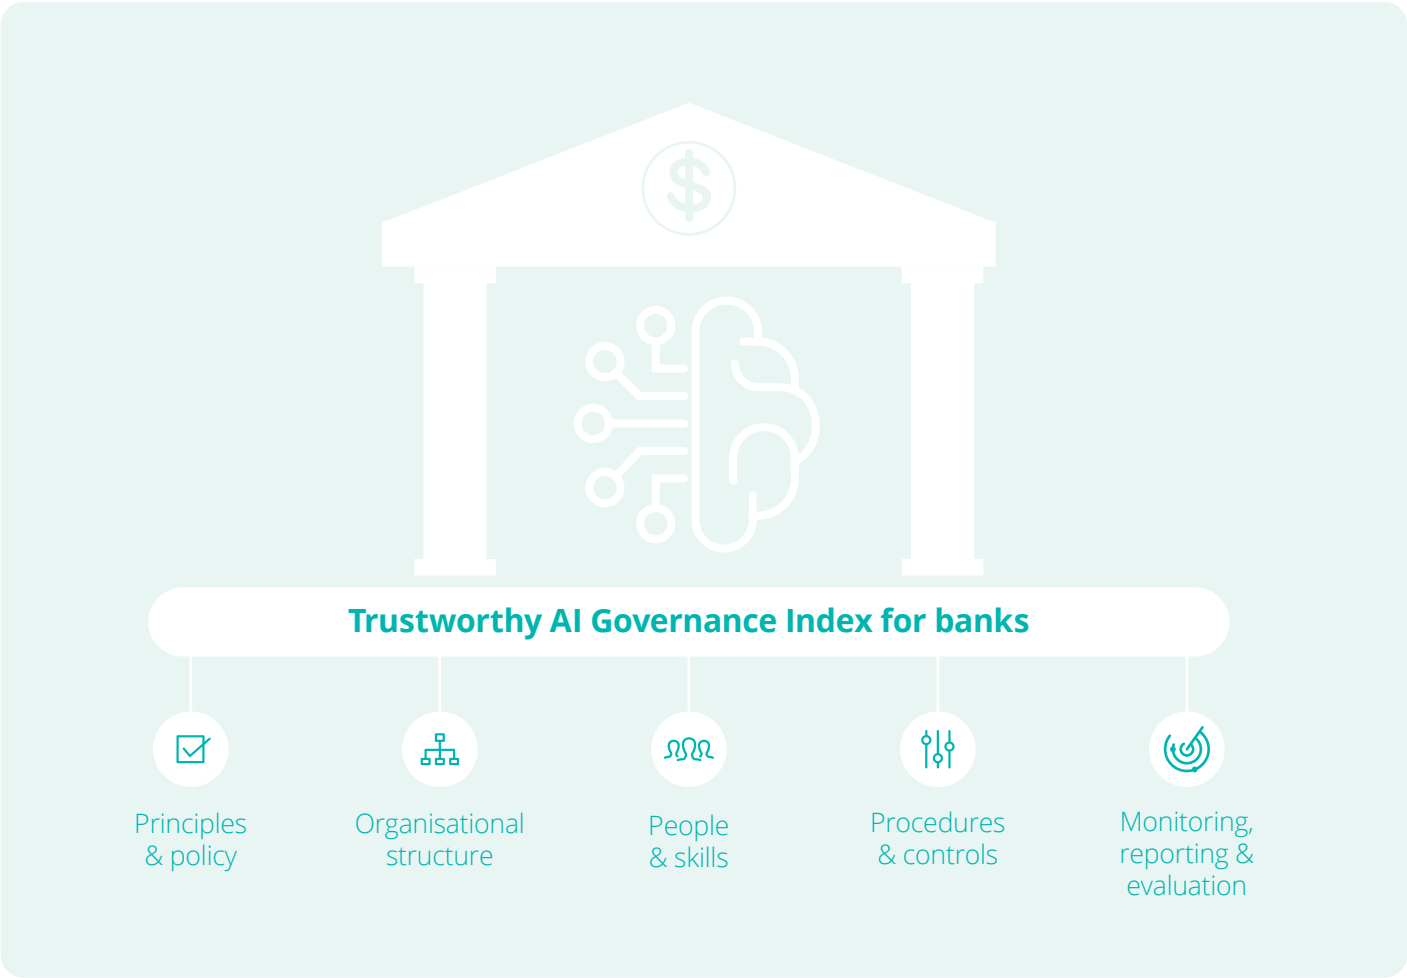


Table 3.1 Deloitte's AI Governance Index

	What this pillar measures	Ad hoc	Rudimentary	Established	Optimised
Principles & policy	*The formal rules, ethical principles and policies set out on the use of AI and the extent to which they guide AI risk management and decision making.*	No AI policy is in place, and responsible AI principles are not defined.	High-level AI policy and responsible AI principles exist but are incomplete and inconsistently applied.	A formal AI policy and responsible AI principles are defined and are referenced in decisions regarding AI use, investment and risk decisions.	AI policy and responsible AI principles actively shape culture and decision making. AI policy is responsive to emerging risks and expectations.
Organisational structure	*The articulation of roles, accountability and ownership of decisions and actions related to AI governance and risks.*	Roles and responsibilities for AI risk management are not defined for delivery teams or risk/compliance teams. Executive leadership do not take ownership.	Ownership is informally assigned for project teams and oversight from risk/compliance teams is unclear or ad hoc. Executive leadership is aware of risks but accountability is unclear.	Responsibilities and roles for project teams, risk/compliance teams and executive leadership are documented but not consistently applied in practice.	Delivery and oversight roles are fully integrated into governance processes and enable proactive risk management. Executive leadership consistently and visibly exercise accountability for AI risk decisions.
Procedures & controls	*The mechanisms and controls applied to AI systems to manage AI risks and outcomes across the design and implementation lifecycle.*	AI risks are not addressed, and formal mechanisms have not been developed to manage AI systems.	AI risks are considered informally or reactively. Some mechanisms and controls are defined for AI systems.	Responses to AI risks are defined formally and are applied consistently. Mechanisms and controls are defined but inconsistently applied across AI systems and across the lifecycle.	A proactive approach to AI risks is embedded in all work practices. Mechanisms and controls are defined and consistently applied across AI systems and across the lifecycle.
People & skills	*The capabilities and knowledge of delivery teams and AI risk and governance roles and associated training.*	AI governance is handled by traditional risk or compliance staff with little AI/ML literacy. AI risk management training is not provided.	Most AI governance roles are allocated to traditional risk staff. Ad hoc training is provided when required.	AI governance roles mostly have strong AI literacy. Role specific AI risk management training is provided across departments.	A multidisciplinary team combines deep AI expertise with risk capability. AI risk management training is regularly refreshed across the organisation.
Monitoring, reporting & evaluation	*The ongoing monitoring and tracking of AI system outcomes and risks using key metrics or indicators.*	AI risks are not identified, measured or tracked and review processes are not updated.	Risk metrics are used inconsistently and are not fit-for-purpose. Reviews occur irregularly and usually after issues or incidents arise.	Fit-for-purpose risk metrics are defined but inconsistently applied across teams. AI risk management process is reviewed using defined roles and schedules.	Risk metrics are constantly reviewed and updated to remain best practice. AI risk management is continuously reviewed and improved based on performance.

Source: Deloitte (2026)

3.1

Nearly all banks can improve their AI governance

Our Trustworthy AI Governance Index shows that the majority (87%) of global banks have clear scope to strengthen how they govern AI in their organisation. More than two thirds of banks are classified as rudimentary or established, while only 13% are at the optimised level. More than one in 10 have governance arrangements that remain ad hoc (Chart 3.1).

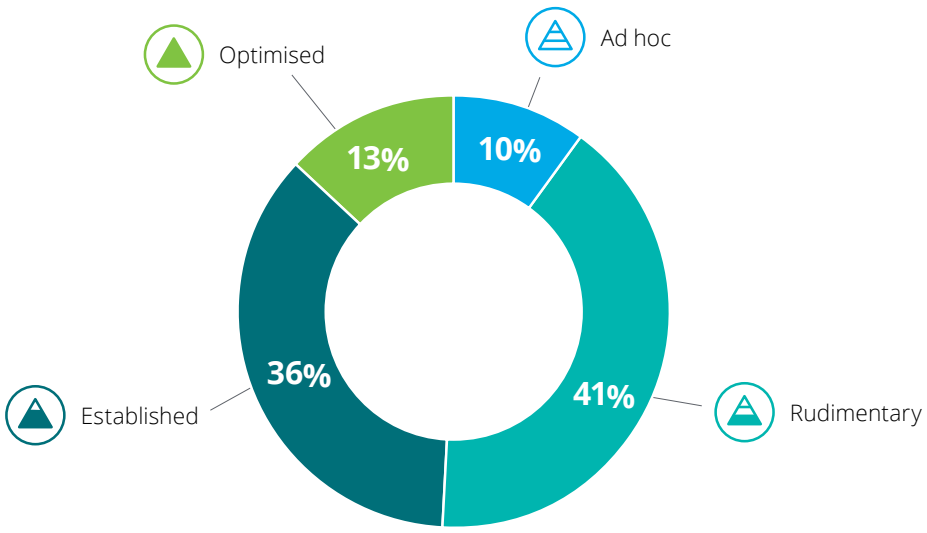
This distribution suggests that, although many global and major banks have established some elements of AI governance, only a minority have the foundations, accountabilities and control settings needed to support regulatory compliance, promote the appropriate use of AI, and respond effectively when issues arise.

Looking across the underlying pillars of AI governance, organisational structure is the weakest area for most banks. This pillar is critical because effective AI governance depends on clear accountability, defined decision-making rights and visible ownership across the organisation. More than a quarter of banks (27%) report that their organisational structure for AI governance is ad hoc. There is substantial room for improvement in the people and skills pillar, which reflects whether employees and leaders have the capability to apply governance in practice. More than half of banks are categorised as either ad hoc or rudimentary in this area, indicating that capability uplift has not occurred or is incomplete.

By contrast, procedures and controls is the strongest pillar, with nearly a quarter of banks achieving optimised status. Principles and policies, and monitoring, reporting and evaluation, also perform relatively well. This could indicate that banks have established centralised capabilities, potentially in relatively small teams or through pilots only (e.g. AI Safety, AI development and monitoring, AI and Model Risk Management), with the next step being to scale controls and good AI risk management practices across the organisation.

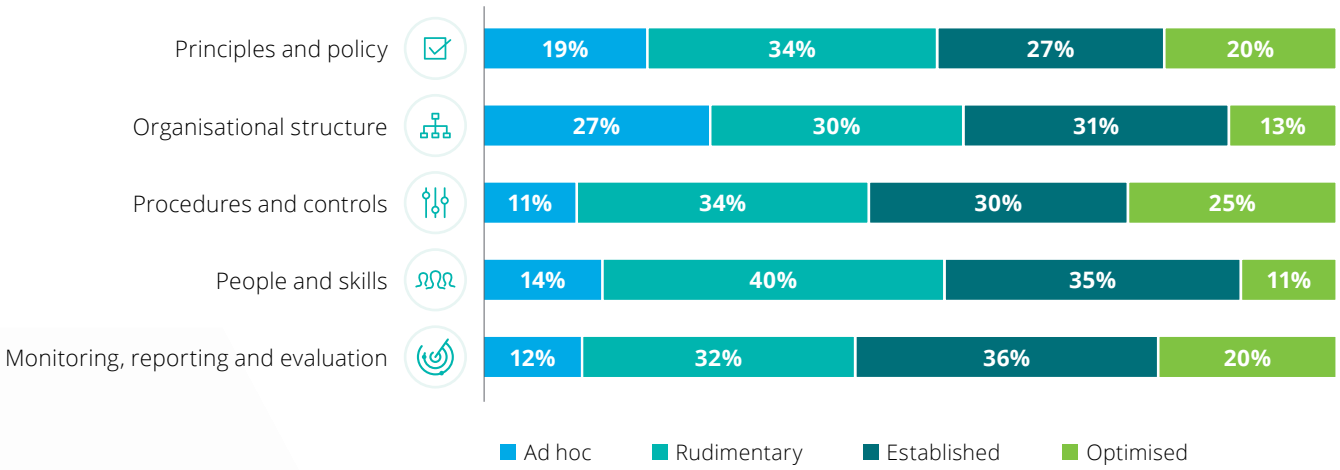
Taken together, these findings indicate that the main gaps in AI governance are human and organisational rather than procedural. Many banks have begun to establish policies and controls, but fewer have equipped their people with the necessary skills or clarified who is responsible and accountable for ensuring appropriate AI outcomes.

Chart 3.1 Global banks by Trustworthy AI Governance category

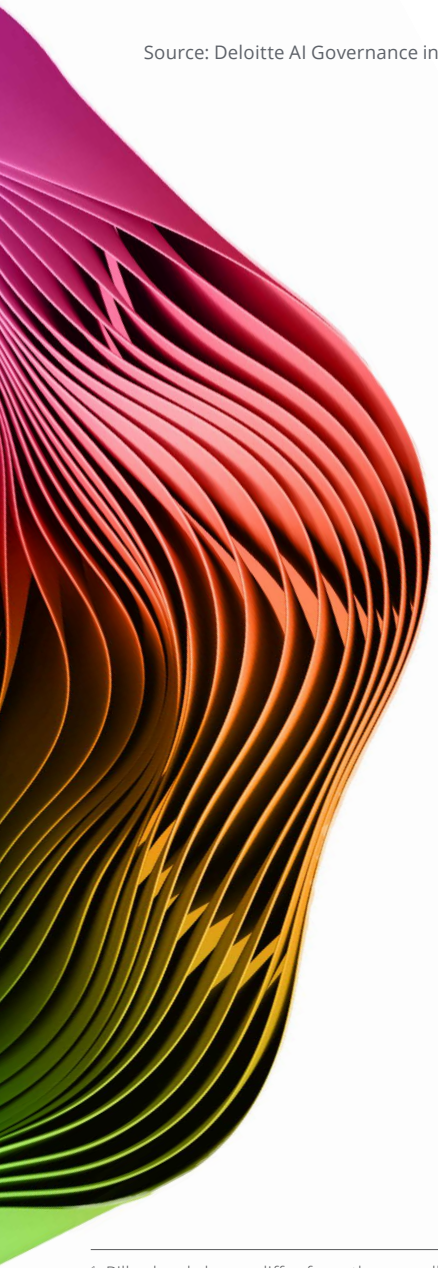


Source: Deloitte AI Governance in Banking Survey (2026)

Chart 3.2 Distribution of AI governance across global banks, by pillar[1]



Source: Deloitte AI Governance in Banking Survey (2026)



These findings indicate that the main gaps in AI governance are human and organisational rather than procedural

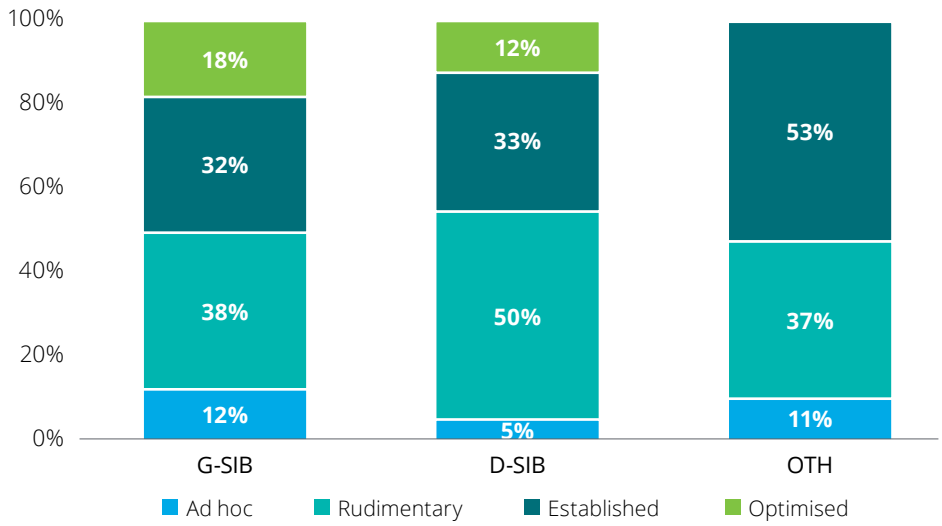
[1] Pillar-level shares differ from the overall governance results as the overall results rely on performance across the five pillars. For instance, the share of banks with "optimised" levels in individual pillars is higher than the share of banks achieving that level overall because banks may perform well in individual pillars, whereas achieving an optimised overall rating requires consistently strong performance across all pillars.

How do G-SIBs and D-SIBs compare with other major banks?

G-SIBs and D-SIBs sit at the centre of the global financial system and operate under the most intensive supervisory expectations for model risk, operational resilience and the use of AI. Their approach to AI governance therefore matters not only for their own institutions, but as a benchmark for the wider banking sector.

Our results show that when comparing AI governance maturity of G-SIBs and D-SIBs to other major banks, the results are mixed (Chart 3.3). Around one in six G-SIBs (18%) and one in eight D-SIBs (13%) reach the optimised tier, while no other major bank surveyed achieves this level of maturity. In contrast, the share of G-SIBs and other major banks at an ad hoc level of AI governance are similar in level.[3]

Chart 3.3 Distribution of AI governance across global banks, by tier



Source: Deloitte AI Governance in Banking Survey (2026)

The governance gap is evident particularly in principles and policy and organisational structure. G-SIBs and D-SIBs are more likely than other major banks to have embedded responsible AI principles in day-to-day decision-making, to align governance with ethical, legal and regulatory requirements, and to define roles clearly between AI delivery teams and risk and compliance teams.

The governance gap is narrower in some more operational areas. The three cohorts are relatively closer on the share of staff receiving regular refresher training and on having mandatory governance mechanisms in place for generative and agentic AI across the implementation lifecycle. In some cases, the pattern even reverses: other major banks report a slightly higher average share of AI use cases captured in a central register than G-SIBs and D-SIBs (73% compared with 68%).

Overall, the findings suggest that G-SIBs and D-SIBs have stronger overall foundations for AI governance, especially in principles, oversight and accountability. As AI adoption deepens, G-SIBs and D-SIBs will need to convert stronger governance foundations into more consistent execution, while other major banks will need to close the governance gap to remain competitive in a more demanding operating environment.

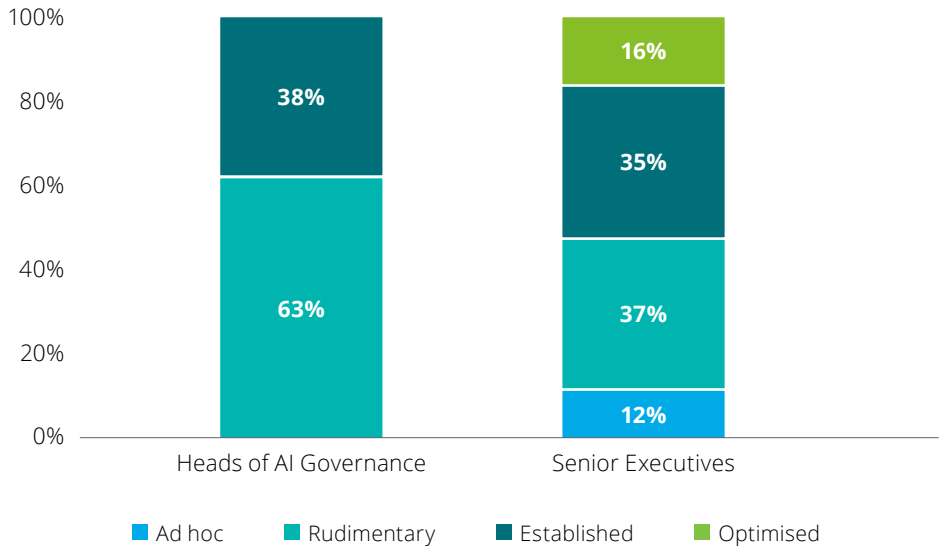
[3] These results should be interpreted with caution as G-SIB, D-SIB and other major bank classification was self-identified by the external survey provider and the 'other major banks' cohort is based on a relatively small sample size of 19 responses.

How perspective shapes views of AI governance maturity

Assessments of AI governance maturity can differ materially depending on who is asked. The analysis in this report draws on two respondent cohorts: 24 intensive interviews and surveys with Heads of AI governance at G-SIBs and D-SIBs, and 111 survey responses from senior executives across the broader bank sample. The comparison suggests that proximity to governance in practice shapes how maturity is perceived.

Heads of AI governance gave much more concentrated assessments of their banks' overall maturity than the broader senior executive sample. All Heads of AI governance placed their bank in the middle of the maturity index, with 63% rating governance as rudimentary and 38% as established. None rated their bank as either ad hoc or optimised. By contrast, senior executive responses were spread across all four tiers, with 12% rating their bank as ad hoc and 16% as optimised.

Chart 3.4 Distribution of AI governance between respondent groups, by tier

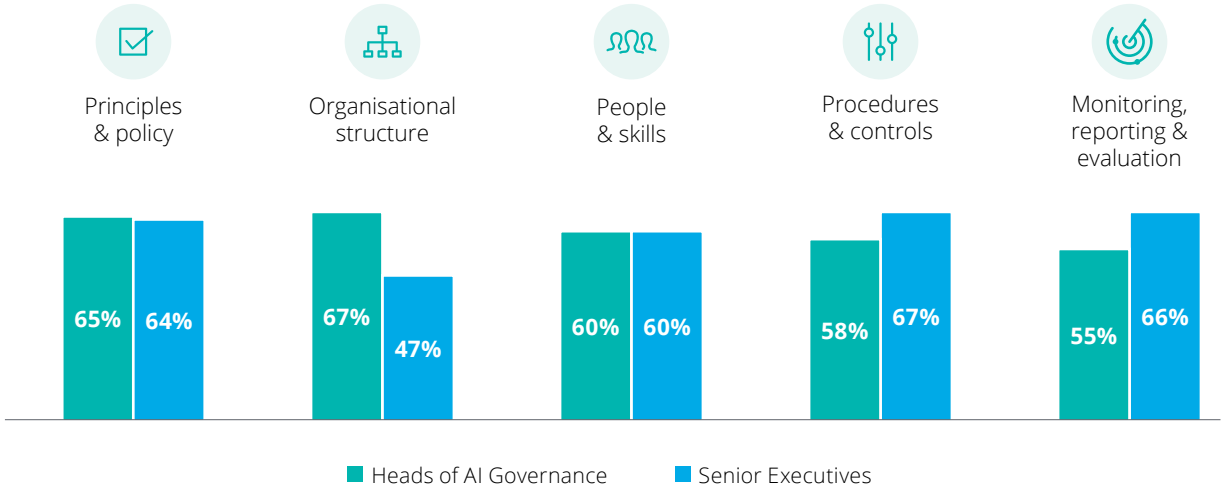


Source: Deloitte AI Governance in Banking Survey (2026)

The two cohorts also differed in how they assessed maturity across the five pillars of the Trustworthy AI Governance Index. Average scores were similar for principles and policies and for people and skills, suggesting broad alignment on the strength of formal policy settings and workforce capability. The largest difference was in organisational structure, where heads of AI governance scored their banks around 20 points higher than other senior executives on average (67 compared with 47), potentially reflecting greater visibility of roles, responsibilities and accountabilities.

In contrast, Heads of AI governance reported lower scores on the more operational pillars: procedures and controls (58 compared with 67) and monitoring, reporting and evaluation (55 compared with 66). This may reflect a more cautious view of how governance operates in day-to-day practice for those responsible for AI governance, particularly where formal arrangements are harder to execute consistently at scale.

Chart 3.5 Average AI Governance Index score between respondent groups, by pillar



Source: Deloitte AI Governance in Banking Survey (2026)

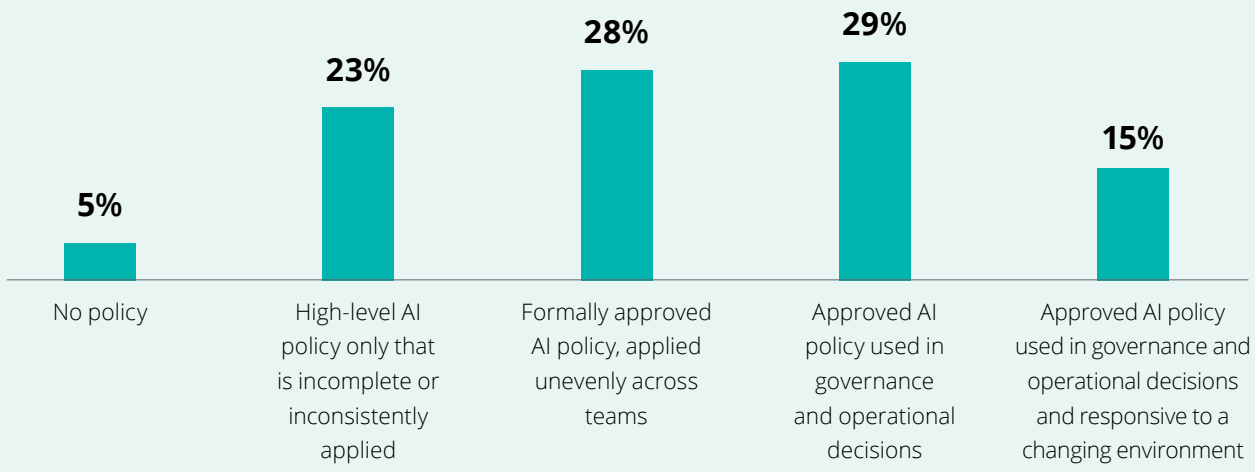
Overall, the findings suggest that a bank's view of its own AI governance maturity depends in part on respondent perspective. Those closest to governance execution appear less likely to overstate maturity and more likely to identify weaknesses in implementation. For boards and senior leaders, the implication is straightforward: understanding the true state of AI governance requires input not only from senior decision-makers, but also from those responsible for making governance work in practice.

3.2
Pillar 1 | Principles and policies

Clear and well-understood principles and policies are a prerequisite for effective AI governance. They set expectations for how AI should be designed, deployed and overseen, and provide a foundation for consistent decision-making across the organisation.

While most banks have an AI policy in place, many apply it inconsistently across the organisation or between teams. For example, a quarter of banks have a high-level AI policy that may be incomplete, and a further 28% have an AI policy that is approved but used inconsistently across teams (Chart 3.6). This limits the policy's ability to shape behaviour and influence outcomes. Only 15% of banks have an AI policy that actively informs AI strategy and is responsive to emerging risks and stakeholder expectations.

Chart 3.6 Maturity of banks' AI policies

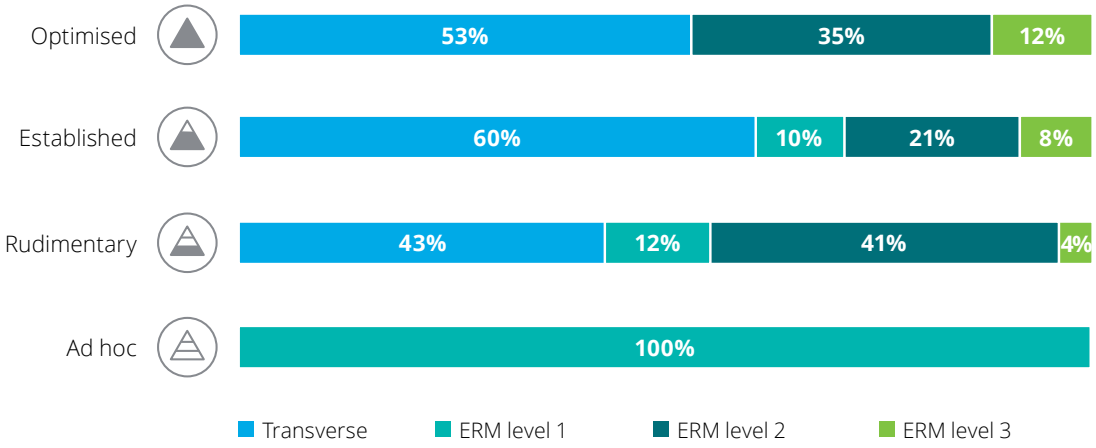


Source: Deloitte AI Governance in Banking Survey (2026)

How a bank classifies AI risk appears to be related to the maturity of its governance. This matters because risk classification influences who is accountable for oversight, and whether AI-related issues are managed in an integrated or isolated way. Treating AI as a transverse risk which typically supports a more integrated governance approach, as it recognises that AI-related risks cut across multiple risk classes. The alternative is to treat AI as its own risk class within the taxonomy, which is duplicative and does not recognise AI as a driver of existing risks.

Banks that have optimised levels of AI governance are more likely to classify AI as a transverse risk, embedded across existing risk domains such as model, data, cyber and legal risk, than banks that treat AI as an Enterprise Risk Management (ERM) Level 1, 2 or 3 in their enterprise risk taxonomy (Chart 3.7). Just over half of optimised banks (53%) classify AI as a transverse risk driver, with most of the remainder treating it as an ERM level 2 (24%) or Level 3 (19%) risk. By contrast, rudimentary banks are less likely to classify AI as a transversal risk driver (43%), and banks in the ad hoc category (100%) treat AI as an ERM Level 1 risk. This suggests that as AI governance maturity increases, banks are more likely to embed AI within broader cross-cutting risk frameworks.

Chart 3.7 AI governance maturity, by how AI risk is classified



Source: Deloitte AI Governance in Banking Survey (2026)
Note: ERM Level 1 captures broad enterprise risk categories, such as operational, business, strategic, reputational, financial and ESG risk; Level 2 defines standardised sub-risk categories that enable aggregation and executive-level reporting; and Level 3 captures more granular, business-specific risks that map back to Level 2 while reflecting local terminology and context.

3.3
Pillar 2 | Organisational structure

A clear organisational structure is fundamental to effective AI governance because it defines who is accountable for decisions, how issues are escalated, and where responsibility sits when risks emerge. Without clear ownership, even well-designed policies and controls can be applied inconsistently or seen as a tick box exercise.

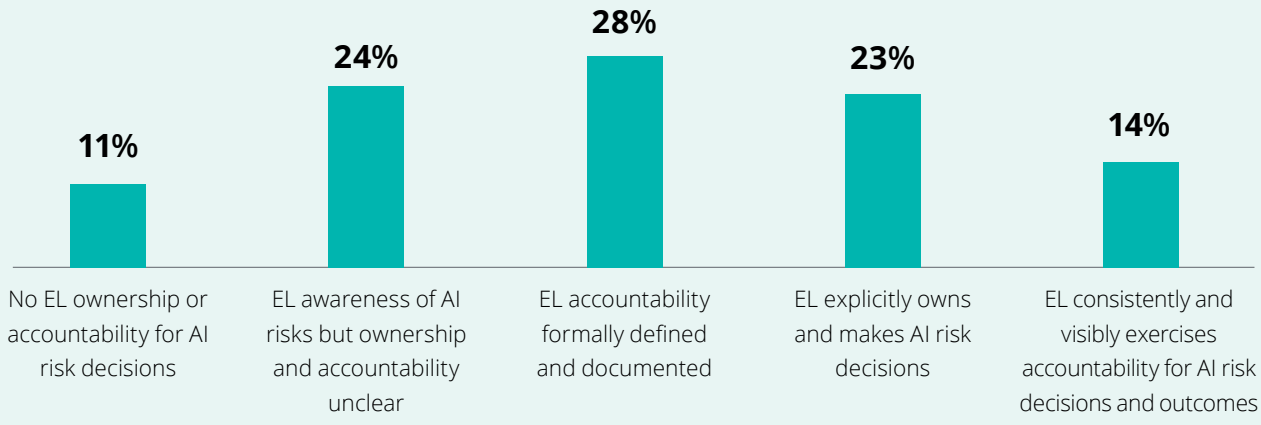
One key component of organisational structure is executive leadership ownership and accountability for AI risk. This matters because visible ownership at senior levels helps ensure AI-related decisions are governed consistently and that emerging issues are identified, escalated and addressed appropriately. Only 37% of banks report that senior leaders explicitly own or visibly exercise accountability for AI risk decisions (Chart 3.8). A further 28% report that accountability is formally assigned within the executive team, but not necessarily demonstrated through visible leadership or decision-making. In the remaining 36% of banks, ownership and accountability are either unclear or not formally defined.

Another important question for organisational structure is who leads AI governance within the organisation. This matters because the location of leadership can influence how strategically AI is governed, how effectively decisions are coordinated across functions, and how clearly accountability is established.

The survey points to a clear pattern. Among banks at the optimised level of AI governance, leadership most commonly sits with a Chief AI Officer or Head of AI, or with the CIO or CTO, each optimised AI governance at 29% of optimised banks (Chart 3.9). A further 17% of optimised banks are led by a Chief Risk Officer. Banks at the rudimentary level look markedly different with half (50%) being led by a Chief Data or Analytics Officer, a quarter (25%) by a Chief Risk Officer, and only 8% by a Head of AI.

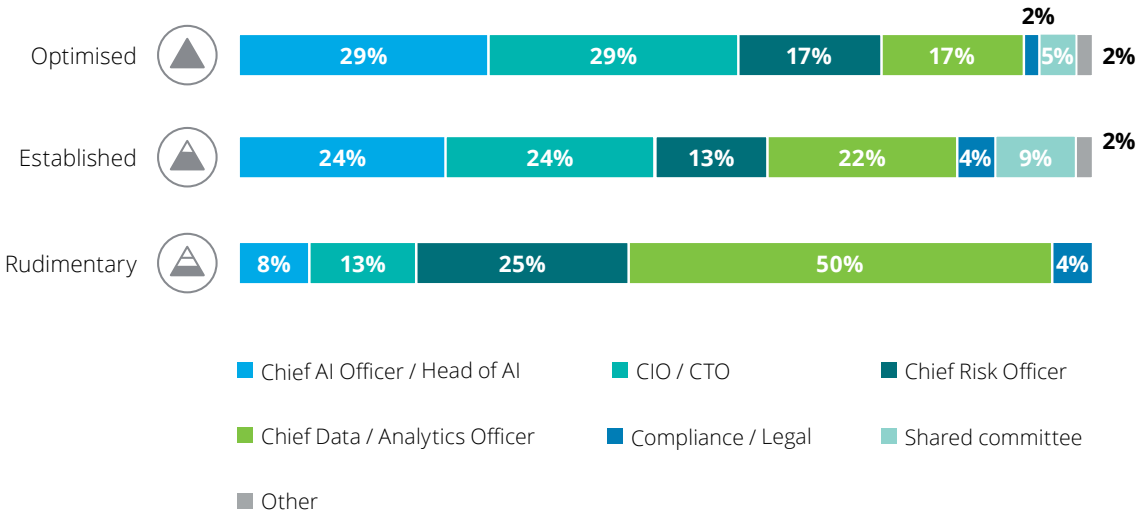
As AI systems become more autonomous and agentic, the First Line of Defence will increasingly need to shift from primarily human-led oversight to technology-enabled governance, with embedded controls, continuous monitoring and automated decision oversight operating at the point of use. Despite tech-enabled governance playing an increased role, human First Line of Defence roles will also need greater fluency of what responsible AI use means for their role and application. The growth of both AI adoption and technology-enabled governance will likely make the impact of a dedicated AI leader or executive even more pronounced over time.

Chart 3.8 Executive leadership ownership and accountability of AI



Source: Deloitte AI Governance in Banking Survey (2026)

Chart 3.9 Who leads AI governance, by governance maturity level



Source: Deloitte AI Governance in Banking Survey (2026)

This suggests that higher levels of AI governance maturity are more often associated with organisational models where a senior, dedicated AI leader or executive is in the lead. By contrast, banks where AI leadership sits primarily within the data or analytics function are more commonly found at lower levels of AI governance maturity.

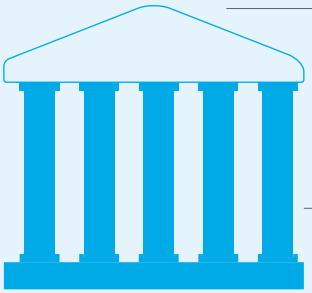
3.4
Pillar 3 | Procedures and controls

Procedures and controls translate high-level AI principles into operational practice. These matter because they help ensure that AI systems are subject to consistent oversight and risk management across the lifecycle, from initial design through to ongoing monitoring and into production. This is one of the strongest areas of AI governance for global banks on average. Nearly six in 10 banks apply a mandatory and consistent approach to determining how much oversight each AI system requires across all stages of the AI lifecycle: design, build, test, deploy and monitor.

However, these controls are applied more consistently in the earlier stages of the AI lifecycle than in the later ones. Around seven in 10 banks apply mandatory governance at the design stage, and roughly two-thirds do so at the build, test and deploy stages (Chart 3.10). This falls to just over half (55%) at the monitoring stage. This gap matters because governance cannot end once a system is deployed. As AI pilots are increasingly moved into production, consistent mandatory controls at the monitoring stage will be essential to identify performance issues, emerging risks and unintended outcomes over time.

Maintaining visibility over where AI is used across the organisation is another important test of governance maturity. Central use-case registers support oversight by helping banks identify different AI applications, where they sit, and whether they are operating within formal governance arrangements. Yet less than three in ten banks record more than half their AI use cases in a central register. This suggests that, in most banks, material pockets of AI activity may still sit outside formal oversight.

Chart 3.11 Share of AI use cases recorded in the central register

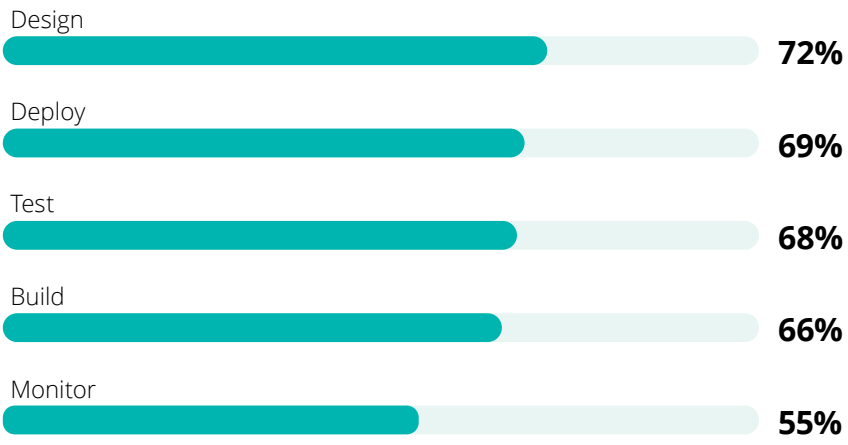


28%
of banks have over half of use cases recorded in the central register.

72%
of global banks have less than half of their AI use cases recorded in a central register.

Source: Deloitte AI Governance in Banking Survey (2026)

Chart 3.10 AI lifecycle phases with mandatory governance controls



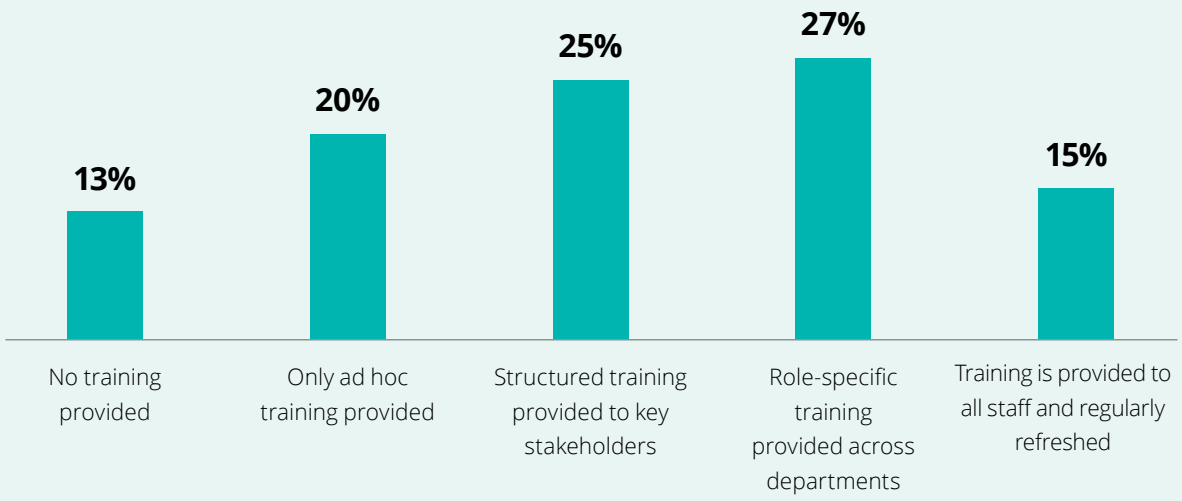
Source: Deloitte AI Governance in Banking Survey (2026)

3.5
Pillar 4 | People and skills

People and skills are central to trustworthy AI because governance only works when employees can apply it in practice. Staff need to understand how to identify AI-related risks, follow governance requirements and escalate issues appropriately. Without this capability, even well-designed policies and procedural controls can fail to produce consistent outcomes.

This remains one of the areas where banks have the most ground to make up. According to the Trustworthy AI Governance Index, half of global banks are at either ad hoc or rudimentary levels on the people and skills pillar. One area where this is apparent is AI governance training. While most banks provide some form of structured training on AI governance, the depth and reach of this training varies across organisations (Chart 3.12). However, one-third of global banks still rely on ad hoc training or provide no AI governance training at all. As AI use spreads across the bank, building consistent, role-relevant capability will be essential.

Chart 3.12 How banks train staff to manage AI risk



Source: Deloitte AI Governance in Banking Survey (2026)

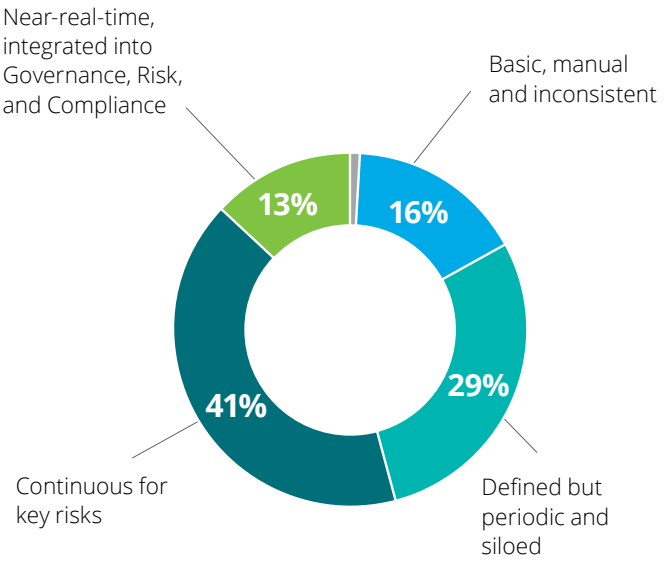
3.6
Pillar 5 | Monitoring, reporting and evaluation

Monitoring, reporting and evaluation becomes increasingly important as banks move from experimentation into scaled deployment. Effective governance requires ongoing visibility over how AI systems perform, what risks are emerging, and whether intended benefits are being realised in practice. The AI Governance Index finds that across the whole sample this pillar is relatively well established, with more than half (57%) reporting either established or optimised levels. This is the highest share of banks achieving these two categories out of any pillar. Yet as discussed in the box around perspectives of AI maturity, heads of AI governance were more likely to report lower levels of AI governance maturity than responses from senior executives, with an average score of 55 for heads of AI governance compared to 66 for senior executives.

Banks are comparatively confident once AI is live. More than half of the banks asked how effectively they monitor AI systems after deployment, reported continuous or near-real-time monitoring: 41% monitor key risks continuously against defined thresholds, and a further 13% run near-real-time monitoring integrated into enterprise risk systems. However, around a third (29%) have only periodic, siloed monitoring, and about one in six (16%) still rely on rudimentary, manual checks. Strong, continuous monitoring gives senior leaders the evidence to scale the AI that works and step back from what does not, and strengthening automated monitoring is the main opportunity here as AI scales.

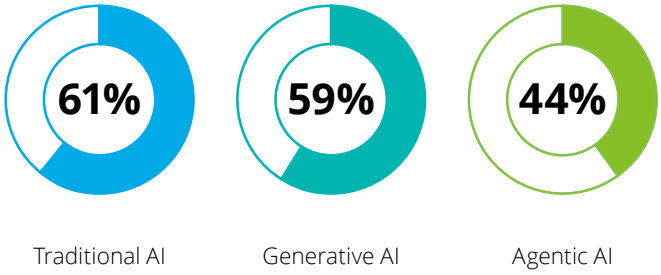
There are, however, clear opportunities to strengthen oversight of newer forms of AI. Maturity depends on the type of AI involved. While around six in 10 banks have risk management monitoring across the implementation lifecycle for traditional AI (61%) and generative AI (59%), this drops to 44% for agentic AI systems. This gap matters because governance frameworks that are effective for established AI applications may not yet be fully adapted to newer, more autonomous systems. These AI systems may operate continuously and do not possess human motivators such as fear of penalties or ethical considerations beyond existing guardrails. While current adoption of Agentic AI may be currently lower than traditional AI or Generative AI, as AI use cases spread across global banks, closing this gap by having careful guardrail design, real-time monitoring and human operators on-the-loop will be essential to maintaining effective oversight.

Chart 3.13 How effectively banks monitor AI systems after deployment



Source: Deloitte AI Governance in Banking Survey (2026)

Chart 3.14 Share of banks with risk monitoring across implementation lifecycle, by type of AI



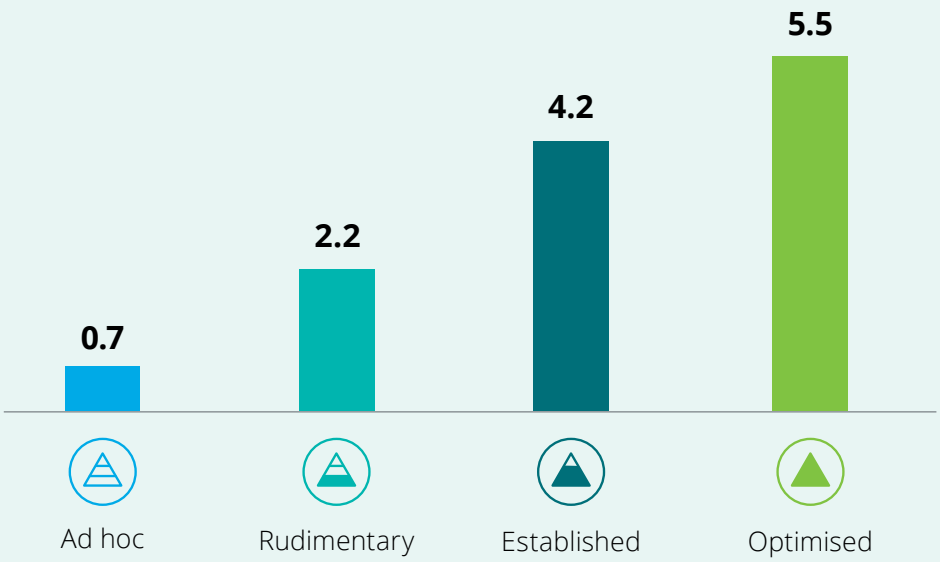
Source: Deloitte AI Governance in Banking Survey (2026)

04
The AI governance dividend

Stronger AI governance is associated with both operational and commercial benefits for banks. Rather than acting as a brake on innovation, governance appears to be an important enabler of AI deployment at scale. Among senior leaders surveyed, the most common benefits of more effective AI governance were faster deployment of AI solutions (48%), greater regulatory compliance (46%) and improved customer reputation (43%).

This link between governance and deployment is evident in how widely AI has been implemented across business units. Banks with optimised AI governance have nearly eight times more fully implemented AI solutions across their business areas – such as sales and marketing, operations, IT, legal and HR – compared with banks with ad hoc governance levels. There is also a substantial gap between optimised and rudimentary banks, with the most mature institutions having more than double the number of fully implemented AI solutions on average. Banks that invest in governance are not trading away adoption; they are creating the enabling conditions for widespread use. Good governance creates the pathway to production.

Chart 4.1 Number of fully implemented AI solutions across business units, by maturity

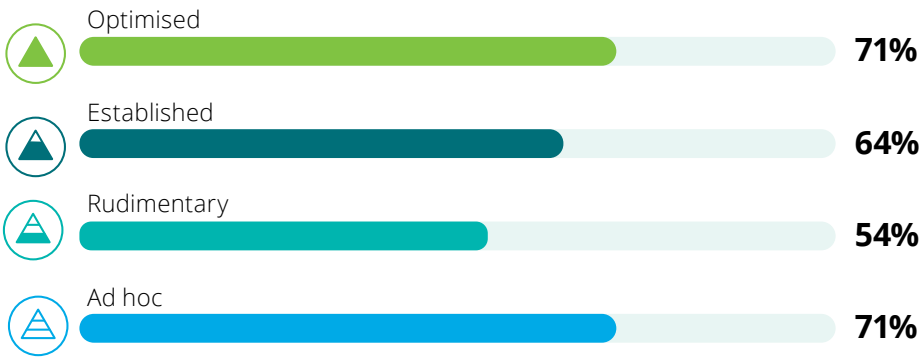


Source: Deloitte AI Governance in Banking Survey (2026)
Note: The numbers above are based on regression averages after controlling for workforce size and location of bank.

The relationship between governance maturity and the share of employees using AI is telling. Average weekly AI usage is lower among established (64%) and rudimentary (54%) banks than among optimised and ad hoc banks, which report similarly high usage rates of 71%. Lower levels of governance maturity can be difficult for employees to navigate, where they may decide that use of AI is not worth the compliance burden. Or they may circumnavigate poorly designed governance through use of shadow AI and therefore use is not visible to the bank. Optimised governance supports use through well designed processes and structured governance, reducing the number of blockers for using AI.

There is growing evidence that banks can improve financial performance by integrating AI into the business, but the size of the opportunity depends on how effectively it is deployed. Industry estimates suggest that generative AI could create between US$200 billion and US$340 billion in additional value for the global banking industry, equivalent to around 3% to 5% of industry revenue. This uplift is expected to come from a combination of higher revenue and lower operating costs.[28] For European banks, the immediate focus is on cost efficiency, with some estimates placing the savings at 4% to 9% of total operating costs, rising to much larger savings in structured, data-intensive processes such as know-your-customer checks, loan processing and internal reporting.[29]

Chart 4.2 Average weekly AI user share, by AI governance maturity



Source: Deloitte AI Governance in Banking Survey (2026)

Yet realising these revenue gains is far from automatic. A 2025 study by the Massachusetts Institute of Technology found that around 95% of enterprise generative AI investments had produced no measurable effect on profit.[30] Another study found that, among financial services businesses, only about 40% report higher profitability from AI so far, while a similar share report no change. This suggests that AI's commercial promise is real, but unevenly realised in practice.[31]

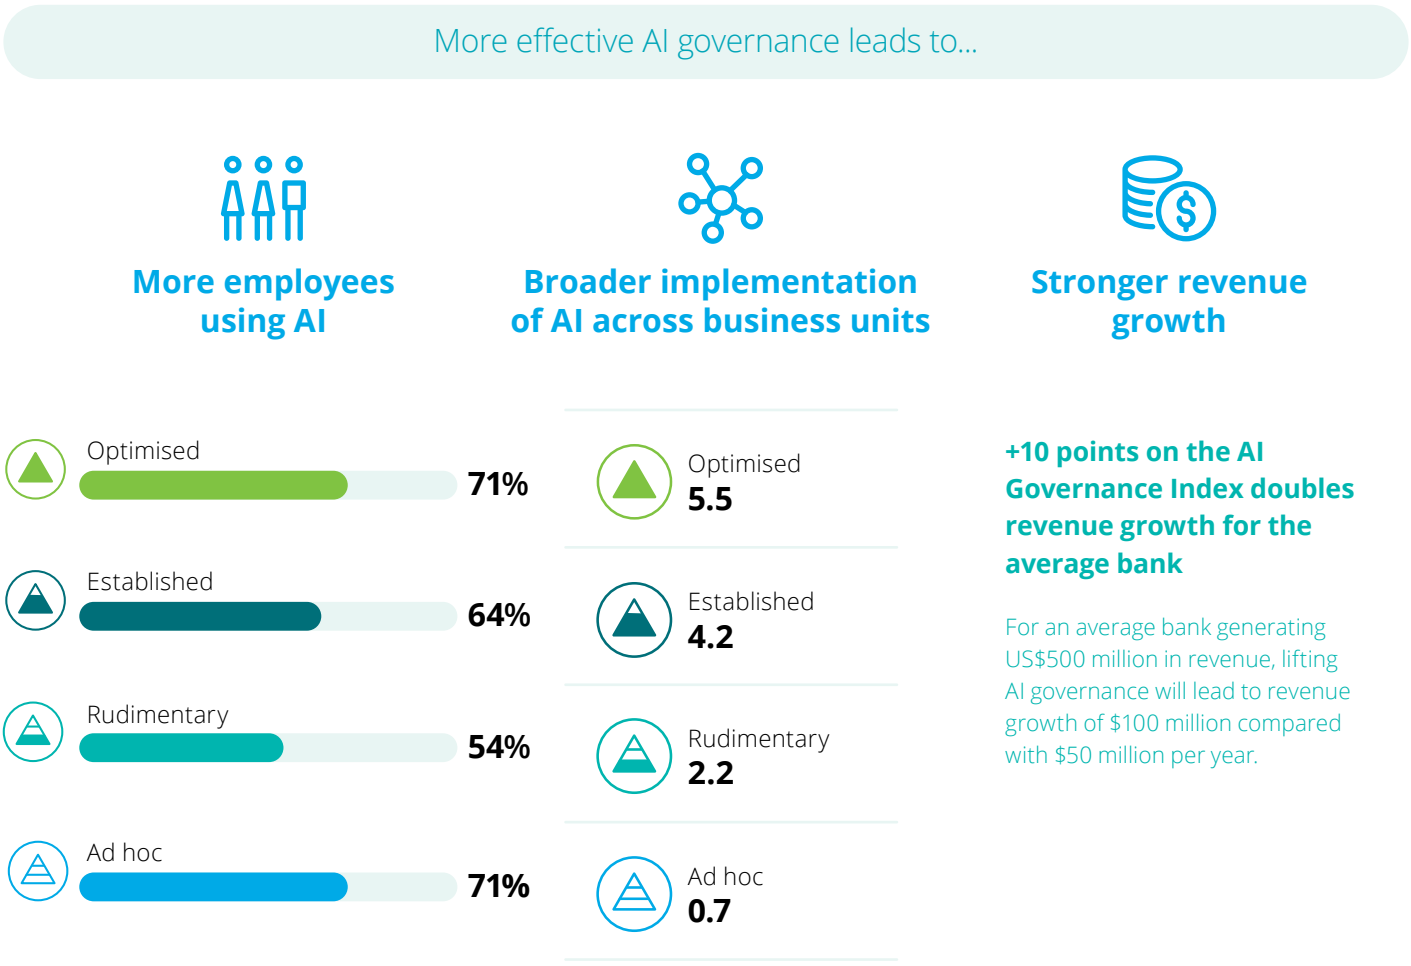
This is where governance matters. The financial value of AI depends not just on the technology itself, but on whether banks can deploy it at scale in a way that aligns with the business strategy, embed it into core processes and manage the associated risks. Banks with higher levels of governance could also achieve higher revenue as customers move away from businesses that experience an AI or data incident. Our analysis provides evidence to support this link. Econometric modelling indicates that banks with higher Trustworthy AI Governance Index scores tend to record stronger revenue growth, even after controlling for the extent of AI use, headquarters location, workforce size and years in operation. On average, a 10-point increase in the Index score is associated with a 10 percentage-point increase in revenue growth.

The economic significance of this relationship is material. For the average bank in our survey, revenue grew by 10% between FY2023–24 and FY2024–25. The modelling suggests that, if AI governance had been 10 points higher, revenue growth could have been around 20% over the same period from unlocking the potential of AI. For a bank generating US$500 million in annual revenue, this would imply that instead of revenue growing by US$50 million, it would grow by US$100 million.

These results should be interpreted with caution. First, the modelling captures the relationship between governance and revenue growth but does not reflect the costs of implementing stronger AI governance or AI systems into a business. The regression analysis was estimated using responses from the external survey panel where revenue data was requested. This may mean the results could reflect variances between the scale of banks and may not hold for all systemically important banks. The modelling does not capture potentially significant costs, such as cyber incidents, regulatory fines or remediation expenses arising from a lack of effective AI governance. Further details on the modelling are set out in Appendix B.

Taken together, the evidence suggests that governance, alongside strategic planning and effective implementation, have an important enabling role to realising the commercial benefits from AI. Banks with these factors in place are more likely to deploy AI more widely, use it more strategically and achieve stronger revenue growth. For senior leaders, the message is clear: effective governance is not only a control mechanism, but an enabler of scale, confidence and commercial value.

Figure 4.1 The AI governance dividend



Source: Deloitte analysis based on Deloitte AI Governance in Banking Survey (2026)

05

How banks can close the AI governance gap

Most banks already recognise the need for more effective AI governance. Nearly two-thirds report that their current governance arrangements do not adequately manage the way AI is being used across the organisation, and only a small minority have reached optimised maturity. That matters because the findings in this report are clear: stronger governance is associated with broader AI deployment, greater organisational confidence and stronger revenue growth. In other words, the gap in AI governance is no longer just a control issue. It is becoming a strategic and commercial gap.

The challenge for senior leaders is therefore not whether to strengthen governance, but what to focus on first and how quickly to do so while AI deployment continues. The evidence in this report suggests that the answer is not to separate innovation from control, but to build governance that is designed for scale from the outset. Banks that do this well are better able to move beyond pilots, deploy AI more widely across the enterprise and respond more effectively as risks evolve. The priorities below are intended to help banks close the gap between AI ambition and governance capability.

1

Make governance a precondition for scale, not a check at the end

Banks should treat AI governance as a key enabler for transformation rather than as a compliance layer applied after deployment decisions have already been made. Too often, governance steps are added too late, once use cases are already in development or moving into production. This creates delay, rework and inconsistent control application across the organisation.

A better approach is to design AI governance as an enterprise-wide operating capability, not a series of checks applied to individual use cases. Banks should establish common standards, decision rights and escalation mechanisms and minimum controls across all business units and geographies, with oversight calibrated to risk. This is especially important as more banks move from isolated pilots to multiple parallel deployment cycles across business units. Governance needs to be designed to accelerate safe adoption, not to react to it.



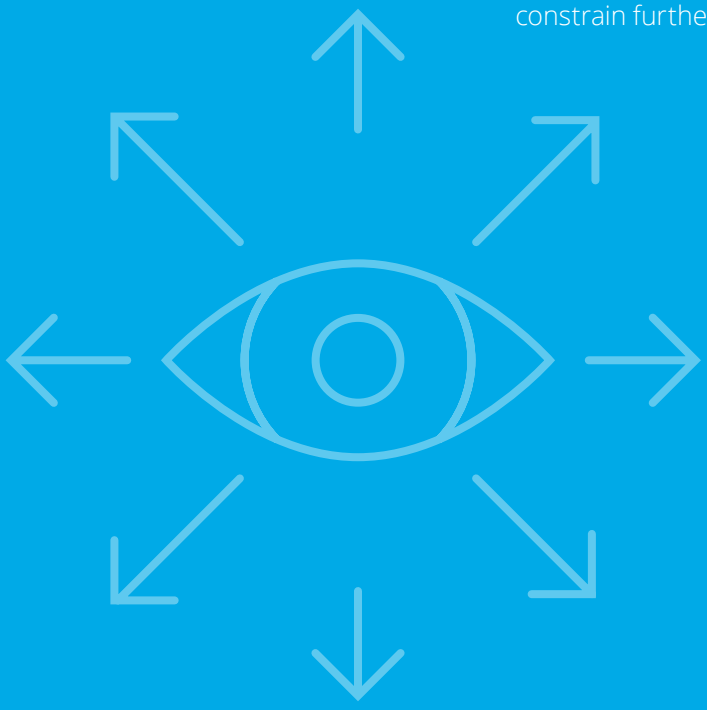
2

Put visible executive accountability at the centre

The weakest pillar in the Trustworthy AI Governance Index for many global banks is organisational structure, and one of the key differentiators is the presence of visible executive ownership. Banks should make accountability for AI risk management and AI strategy explicit at executive level, with clearly defined responsibilities across the first and second lines of defence.

This does not necessarily require a single universal operating model, but it does require unambiguous leadership. In many higher-maturity banks, AI governance is led by a Chief AI Officer, Head of AI, or senior technology executive with a deeper understanding of the technology itself and enough authority to coordinate across the business. Whatever model is chosen, senior leaders should ensure that accountability is visible, decisions are traceable and issues can be escalated quickly. AI governance cannot remain diffused across data, technology, risk and business teams without clear ownership and accountability.

Boards also have clear responsibilities to ensure AI is governed in a way that is prudent, accountable and consistent with the bank's risk appetite and legal obligations. While they do not need to oversee every use case, they should set and regularly review the overall risk appetite for AI, receive clear reporting on material exposures and challenge whether governance capability is keeping pace with deployment ambition. If they are not, boards should require remediation and be prepared to constrain further scaling until governance capability catches up.



3

Treat AI as a transverse risk, not a standalone issue

Banks with stronger governance are more likely to treat AI as a transverse risk driver embedded across existing risk domains, rather than as a narrow standalone category in their Enterprise Risk Management framework. This approach should become the norm. AI-related risks frequently cut across model risk, data risk, cyber security, legal and compliance risk, conduct risk and operational resilience. Managing them in isolation is less likely to be effective based on our survey of banks.

Banks should therefore integrate AI across enterprise risk frameworks rather than building entirely separate structures wherever possible. This includes aligning AI governance with the risk management frameworks governing model risk management, third-party risk, privacy, cyber controls and conduct oversight. Doing so helps avoid fragmented ownership and supports a more coherent view of where risks sit and how they should be managed.

Simultaneously, banks must invest in the systems and tools to allow them to effectively manage AI as a transverse risk. Whilst this is the most common approach by the more advanced banks, there are few who have visibility of AI risk exposure distinct from the transverse risk classes.



4

Encourage tech-enabled governance with 'human on the loop'

As agentic AI becomes more capable and used in the banking sector, governance models built around constant manual approval at every decision point will become impossible to sustain. For many use cases, especially those involving autonomous or semi-autonomous systems, banks will need to shift from "human in the loop" to "human on the loop" - where people oversee systems, define boundaries, monitor performance and intervene when needed, rather than manually reviewing every output.

This requires more technology-enabled governance steps and controls. Banks should invest in automated monitoring, exception reporting, usage logging, access controls and real-time alerts that can keep pace with AI systems operating at speed and scale. Human judgement remains essential, but it must be directed to the points of greatest risk rather than used as a blanket control that slows decision-making without improving oversight.

This shift will be especially important for agentic AI, where governance maturity currently lags traditional and generative AI. Banks should not assume that existing controls will be sufficient for more autonomous systems. They will need new guardrails, stronger monitoring and clearer intervention triggers.



5

Create a risk-reward culture around AI

One of the barriers to effective AI governance is organisational separation: innovation teams are incentivised to move quickly, while risk and compliance teams are asked to slow things down. In practice, this can create friction, encourage workarounds and slow decision-making. Banks should build a risk-reward culture in which business, innovation, risk and compliance functions share accountability for scaling AI safely and effectively.

A core principle of this culture should be that ownership of AI risk sits first with the business function that stands to benefit from the use case. If a business line expects the commercial, operational or customer benefit from deploying AI, it should also own the associated risks, controls and outcomes, rather than treating risk management as something transferred to second-line teams. This helps align incentives, improve decision quality and ensure AI use cases are pursued only where the business is prepared to manage the risks that come with the reward. Establishing and embedding a stronger risk-reward culture helps avoid two common failures identified in this report: excessive caution that delays deployment, and rushed adoption under competitive pressure without adequate safeguards. The goal is balanced execution – moving at pace, but with discipline.



Appendices



Appendix A
Global bank survey

Data was collected using two purpose-built survey instruments designed to capture the perspective across the global banks, securing a total of 135 respondents from across 16 countries. The surveys aimed to assess the maturity level of AI governance structures and understand the benefits of good AI governance.

In 2026, Deloitte's Trustworthy AI team surveyed 24 leaders from G-SIBs and D-SIBs across 14 countries to assess the maturity level of AI governance structures. Following this, a survey panel provider was used to survey an additional 111 employees of G-SIBs, D-SIBs and other large banks in May 2026. Respondents were specifically targeted to be in senior roles like Chief Technology Officers, Chief Information Officers and Chief Data Officers in the banking industry. This secondary sample asked questions related to AI governance, AI usage and operational information. Chart A.1 presents the role of respondents, and Table A.1 presents the country of respondents' headquarters.

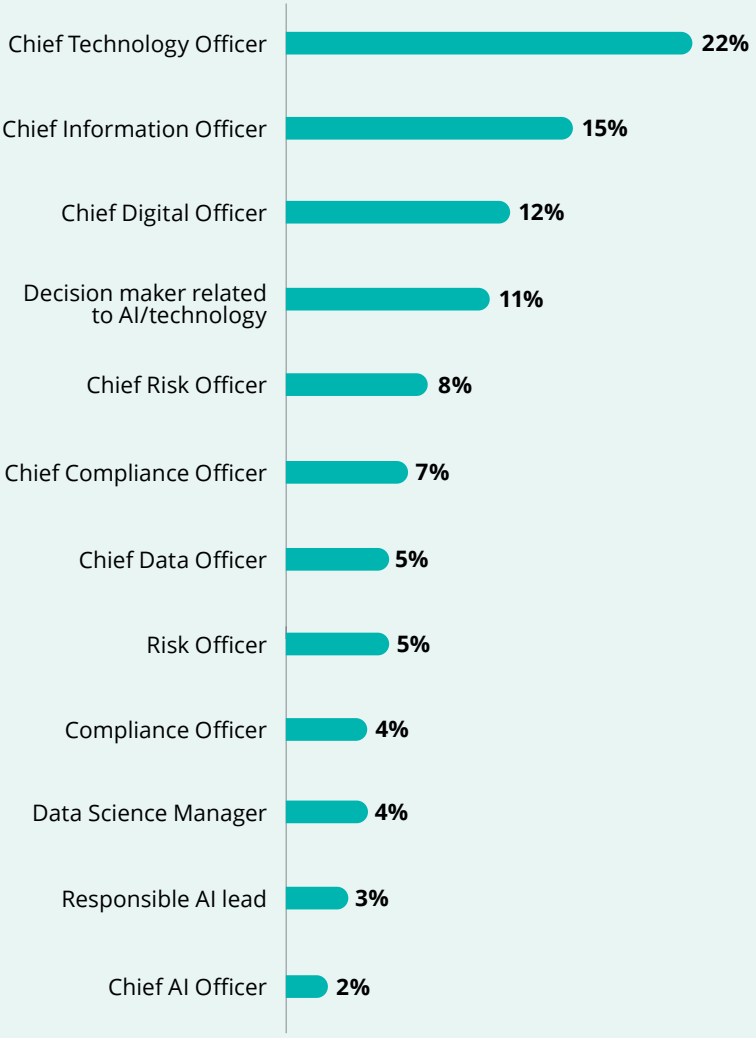
Table A.1 Location of survey respondents

Location of headquarters	Number of respondents
United States	24
Canada	1
Europe	93
United Kingdom	23
Ireland	1
Spain	1
France	1
The Netherlands	1
Germany	18
Denmark	22
Sweden	24
Norway	1
Finland	1
Australia	9
China	4
Japan	3
Singapore	1
Total	**135**

Source: Deloitte global bank survey (2026) and Deloitte Trustworthy AI survey (2026)



Chart A.1 Role of survey respondents (Deloitte Survey, n=111)



Source: Deloitte Global Bank Survey (2026)

Appendix B
Index development and econometric modelling

This appendix sets out how the AI Governance Index was constructed and how the link between AI governance and bank financial performance was estimated. The methodology follows three steps: building the index from the survey responses, estimating a regression of financial performance on governance, and interpreting the results for revenue and revenue growth.

B.1
Building the AI Governance Index

The AI Governance Index scores each bank on a 0 to 100 scale built from the five pillars of AI governance: principles and policy, organisational structure, people and skills, procedures and controls, and monitoring, reporting and evaluation. The index is built in four steps. First, each underlying survey question is scored from 1 (early-stage practice) to 5 (fully embedded practice). Second, the scores for a bank are averaged across all valid responses to give an overall mean on the 1 to 5 scale. Third, the mean is rescaled to a 0 to 100 index using the formula (mean minus 1) divided by 4, multiplied by 100. Fourth, banks are grouped into four tiers of governance maturity based on the index score. The index covers all 135 banks in the survey.

Index = (mean survey score minus 1) / 4 × 100

Each bank is assigned to one of four maturity tiers based on its index score. The ad hoc tier covers scores of 25 or below and contains 15 banks (14%) of the sample. The rudimentary tier covers scores from 26 to 50 and contains 14 banks (13%). The established tier covers scores from 51 to 80 and contains 59 banks (54%). The optimised tier covers scores above 80 and contains 22 banks (20%). The mean overall score is 61, the median is 65, and the standard deviation is 23. The distribution is broadly bell-shaped, with most banks sitting in the established tier where AI governance is real but uneven.

Figure B.1 AI Governance Index: pillar composition and average scores

Pillar (weight)	What it captures	Average score (0–100)
Principles and policy (20%)	Whether AI principles and policies are documented, approved and applied across the bank.	58
Organisational structure (20%)	Whether clear roles, committees and accountability for AI governance are in place.	50
People and skills (20%)	Whether staff are trained in AI risk and the bank has the skills to govern AI.	56
Procedures and controls (20%)	Whether controls operate across the AI lifecycle, including registers and approvals.	61
Monitoring, reporting and evaluation (20%)	Whether AI use is monitored, reported to leadership and evaluated against outcomes.	61
Overall index	Equal-weighted average of the five pillars	57

Source: Deloitte global bank survey (2026)

B.2
Modelling the link to financial performance

To test whether stronger AI governance is associated with better financial results, revenue growth was modelled against the AI Governance Index for the banks that reported financial figures. Reported revenue was first converted to a common currency, United States dollars, using exchange rates of AUD 0.66, GBP 1.27 and EUR 1.08. The reported figures were then cleaned: where a bank appeared to have stated revenue in a different unit to the one requested, the figure was rescaled; a small number of implausibly low values were set to a minimum threshold; and the top and bottom 7.5% of the revenue distribution were excluded to limit the influence of outliers. This left 87 banks in the estimating sample. Revenue growth was measured year on year, from the latest reported year to the following expected year.

The analysis used ordinary least squares regression with standard errors corrected for differences in variance across banks. The AI Governance Index enters as a continuous 0 to 100 score. The model controls for the extent of AI use (the number of business areas with fully implemented AI), headquarters country, bank size (the natural logarithm of employee count) and the number of years the bank has been in operation. Revenue growth is modelled in log-difference form, as ln(1 plus the growth rate), because this is less sensitive to extreme values than the raw rate.

The variables in each regression are defined as follows.

Outcome variable

- Revenue growth (log) is the log difference in revenue from FY2024–25 to FY2025–26, calculated as ln(1 plus the growth rate).

Independent Variable

- Governance variable is the Trustworthy AI Governance Index maturity score, which can range from 0 to 100.

Control Variables

- AI use is the number of business areas (0 to 10) in which the bank has fully implemented AI.
- Country is a set of indicators for the headquarters market, with Australia as the omitted reference category.
- ln(Employees) is the natural logarithm of employee count, controlling for bank size.
- Years in operation is the number of years the bank has been in operation.

The estimating equation for the tier specification is:

$$ln(1 + \text{Revenue growth}) = \beta_0 + \beta_1 \text{AI Governance Index} + \beta_2 \text{AI use} + \beta_3 \text{Country} + \beta_4 ln(\text{Employees}) + \beta_5 \text{Years in operation} + \varepsilon$$

Figure B.2 Revenue growth (log), continuous AI Governance Index

Variable	Coefficient	Std. error	p-value
Constant	1.11	(2.27)	0.63
AI Governance Index (0–100)	0.010*	(0.005)	0.04
United States	-0.65	(0.82)	0.42
United Kingdom	-0.89	(0.80)	0.27
Denmark	-0.74	(0.81)	0.36
Germany	-0.81	(0.78)	0.31
Sweden	-0.90	(0.79)	0.26
ln(Employees)	-0.08	(0.27)	0.77
AI use (business areas fully implemented)	-0.00	(0.04)	0.99
Years in operation	-0.00	(0.01)	0.93
Country fixed effects	Yes		
N	87		

Variable	Coefficient	Std. error	p-value
R^2	0.14		
Adjusted R^2	0.04		
Constant	1.11	(2.27)	0.63

Significance: * p < 0.05, ** p < 0.01, *** p < 0.001. Standard errors in parentheses are corrected for differences in variance (HC1). Country controls are five indicators (United States, United Kingdom, Denmark, Germany, Sweden) with Australia as the omitted reference category.

Some of the evidence suggests a positive association between governance and revenue growth. Each additional point on the 0 to 100 Index is associated with faster revenue growth, and this relationship is statistically strong (p = 0.04). Read across a 10-point difference in the Index, the model implies revenue growth around 10 percentage points faster, holding AI use, country, size and age constant. For the average bank in our survey, which recorded revenue growth of about 10% over the year, this points to growth nearer 20% had its governance score been 10 points higher. This should be read as the financial profile of higher-governance banks rather than a forecast of the return any individual bank would receive from a governance programme.

To make the scale tangible, consider a bank with annual revenue of around US$500 million growing at the survey average. Instead of revenue rising by about US$50 million in a year, the model would associate a governance score 10 points higher with an increase nearer US$100 million. This figure is purely illustrative. It describes a statistical association observed across a small sample of mostly smaller banks, not a return that a specific bank could expect to make by raising its score.

B.3
Modelling the link to AI implementation

To test whether stronger AI governance is associated with wider AI implementation, the count of business areas with fully implemented AI was modelled against the AI Governance Index for the 111 banks in the Dynata sample. The outcome ranges from zero to ten, across the ten business areas respondents rated (operations; marketing and sales; finance; human resources; customer service; research and development; information technology; management and administration; legal and compliance; and other). The mean is 4.6 fully implemented areas per bank.

Figure B.3 Fully implemented business areas, continuous index

Variable	Coefficient	Std. error	p-value
Constant	-0.69	(5.89)	0.91
AI Governance Index (0-100)	0.065***	(0.01)	<0.001
United States	1.71	(0.92)	0.06
United Kingdom	1.22	(0.75)	0.10
Denmark	-0.37	(0.69)	0.59
Germany	1.625*	(0.71)	0.02
Sweden	1.18	(0.85)	0.17
ln(Employees)	-0.28	(0.67)	0.68
AI share (%)	0.047***	(0.01)	<0.001
Country fixed effects	Yes		
N	111		
R-squared	0.53		

Significance: * p < 0.05, ** p < 0.01, *** p < 0.001. Standard errors in parentheses are corrected for differences in variance (HC1). Country controls are five indicators (United States, United Kingdom, Denmark, Germany, Sweden) with Australia as the omitted reference category.

Figure B.4 Fully implemented business areas, continuous AI Governance Index

Variable	Coefficient	Std. error	p-value
Constant	-0.69	(5.89)	0.91
AI Governance Index (0–100)	0.065***	(0.01)	<0.001
United States	1.71	(0.92)	0.06
United Kingdom	1.22	(0.75)	0.10
Denmark	-0.37	(0.69)	0.59
Germany	1.625*	(0.71)	0.02
Sweden	1.18	(0.85)	0.17
ln(Employees)	-0.28	(0.67)	0.68
AI share (%)	0.047***	(0.01)	<0.001
Country fixed effects	Yes		
N	111		
R^2	0.53		
Constant	-0.69	(5.89)	0.91
AI Governance Index (0–100)	0.065***	(0.01)	<0.001

Significance: * p < 0.05, ** p < 0.01, *** p < 0.001. Standard errors in parentheses are corrected for differences in variance (HC1). Country controls are five indicators with Australia as the reference.

In this specification a 10-point increase in the Index is associated with around 0.65 more business areas with fully implemented AI, holding bank size, country and AI use constant, and the relationship is statistically strong ($p < 0.001$). The association with implementation is therefore clearer than the association with revenue growth.

B.4
Limitations

The findings in this appendix should be read with considerable caution. Several features of the data and method mean the results describe a result occurring on average, not a guaranteed return that any individual bank could rely on. Some of the limitations of the study include:

- The regressions estimate a statistical relationship between AI governance and financial outcomes, conditional on the observed controls. They do not establish that lifting governance causes higher revenue growth for all banks. Influence may also run both ways: larger, better-resourced banks may find it easier to invest in governance, and well-run banks may invest in both governance and other, unobserved drivers of performance.

- We cannot be certain this holds for very large banks. The estimating sample is made up almost entirely of smaller institutions. The largest bank in the cleaned sample reports under US$4 billion in annual revenue, and none reach US$5 billion. The relationship is therefore estimated well below the scale of the world's largest banks. Some of the evidence suggests the association may extend upwards, but we cannot be certain it would play out among G-SIBs, whose revenues run into the tens of billions of dollars and whose growth is shaped by factors this survey does not observe. Any dollar figure applied to such banks is an out-of-sample extrapolation and should be treated as illustrative only.

- Revenue and growth are self-reported by respondents, not drawn from audited accounts, which introduces measurement error, particularly at the tails of the distribution. The reported revenue of the lowest-governance banks is very low and is likely to reflect incomplete reporting. Cleaning the data reduces, but does not remove, this error.

- The survey provides a single observation per bank at one point in time. It cannot observe a bank moving between maturity levels and track its outcomes, which a repeated survey of the same banks across several years would allow.

- The models control for headquarters country, bank size, AI use and years in operation. They do not control for business mix, regulatory environment or client base, which may differ across banks in ways the survey does not capture.

- The four maturity levels (ad hoc at 25 or below, rudimentary from 26 to 50, established from 51 to 80, and optimised above 80) are analytical choices. The continuous Index results reported here do not depend on these cut-offs.

Appendix C
Regulatory developments in select markets

The table below provides some regulatory developments and approaches to AI for banks or financial services for select markets.

Market	AI regulatory environment and developments
United States	AI regulation in banking is mostly through 'rules through existing' frameworks, with federal banking agencies, Treasury, state regulators and sector-specific laws shaping expectations. Recently, the Federal Reserve, Office of the Comptroller of the Currency (OCC), and Federal Deposit Insurance Corporation (FDIC) issued a revised model risk management guidance, replacing SR 11-7 and emphasising a risk-based approach tailored to banks' model risk profile.[37] Treasury also released a Financial Services AI Risk Management Framework and AI Lexicon in February 2026.[38]
United Kingdom	The UK continues to take a technology-neutral, principles-based approach to AI regulation in financial services. The Prudential Regulation Authority's SS1/23 sets expectations for banks' model risk management across five principles: Model identification and model risk classification, governance, model development, implementation and use, independent model validation and model risk mitigants.[39] It explicitly notes AI and Machine learning risks where they arise in modelling.
European Union	The EU AI Act (2024) provides risk-based AI regulation that is applicable to all industries, including banking. It entered in force on 1 August 2024, with most obligations applying from 2 August 2026 onwards. It imposes compliance obligations on 'high-risk' use cases, such as the use of AI to assess creditworthiness or establish credit scores of natural persons, triggering requirements around risk management, data governance, transparency, human oversight and documentation.[40] Under Article 99 of the Act, non-compliance with certain prohibited or non-compliant AI practices can result in fines of up to €35 million or 7% of annual global turnover, whichever is higher. Other violations, including the provision of incorrect or misleading information, can attract fines of up to €7.5 million or 1% of annual global turnover.[41]
United Arab Emirates	The central bank of the UAE issued a Guidance Note on consumer protection and responsible adoption of AI and Machine Learning by Licensed Financial Institutions.[42] The framework sets out expectations across governance, accountability, fairness / non-discrimination, transparency and explainability, data quality, privacy and security, continuous monitoring, human oversight, outsourcing and third-party risk. The guidance is framed as broad and flexible rather than static.
Japan	Japan has adopted a promotion-oriented, soft-law and coordination focused approach to AI regulation. The Act on Promotion of Research, Development and Utilisation of AI-related Technology establishes principles, an AI Basic Plan and an AI Strategic Headquarters, aimed at promoting AI research and development, and use while addressing risks.[43]
China	China is characterised by a layered, sector specific AI regulatory framework, with a combination of horizontal legislation, including the Cybersecurity Law, Data Security Law, and personal Information Protection Law, with AI-specific regulations such as the Provisions on the Administration of Algorithmic Recommendations (2022), the Deep Synthesis Provisions (2023), and the Interim Measures of the Management of generative AI Services (2023). These measures are complemented by financial-sector oversight from the People's Bank of China (PBOC) and National Financial Regulatory Administration (NFRA), while the Cybersecurity Administration of China (CAC) serves as the lead authority for AI governance. These frameworks place strong emphasis on data governance, cybersecurity, algorithm governance, risk management and accountability.

Market	AI regulatory environment and developments
Hong Kong SAR	Hong Kong SAR has taken a risk-based and technology-neutral supervisory approach, with targeted guidance for financial institutions. The Hong Kong Monetary Authority's (HKMA) August 2024 circular on customer-facing GenAI builds on earlier Big Data Analytics and Artificial Intelligence (BDAI) principles and focuses on governance/accountability, fairness, transparency/disclosure and data privacy/protection.[45] The Security & Futures Commission's November 2024 circular warns that AI language models can amplify risks and treats use cases such as investment recommendations, advice or research as higher-risk activities requiring stronger controls. The HKMA also launched a new Generative Artificial Intelligence ("GenAI") sandbox++ initiative in March 2026, supporting banks to safely develop and test AI use cases, aiming to accelerate responsible AI adoption across the financial sector.[46]
Singapore	Singapore has historically relied on principles-based and industry-collaborative AI governance, including Monetary Authority of Singapore's (MAS) Fairness, Ethics, Accountability and Transparency (FEAT) principles and Veritas initiative.[47] More recently, MAS has moved towards more structured supervisory expectations. MAS issued proposed Guidelines on AI Risk Management in November 2025 for financial institutions, covering governance, AI inventories, lifecycle controls, human oversight, fairness, transparency and third-party risk.
Australia	The Australian Prudential Regulation Authority (APRA) has moved from general principles to clearer AI-specific supervisory expectations. In its 30 April 2026 letter, APRA called for a step-change in how banks, insurers and superannuation trustees manage AI-related risks, warning that governance, risk management, assurance and operational resilience are not keeping pace with AI adoption.[48] APRA expects boards to improve AI literacy and oversight, and entities to strengthen governance, cyber, supplier risk and assurance practices.

Endnotes

1 Bank for International Settlements (2026). Banks from All reporting countries - Consolidated total assets. https://data.bis.org/topics/CBS/BIS%2CWS_CBS_PUB%2C1.0/Q.S.5A.4O.F.F.A.A.TO1.A.5J?view=observations.

2 IMF (2026) World Economic Outlook, https://data.imf.org/en/datasets/IMF.RES:WEO.

3 McKinsey (2026) Global banking annual review 2026: Precision with speed, https://www.mckinsey.com/industries/financial-services/our-insights/global-banking-annual-review#/.

4 International Monetary Fund (2026) World Economic Outlook, https://www.imf.org/en/publications/weo

5 Moody's (2025) Outlooks 2026: Global Banks https://www.moodys.com/web/en/us/insights/credit-risk/outlooks/banking-2026.html

6 Vultr (2025). Unlocking the Power of AI in Financial Services. https://www.vultr.com/marketing-sales-files/ai-maturity-report-financial-services-2025.pdf

7 Bank of America (2024), BofA's Erica Surpasses 2 Billion Interactions, Helping 42 Million Clients Since Launch. https://newsroom.bankofamerica.com/content/newsroom/press-releases/2024/04/bofa-s-erica-surpasses-2-billion-interactions--helping-42-millio.html

8 Deloitte (2026). State of AI in the Enterprise. https://www.deloitte.com/content/dam/assets-zone3/us/en/docs/services/consulting/2026/state-of-ai-2026.pdf

9 J.P.Morgan (2026) Four takeaways on AI's next phase in payments https://www.jpmorgan.com/payments/newsroom/agentic-ai-payments-takeaways-2026

10 Anthropic (2026) Project glasswing: securing critical software for the AI era, https://www.anthropic.com/glasswing

11 OECD AI Observatory (2026). AIM: AI Incidents and Hazards Monitor https://oecd.ai/en/incidents?search_terms=%5B%5D&and_condition=false&from_date=1900-05-27&to_date=2026-05-27&properties_config=%7B%22principles%22:%5B%5D,%22industries%22:%5B%5D,%22harm_types%22:%5B%5D,%22harm_levels%22:%5B%5D,%22harmed_entities%22:%5B%5D,%22business_functions%22:%5B%5D,%22ai_tasks%22:%5B%5D,%22autonomy_levels%22:%5B%5D,%22languages%22:%5B%5D%7D&order_by=date&num_results=20

11 IBM (2025), Cost of a Data Breach Report 2025 The AI Oversight Gap https://www.ibm.com/downloads/documents/us-en/131cf87b20b31c91

12 IBM (2025), Cost of a Data Breach Report 2025 The AI Oversight Gap https://www.ibm.com/downloads/documents/us-en/131cf87b20b31c91

13 APRA (2026) APRA Letter to Industry on Artificial Intelligence, https://www.apra.gov.au/apra-letter-to-industry-on-artificial-intelligence-ai

14 Financial Stability Board (2026) Sound Practices for Responsible Adoption of Artificial Intelligence (AI) https://www.fsb.org/uploads/P100626.pdf

15 IOSCO (2026) IOSCO publishes AI Supervisory Toolkit https://www.iosco.org/news/pdf/IOSCONEWS796.pdf

16 Federal Reserve (2026), SR 26-2: Revised Guidance on Model Risk Management https://www.federalreserve.gov/supervisionreg/srletters/SR2602.htm

17 U.S. Department of the Treasury (2026), Treasury Releases Two New Resources to Guide AI Use in the Financial Sector https://home.treasury.gov/news/press-releases/sb0401

18 Bank of England PRA (2026) Model risk management principles for banks https://www.bankofengland.co.uk/prudential-regulation/publication/2023/may/model-risk-management-principles-for-banks-ss

19 European Banking Authority (2025), AI Act: implications for the EU banking and payments sector https://www.eba.europa.eu/sites/default/files/2025-11/d8b999ce-a1d9-4964-9606-971bbc2aaf89/AI%2520Act%2520implications%2520for%2520the%2520EU%2520banking%2520sector.pdf

20 Central Bank of the U.A.E. (2026) Guidance Note on the Consumer Protection and Responsible Adoption and Use of Artificial Intelligence and Machine Learning by Licensed Financial Institutions in the U.A.E. https://rulebook.centralbank.ae/en/rulebook/guidance-note-consumer-protection-and-responsible-adoption-and-use-artificial-intelligence

21 Public Relations Office Government of Japan (2025), Act on Promotion of Research and Development, and Utilization of Artificial Intelligence-related Technology Now in Full Effect https://www.gov-online.go.jp/hlj/en/november_2025/november_2025-08.html

22 The Asian Banker (2026), Can China's banks turn AI governance from a board mandate into a competitive advantage? https://www.theasianbanker.com/updates-and-articles/can-china-s-banks-turn-ai-governance-from-a-board-mandate-into-a-competitive-advantage

23 Hong Kong Monetary Authority (2026) Regulators launch GenA.I. Sandbox++ to foster A.I. innovation across financial services https://www.hkma.gov.hk/eng/news-and-media/press-releases/2026/03/20260305-3/

24 Monetary Authority of Singapore (2026), Artificial Intelligence https://www.mas.gov.sg/development/fintech/artificial-intelligence

25 APRA (2026) APRA Letter to Industry on Artificial Intelligence, https://www.apra.gov.au/apra-letter-to-industry-on-artificial-intelligence-ai

26 Harvard Law School Forum on Corporate Governance (2025) AI Risk Disclosures in the S&P 500: Reputation, Cybersecurity, and Regulation https://corpgov.law.harvard.edu/2025/10/15/ai-risk-disclosures-in-the-sp-500-reputation-cybersecurity-and-regulation/

27 Salesforce (2025), The Connected Financial Services Report. https://www.salesforce.com/content/dam/web/en_us/www/documents/industries/financial-services/Connected-FINS_Report_Final.pdf

28 McKinsey & Company, (2023), Capturing the full value of generative AI in banking,https://www.mckinsey.com/industries/financial-services/our-insights/capturing-the-full-value-of-generative-ai-in-banking

29 CIO Dive (2025), AI Adoption will trim banking industry costs by up to 20%, https://www.ciodive.com/news/ai-trim-banking-industry-costs/804341/

30 MIT NANDA, (2025), The GenAI Divide: State of AI in Business 2025 (reported via Fortune), https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/

31 Cambridge Centre for Alternative Finance, (2026), 2026 Global AI in Financial Services Report, https://www.jbs.cam.ac.uk/faculty-research/centres/alternative-finance/publications/2026-global-ai-in-financial-services-report/

32 Salesforce (2025), The Connected Financial Services Report. https://www.salesforce.com/content/dam/web/en_us/www/documents/industries/financial-services/Connected-FINS_Report_Final.pdf

33 McKinsey & Company, (2023), Capturing the full value of generative AI in banking, https://www.mckinsey.com/industries/financial-services/our-insights/capturing-the-full-value-of-generative-ai-in-banking

34 CIO Dive (2025), AI Adoption will trim banking industry costs by up to 20%, https://www.ciodive.com/news/ai-trim-banking-industry-costs/804341/

35 MIT NANDA, (2025), The GenAI Divide: State of AI in Business 2025 (reported via Fortune), https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/

36 Cambridge Centre for Alternative Finance, (2026), 2026 Global AI in Financial Services Report, https://www.jbs.cam.ac.uk/faculty-research/centres/alternative-finance/publications/2026-global-ai-in-financial-services-report/

37 Federal Reserve (2026), SR 26-2: Revised Guidance on Model Risk Management https://www.federalreserve.gov/supervisionreg/srletters/SR2602.htm

38 U.S. Department of the Treasury (2026), Treasury Releases Two New Resources to Guide AI Use in the Financial Sector https://home.treasury.gov/news/press-releases/sb0401

39 Bank of England PRA (2026) Model risk management principles for banks https://www.bankofengland.co.uk/prudential-regulation/publication/2023/may/model-risk-management-principles-for-banks-ss

40 European Banking Authority (2025), AI Act: implications for the EU banking and payments sector https://www.eba.europa.eu/sites/default/files/2025-11/d8b999ce-a1d9-4964-9606-971bbc2aaf89/AI%2520Act%2520implications%2520for%2520the%2520EU%2520banking%2520sector.pdf

41 EU Artificial Intelligence Act (2025). Article 99: Penalties https://artificialintelligenceact.eu/article/99/

42 Central Bank of the U.A.E. (2026) Guidance Note on the Consumer Protection and Responsible Adoption and Use of Artificial Intelligence and Machine Learning by Licensed Financial Institutions in the U.A.E. https://rulebook.centralbank.ae/en/rulebook/guidance-note-consumer-protection-and-responsible-adoption-and-use-artificial-intelligence

43 Public Relations Office Government of Japan (2025), Act on Promotion of Research and Development, and Utilization of Artificial Intelligence-related Technology Now in Full Effect https://www.gov-online.go.jp/hlj/en/november_2025/november_2025-08.html

44 The Asian Banker (2026), Can China's banks turn AI governance from a board mandate into a competitive advantage? https://www.theasianbanker.com/updates-and-articles/can-china-s-banks-turn-ai-governance-from-a-board-mandate-into-a-competitive-advantage

45 Hong Kong Monetary Authority (2024), Consumer Protection in respect of Use of Generative Artificial Intelligence https://brdr.hkma.gov.hk/eng/doc-ldg/docId/getPdf/20241107-1-EN/20241107-1-EN.pdf

46 Hong Kong Monetary Authority (2026) Regulators launch GenA.I. Sandbox++ to foster A.I. innovation across financial services https://www.hkma.gov.hk/eng/news-and-media/press-releases/2026/03/20260305-3/

47 Monetary Authority of Singapore (2026), Artificial Intelligence https://www.mas.gov.sg/development/fintech/artificial-intelligence

48 APRA (2026) APRA Letter to Industry on Artificial Intelligence, https://www.apra.gov.au/apra-letter-to-industry-on-artificial-intelligence-ai

Contact us

Key contacts



Dr. Elea Wurth
Lead Partner, Trustworthy AI
Deloitte Asia Pacific & Australia
ewurth@deloitte.com.au



Thomas Clifford
Lead Partner, Trustworthy AI
Deloitte EMEA
thclifford@deloitte.dk



Clifford Goss
Lead Partner, AI Risk Management and Trustworthy AI
Deloitte US
cgoss@deloitte.com



John O'Mahony
Partner, Deloitte Access Economics
Deloitte Australia
joomahony@deloitte.com.au



Neil Tomlinson
Global Banking & Capital Markets Leader
Deloitte Global
ntomlinson@deloitte.co.uk



Stuart Johnston
Asia Pacific Financial Services Leader
Deloitte Asia Pacific
stujohnston@deloitte.com.au



David Wai Kit Wu
Asia Pacific Financial Services Market Activation Leader & Hong Kong Financial Services Leader
Deloitte Asia Pacific & China
davidwwu@deloitte.com.hk



Julian Leake
EMEA Financial Services Leader
Deloitte EMEA
jileake@deloitte.co.uk

Contributors



Mariette Van Niekerk
Director, Trustworthy AI
Deloitte New Zealand
mvanniekerk@deloitte.co.nz



Nick Hull
Director, Deloitte Access Economics
Deloitte Australia
nhull@deloitte.com.au



Lester Gunnion
Manager, Deloitte Access Economics
Deloitte Australia
legunnion@deloitte.com.au



Elvis Ferreira
Analyst, Deloitte Access Economics
Deloitte Australia
elvferreira@deloitte.com.au



Reid Quekett
Analyst, Deloitte Access Economics
Deloitte Australia
rquekett@deloitte.com.au



Adelle Thomas
Analyst, Deloitte Access Economics
Deloitte Australia
adthomas@deloitte.com.au



Alex MacMillian
Senior consultant, Trustworthy AI
Deloitte New Zealand
almacmillan@deloitte.co.nz



Charlotte Stevens
Senior consultant, Trustworthy AI
Deloitte Australia
charstevens@deloitte.com.au

Deloitte.

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited ("DTTL"), its global network of member firms, and their related entities (collectively, the "Deloitte organization"). DTTL (also referred to as "Deloitte Global") and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/about to learn more.

Deloitte Asia Pacific Limited is a company limited by guarantee and a member firm of DTTL. Members of Deloitte Asia Pacific Limited and their related entities, each of which is a separate and independent legal entity, provide services from more than 100 cities across the region, including Auckland, Bangkok, Beijing, Bengaluru, Hanoi, Hong Kong, Jakarta, Kuala Lumpur, Manila, Melbourne, Mumbai, New Delhi, Osaka, Seoul, Shanghai, Singapore, Sydney, Taipei and Tokyo.

This communication contains general information only, and none of DTTL, its global network of member firms or their related entities is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser.

No representations, warranties or undertakings (express or implied) are given as to the accuracy or completeness of the information in this communication, and none of DTTL, its member firms, related entities, employees or agents shall be liable or responsible for any loss or damage whatsoever arising directly or indirectly in connection with any person relying on this communication.

© 2026 Deloitte Asia Pacific Services Limited