Cambridge
Centre
for Alternative
Finance

UNIVERSITY OF
CAMBRIDGE
Judge Business School

In partnership with

BIS

International Monetary Fund

World Economic Forum

IDB

CGAP

صندوق النقد العربي
ARAB MONETARY FUND

The 2026 Global AI in Financial Services Report

Adoption, impact and risks

April 2026

In collaboration with

Financial Innovation for Impact

With the support of

Foreign, Commonwealth & Development Office

Suggested Citation

The 2026 Global AI in Financial Services Report: Adoption, impact and risks,
Cambridge Centre for Alternative Finance, University of Cambridge

Contents

Foreword

Artificial intelligence is no longer a peripheral experiment in financial services; it is rapidly becoming a central, structural component of the global financial ecosystem. As the industry accelerates from classical machine learning towards generative and increasingly autonomous agentic AI systems, it faces a critical inflection point. This transition promises to enhance operational efficiency, unlock new value, and reshape market dynamics. Yet, as in earlier phases of financial innovation, this rapid pace of technological change risks outpacing organisational readiness and regulatory capacity. The sector currently lacks a consistent, global evidence base to navigate this landscape, leaving financial institutions and policymakers struggling to separate tangible value from technological promise, and working to manage emerging systemic risks such as data privacy, model hallucinations, critical third-party dependencies and rapidly evolving cyber risks.

To help address this empirical blind spot, the Cambridge Centre for Alternative Finance (CCAF) at the University of Cambridge Judge Business School, in collaboration with Financial Innovation for Impact (Fii) and a consortium of global institutions including; the Bank for International Settlements (BIS), the International Monetary Fund (IMF), the World Economic Forum (WEF), the World Bank Group, the Inter-American Development Bank (IDB), CGAP, and the Arab Monetary Fund (AMF), has produced the 2026 Global AI in Financial Services Report.

Building upon our inaugural 2020 CCAF-WEF AI Report, this research captures the intersecting perspectives of 628 financial institutions, AI vendors, and regulatory authorities operating across 151 jurisdictions. Rather than simply cataloguing adoption rates, this report provides a broad perspective on the complex realities of AI integration. The analysis reveals an industry in mid-transition, marked by a significant execution gap between widespread early-stage adoption and true strategic transformation. It highlights profound structural divides driven by persistent constraints in data quality, legacy infrastructure, and talent access. Furthermore, the study underscores the urgent need for aligned governance, as the deployment of advanced AI systems by the private sector currently outpaces the supervisory frameworks and technical capacities required to oversee them.

We hope the insights presented in this landmark study provide a robust foundation for forward-looking dialogue among financial institutions, technology providers, and policymakers. By moving beyond isolated efficiency plays to measurable value creation and responsible governance, the financial sector can realise AI's full transformative potential while safeguarding market stability and protecting consumers.

Finally, we extend our gratitude to Keith Barnes and the UK Foreign, Commonwealth and Development Office (FCDO) for their support, and to all survey participants, partners, and collaborators whose perspectives made this global research possible. We would also like thank Jon Frost, Bénédicte Nolens, Drew Propson, Harish Natarajan, Gian Boeddu, Miguel Segoviano, Diego Herrera, Henrique Chitman, Ana María Zárate Moreno, Haocong Ren, Eric Duflos, Camila Adriana Quevedo Vega, Nouran Youssef and Karim Mouaffak for their partnership.

Bryan Zhang
Co-Founder and Executive Director, Cambridge Centre for Alternative Finance Judge Business School, University of Cambridge

Kieran Garvey
AI Lead, Financial Innovation for Impact, Strategic AI Advisor, Cambridge Centre for Alternative Finance Judge Business School, University of Cambridge.

Research Team and Acknowledgements

Research Team

Kieran Garvey (CCAF/Fii), Bryan Zhang (CCAF), Innes Roberts (CCAF), Farah Abdou (CCAF), Hamidu Midlah Barrie (CCAF), Douglas Arner (CCAF), Benedicte Nolens (BISIH), Jon Frost (BIS), Miguel Segoviano (IMF), Drew Propson (World Economic Forum), Felipe Ferri de Camargo Paes (CCAF), Krishnamurthy Suresh (Fii), Alexander Apostolides (CCAF), Waleed Nadeem (CCAF), Ashis Mittal (CCAF), Lucas Elletore (CCAF), Marcos Proença (CCAF), Lukas Ryll (CCAF), Eric Duflos (CGAP), Haocong Ren (CGAP), Diego Herrera (IDB), Keith Bear (CCAF), Yue Wu (CCAF), Stanley Mutinda (CCAF), Senanu Dekpo-Adza (CCAF), Shashank Dubey (CCAF), Andres Lehtmets (CCAF), Patricia Evite (UNINA), Caroline Malcolm (Fii), Hugo Coelho (Fii), Nadia Hazeveld (CCAF, Fii), Richard Kithuka (CCAF), Nolwazi Hlophe (FSCA), Roman Rampa, Jack Lee (BISIH), Vatsala Shreeti (BIS), Kumar Rishabh (BIS), Henrique Chitman (IDB), Ana María Zárate Moreno (IDB), Camila Quevedo Vega (CGAP), Arisha Salman (CGAP), Nouran Youssef (AMF), Karim Mouaffak (AMF) and Bretchen Hoskins (CMF).

Acknowledgments

On behalf of the Cambridge Centre for Alternative Finance (CCAF) at the University of Cambridge Judge Business School, the research team would like to thank Keith Barnes and the UK Foreign Commonwealth and Development Office (FCDO) for making this report possible, and for their continued support of the CCAF and Financial Innovation for Impact (Fii).

We are grateful for the support and contributions of the wider research team for their dedication and support; Jelena Novakovic, Denny Justin, Laone Otukile, Sarah Ombija, Waleed Samarah, Shubhi Agrawal, Vidhi Arora, Bridget Menyeh, Moriti Omondi, Aaryani Dogra, Matt Jones, Leigh Shlomovic, Sahan Bulathwela and Derek Snow.

We are grateful to the CCAF and Fii teams for their support; Kate Belger, Hunter Sims, Pavle Avramovic, Nick Clark, Ana Odorovic, Philip Rowan, David Allen, Clare Bannister, Tim Brown, Yvona Duncan, Lindis Oma, Yassar Nasser, Kimberley Benwell, Sara Coupe, Pilar Recarte, for their valuable support throughout the study.

In addition, we would like to thank our partners at the BIS, IMF, World Economic Forum, IDB, CGAP and the AMF for their immense help in disseminating this study as well as their perspectives outlined within the report. We would also like to thank WBG staff (Gian Boeddu , Ezio Caruso, Erik Feyen, Serafin Jaramillo, Karol Karpinski, Sergio Mesquita, Yasemin Palta, Srishti Sinha and Alexandra Traverso) for the survey design inputs and preparation of the WBG sub-section. We would also like to acknowledge Agustin Carstens, Gillian Tett, and Michael Vessey for their support with the launch event.

The research team would like to thank the following individuals for their considerable support in survey dissemination and data collection (in no particular order): Samantha Lewis (ECSRC), María José Baqueiro (ASBA), Joanna Aleksandrowicz (SoftwareMind), Ewa Ward and Kimberley Moran (Bank of England), Hayden Cole and Ian Moloney (FinTech Council), Michael Vessey (CBUAE), Mohammed Moradi, Raja Adam and Nivedita Shrikent (DFSA), Herman Smit (NED Bank), Simran Singh and Matt Lowe (FCA), Paul Diggle (Aberdeen), Lawrence Sackey (GAB), Jochen Biedermann (WAIFC), Zinhle Mkhari, Glory Kasasi and Jacob Mkandawire, Valerie Muhwati (MEFMI) and Hernán Colmán.

We appreciate these individuals' significant contribution and willingness to participate as research partners to make this study possible.

We would like to extend our gratitude to Jo Barr for proofreading the report, Mike Green for designing the report and Katherine Selby and Charles Goldsmith for press and communications support.

Finally we would like to thank all the survey respondents who made this report possible

Images throughout this report used from Unsplash and Pexels

Executive Summary

Artificial intelligence (AI) is accelerating transformation in financial services globally – shaping how financial institutions operate, how vendors scale solutions, and how regulators safeguard prudential soundness, market stability, conduct of business and consumer protection. To help navigate this rapidly evolving landscape, the sector requires more substantive, global evidence.

The 2026 Global AI in Financial Services Report provides a unique contribution towards our collective understanding of the adoption and impact of AI in the financial services sector, at the intersection of financial services providers, AI vendors, regulators as well as users and consumers.

This global research initiative was conducted by the Cambridge Centre for Alternative Finance (CCAF) at the University of Cambridge Judge Business School, in partnership with Financial Innovation for Impact (Fii), the Bank for International Settlements (BIS), the International Monetary Fund (IMF), the World Economic Forum (WEF), the Inter-American Development Bank (IDB), the CGAP, the Arab Monetary Fund (AMF) and with the support of the UK Foreign, Commonwealth and Development Office (FCDO).

This report examines the current adoption of AI in financial services from industry, technology and vendor perspectives. It also assesses the impact of AI on the financial services industry, including fintechs and traditional financial institutions, from efficiency to productivity and profitability, as well as persistent organisational adoption and scaling challenges. The study then analyses the emerging risks associated with AI adoption in finance and discusses a range of the correlated governance issues. This report highlights the developing policy and regulatory approaches regarding the use of AI in financial services as well as the internal adoption of AI by central banks, financial regulators and supervisors across the world. Finally, this report looks ahead to 2030 and discusses some potential long-term trends relating to the increasing adoption of AI in financial services, for instance, relating to jobs and competition dynamics.

This study builds upon our 2020 CCAF-WEF AI[1] report that analysed early adoption of AI in finance among both fintechs and traditional financial services providers. It expands the previous survey database by capturing insights from 628 respondent organisations, including 203 fintechs, 149 traditional financial institutions, 146 AI vendors and 130 central banks and other financial regulators across 151 jurisdictions around the world.

We also invited our research partners to share their perspectives on the adoption of AI in financial services from their respective vantage points and based on our survey findings, covering a range of topics from macroeconomic impact, AI supply and value chains, model explainability, to consumer protection and supervision implications.

Key findings of the report include:

Adoption

The financial services industry is ahead of regulators in AI adoption, and fintechs are ahead of incumbents. 81% of surveyed financial services firms are adopting AI at some level, with 40% of industry respondents reporting advanced AI adoption ('Scaling' or 'Transforming'), more than double that of regulators (with only 20% reporting having advanced AI adoption).

Among surveyed industry respondents, however, only 14% currently see AI as transformational to their organisational strategy and competitive advantage, signaling a potentially significant execution and business integration gap. Within the financial services industry, fintechs lead incumbents by 47% to 30% in the adoption of advanced AI, and in reaching a Transforming stage of adoption (19% versus 6%). Workforce preparedness and AI investment levels emerge as key differentiators between fintechs and incumbents. Among the 130 regulatory authorities surveyed, 48% of them reported that they are still in the 'Exploring' stage for AI adoption or not engaged with AI at all. 33% of surveyed regulators are Piloting AI and 18% are Scaling.

GenAI and Agentic AI are the clearest frontiers with lower barriers for adoption than traditional machine learning methods. AI adoption remains highest for Classical Machine Learning (75%) and GenAI (71%) (despite only gaining traction since 2022), among surveyed industry respondents. However, agentic AI is already in active adoption among 52% of industry respondents, demonstrating rapid uptake in a relatively short period of time. 23% of industry respondents are at more mature stages (Scaling or Transforming) of adopting agentic AI, while 29% remain in piloting. Fintechs are again ahead of traditional financial institutions (traditional FIs) in agentic AI adoption (57% versus 45%).

Looking ahead, 81% of surveyed industry respondents state that agentic AI will be meaningfully achieved by 2030, making it the clearest growth frontier in AI technology. GenAI and agentic AI are now reported as more widely used than supervised learning, unsupervised learning, reinforcement learning, and time-series techniques, which may reflect lower engineering (adoption) barriers than newer, provider-packaged methodologies. 53% of surveyed industry respondents spend under USD 100,000 annually on AI yet still report high maturity in GenAI and agentic AI. Industry respondents in emerging markets and developing economies (EMDEs) are also reporting higher levels of deployment ('Scaling' and 'Transforming') than firms in advanced economies (AEs).

Current AI deployment remains concentrated in internal operations rather than business model reinvention. Four of the top five financial services industry AI use cases are back-office functions. The most common use cases at Pilot stage or beyond are internal: process automation (79%), data visualisation (75%), software engineering (75%), and data and knowledge management (69%). The leading front office use case is AI-powered customer support (74%), where fintechs lead at 82% versus 67% among incumbents. Fraud detection (57%) and credit risk modelling (54%) lead among risk and compliance applications. Overall, AI is primarily being used currently to improve execution rather than to fundamentally reconfigure business models, though 51% of more mature AI adopters are Piloting or deploying new financial products powered by AI versus 28% among less mature institutions.

Cloud and foundation model choices reveal a meaningful architectural divide across stakeholder groups. Most organisations are building on external models rather than training their own models from scratch: 63% of industry and 65% of regulators use internal workflows built on external foundation models. However, many also customise or develop some AI systems in-house. At the time of the survey, OpenAI was the most-used foundation model provider across all groups (76% of industry and 48% of regulators), followed by Google (57% of industry) and Anthropic (35% of industry and 33% of vendors). DeepSeek is used by 15% of industry respondents.

24% of surveyed regulators do not use any foundation models. In terms of cloud infrastructure, Amazon Web Services (AWS) leads industry (46%) and vendors (55%), while 46% of regulators report using no cloud infrastructure at all. According to our survey data, the top three cloud providers serve more than 80% of industry. Among regulators who do use cloud, Azure leads at 39%. Traditional FIs remain more reliant on on-premises or local cloud deployments (39%) than fintechs (23%), which is a finding that is consistent with the analysis of the 2020 CCAF-WEF AI report.

Impact

Productivity effects are already felt, but enterprise value remains harder to evidence. Positive productivity impacts brought about by AI are perceived to be highest in technology, data and product functions (79%), followed by back office and operations roles (75%) and front office roles (69%). However, 55% of industry respondents and 63% of surveyed regulators find it difficult to measure the value of AI deployment, rising to 76% among large financial institutions.

Profitability outcomes are positive but uneven, correlating with AI investment and workforce preparedness. 40% of respondents report increased profitability from AI, while 43% report no change. Higher spend appears strongly associated with greater impact: 62% of organisations spending more than USD 100,000 annually on AI have reached advanced maturity, and 62% of that group report increased profitability, compared with 39% among lower-spending organisations. Fintechs again outperform, with 56% reporting higher profitability versus 34% of traditional FIs.

There is meaningful convergence across industry, vendors, and regulators on several core themes. All three groups identify greater operational efficiency as the top expected benefit of AI by 2030 (73% of industry, 66% of vendors, and 56% of regulators). There is also broad alignment on the need for clearer regulatory guidance, ranked as a top priority by 69% of industry, 67% of vendors and 79% of regulators. More broadly, all groups place high importance on privacy, accountability, and the need for human oversight, suggesting that while institutional perspectives differ, substantial common ground already exists on the governance frameworks needed for responsible AI deployment.

Challenges

Data quality, talent and legacy architecture remain the core constraints to adoption and scaling. These bottlenecks are not new: data quality and talent access were already identified as the top two barriers in the 2020 CCAF-WEF AI report. Data availability and quality remain the leading pain point hindering AI adoption, cited by 66% of AI vendors, 46% of regulators, 40% of industry (of those, 49% of traditional FIs and 34% of fintechs). Vendors also report specifically acute data-related challenges when working with their clients: 72% cite data quality and completeness, 46% legacy systems and siloed environments, and 41% report data-sharing restrictions. For surveyed regulators, lack of AI training and capacity building (48%), talent (47%), and technology and infrastructure (45%) were also core constraints for AI adoption in addition to data issues.

Risks

There is broad consensus on the top risks of AI in financial services. Data privacy and protection (cited by 65% of AI vendors, 74% of industry and 80% of regulators), and model hallucinations and unreliable outputs (cited by 67% of AI vendors, 70% of surveyed industry firms and 70% of regulators) were rated as the top two risks by all stakeholder groups. Operational and cyber resilience (59% of regulators and 46% of industry), model opacity and lack of explainability (56% of regulators), loss of human oversight (55% of industry and 51% of AI vendors), adversarial AI-related cyber threats (50% of industry and 57% of regulators) and algorithmic bias and fairness (43% of vendors), also feature prominently among the top risks.

Important divergences remain in risk perception, accountability, and market expectations. Regulators are markedly more concerned than vendors about cyber and operational resilience (59% versus 32%), critical third-party risk (43% versus 23%) and consumer protection and bias (41% versus 21%). Notably, industry, especially traditional FIs, is more concerned than regulators about the loss of human oversight and collective forgetting (60% versus 42%). Views on accountability are also fragmented: industry (35%) and vendors (39%) most often favour a case-by-case approach, whereas regulators (38%) most often place primary responsibility on the regulated financial institution (only 18% of industry and 16% of vendors do). Vendors (24%) and industry (22%) are more open than regulators (9%) to shared accountability. Regulators are also more concerned with concentration issues than industry (43% versus 28%).

The rapid deployment of agentic AI compounds cyber vulnerabilities, rendering manual oversight increasingly ineffective. Software engineering is the financial industry's most mature AI application (42% fully deployed, 33% in development) and is also a primary cyber risk transmission vector. Broadly, 51% of respondents cite the "loss of human oversight" as the third highest AI risk overall; in software engineering specifically, the unprecedented volume and velocity of AI-generated code make traditional manual reviews increasingly ineffective. This structural vulnerability is compounded by external threats: 48% of respondents flag adversarial AI as a top concern, reinforced by Anthropic's recent "Mythos" disclosures which point to an imminent future where next generations of AI models are set to be incredibly capable at exploiting software vulnerabilities presenting both firm-level cyber resilience as well as systemic financial risks. Further complicating this problem space is a notable perception gap: AI vendors place less priority than industry and regulators on both adversarial AI threats (35% versus 50% industry, 57% regulators) and cyber/operational resilience (32% versus 46% industry, 59% regulators). These intersecting vulnerabilities can also feed into the top perceived risk across all stakeholders – data privacy and protection (73% of respondents) as sensitive data is typically the primary target for the cyber exploits these vulnerabilities enable.

Regulation and supervision

Among regulators, supervision remains the dominant use case for AI, but it is also applied in licensing and policymaking , rule-making, macro-prudential analysis and operational functions. Most of the internal AI use cases for regulators relate to supervisory functions and activities such as market surveillance and misconduct detection (31% of surveyed regulators are either piloting or already deployed), anti-money laundering and counter-financing of terrorism (AML/CFT) supervision (27%) and consumer protection (25%).

However, AI is also being used in other regulatory functions including; licensing and authorisation (for example, fit and proper checks, application screening and ownership structure inspection) and policy and rule-making processes (for example, horizon scanning and risk identification, and consultation analysis and drafting). The most referenced external AI frameworks are the EU AI Act (42%), financial sector-specific guidance from standard-setting bodies (41%), and the ISO/IEC AI standards (for example, ISO 42001) (27%).

A clear explainability-expectation gap looms. 79% of surveyed regulators rate explainability as critical or important to their regulatory objectives, yet 50% of industry have adopted explainable AI methods, and vendors perceive more than half of their clients as having low or no expertise in the use of explainability tools. About two-thirds of industry respondents are not monitoring for bias or arbitrary discrimination, exclusion or systemic bias in AI, and only 37% are concerned about model explainability and opacity as an operational risk.

Regulators are generally optimistic about AI's role in achieving their objectives by 2030. 78% of surveyed regulators view AI as significant or transformative for supporting their objectives by 2030, with 29% rating it as potentially transformative. Surveyed regulators also see that AI usage can have a favourable impact on supporting financial inclusion (49% supportive versus 12% that see it as challenging), fighting financial crime (42% supportive versus 18% challenging) and data sharing via open banking/finance (37% supportive versus 9% challenging). Regulators are less sure about AI's overall impact on consumer protection (with 27% viewing it as supportive versus 27% challenging), competition (23% supportive versus 15% challenging), financial stability (22% supportive versus 13% challenging) and technology/cyber resilience (19% supportive versus 27% challenging). On global cooperation, regulators are cautiously optimistic: 48% say cooperation is challenging today but likely to improve, while only 9% are pessimistic.

2030 outlook

Reskilling, not displacement, is the dominant workforce expectation for now. 10% of industry respondents expect a net increase in jobs and around 25% expect significant reskilling and job transformation without large net losses. Meanwhile, 24% of industry respondents expect a net reduction in roles, more than the last three years (where 13% of surveyed industry respondents saw job losses, but an equal number also saw job gains). Industry respondents suggest that commercial and wholesale banking is most likely to see net increases in jobs (44% of respondents) and payments less likely (10%). Interestingly, 58% of industry respondents state that their own organisation is likely to see a net increase or reskilling in jobs (36% expect a net increase and 22% expect reskilling and transformation).

Competition, market dynamics and consolidation are expected to shift materially by 2030. Vendors expect greater disruption to market dynamics, with 55% anticipating either winner-takes-all or fragmented market outcomes, while regulators are the most cautious, with 52% stating it is too early to tell. Expectations of competitive disruption have shifted dramatically since 2020: 42% of respondents then believed the market status quo would prevail, versus only 8% today.

Artificial General Intelligence (AGI) and Artificial Super Intelligence (ASI) are expected to be meaningfully achieved by 2030 by a material number of respondents. 44% of all respondents expect AGI to be achieved by 2030 (defined as when AI reaches parity with humans on all tasks), with that expectation higher among AI vendors (51%) and industry (50%) than regulators (28%). In terms of perception of priority risks today, all three stakeholder groups ranked the emergence and impact of AGI 21st out of 22 surveyed risks, selected by just 9% of respondents. 28% of all respondents have gone further to expect the realisation of Artificial Super Intelligence (ASI) by 2030.

Introduction and Approach

The study employed three parallel, purpose-built survey instruments designed to capture the intersecting perspectives across the financial services ecosystem, securing a total of 628 respondents drawing perspectives from across 151 countries.

- **Industry survey**: Targeted both fintechs (n=203) and traditional FIs (n=149) across banking, insurance, payments, capital markets, digital assets and other financial services sectors.
- **AI vendor survey**: Targeted technology companies (n=146) providing AI products and services to the financial services sector, including both specialist financial services AI vendors and general-purpose AI platforms.
- **Regulator survey**: Targeted financial regulators, supervisory authorities, and central banks (n=130) with mandates spanning prudential regulation, securities oversight, payments supervision, consumer protection and financial stability.

Figure A.1: Sample demographic of the 2026 AI in Finance Global Report

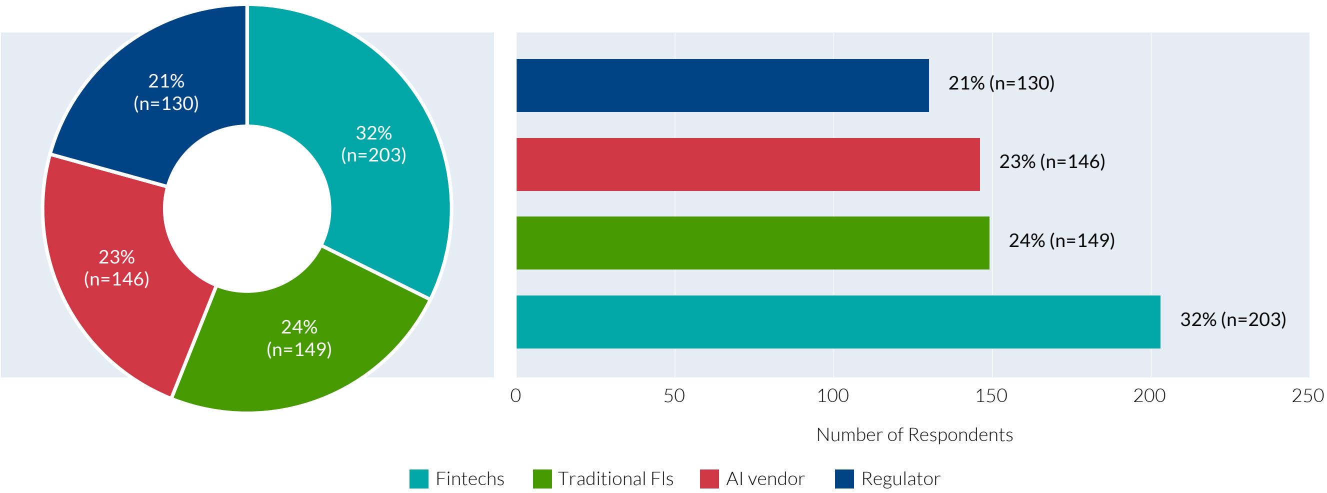

The survey data sample consists of companies operating in many geographies, totalling 124 headquarter countries and 151 operational countries.

The scope of this study expands upon the 2020 CCAF-WEF AI Report to also include AI vendors and regulators; enabling a unique cross-sectoral analysis on shared themes, such as AI adoption maturity and impact, AI technology and tool use, perceptions of risk and governance, accountability and the regulatory environment.

Figure A.2: Survey respondent jurisdictions coverage

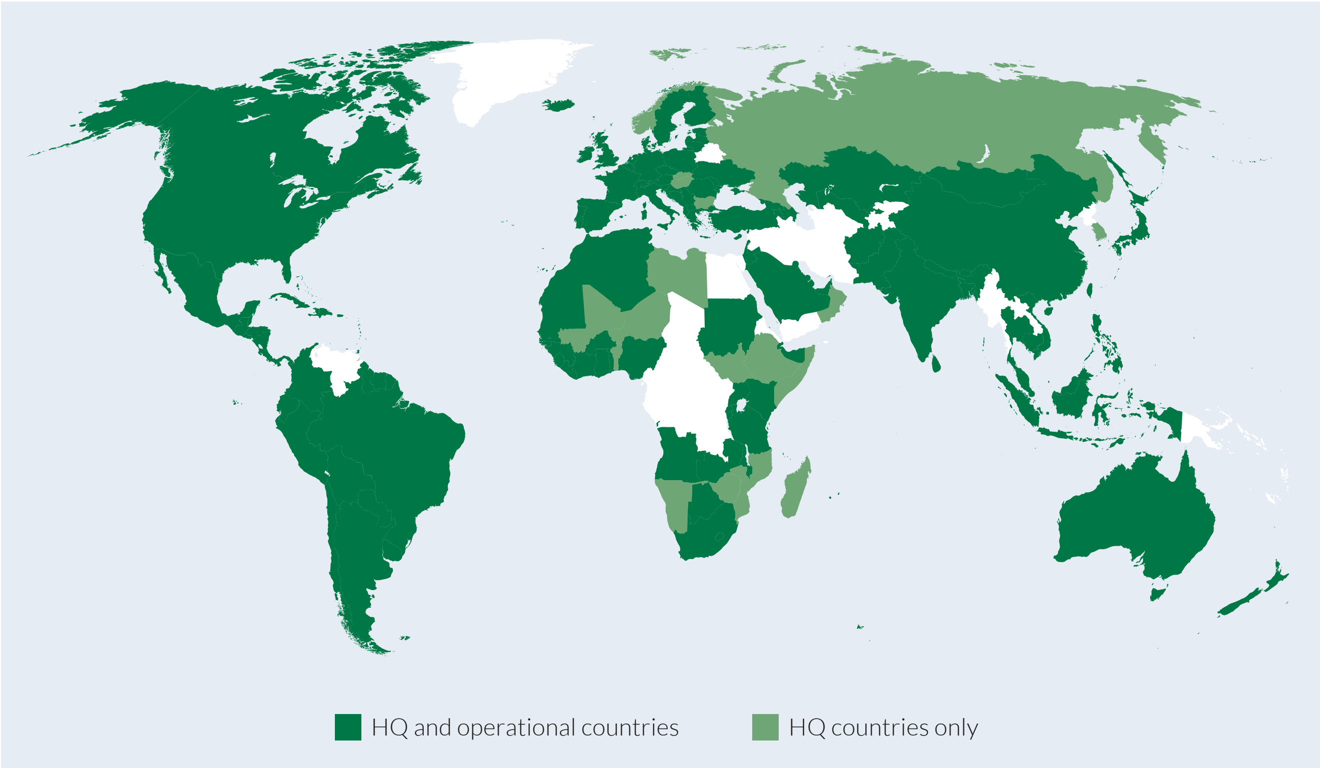

Survey design

The online survey was designed by the CCAF / Fii team and went through an extensive peer review process to ensure clarity and robustness. The Research Partners – the World Bank, IMF, World Economic Forum, CGAP, AMF, IDB and BIS – did not have access to raw data and were involved only in providing survey methodological input, supporting distribution, and contributing to the final report with only anonymized data. The online survey was designed by the CCAF / Fii teams and went through an extensive peer review process to ensure clarity and robustness.

Data analysis and limitations

The data analysis followed a structured process, allowing the research team to draw granular insights as well as make comparisons between different data cuts. Further details of this can be found in the appendix.

Data collection

All three surveys were administered via web-based questionnaires on Qualtrics and distributed through partner networks, direct outreach and LinkedIn. Data collection was conducted between October 2025 and January 2026. The surveys were available in English, Spanish, Portuguese, French, Arabic, Simplified Chinese, Tagalog, Thai, Ukrainian and Vietnamese. The surveys were distributed globally through the partner organisations listed in the appendix and via professional networks, leveraging established relationships with financial institutions, AI vendors and regulatory authorities across multiple jurisdictions.

Participation was voluntary and responses were collected on a confidential basis. Respondents included senior management, C-suite executives, technology leads and regulatory officials with direct knowledge of their organisation's AI strategy, adoption and governance practices.

Data cleaning and validation

To ensure data security, confidentiality and data leakage, all data were anonymised, and analysis was done with aggregated data; these measures avoid disclosing personal and institutional information.

A systematic data cleaning and validation pipeline was applied to all three surveys prior to analysis. This process included:

- **Invalid responses identification**: Incomplete and duplicate responses were removed from each survey.
- **Response validation**: The bot detection function was enabled to flag potential bots and individual responses were reviewed for consistency. Key data fields were checked, such as consistency between countries and regions, revenue outliers and AI spend outliers.
- Percentages: Figures have been rounded for readability and ease of interpretation. As a result, totals in some figures may not sum exactly to 100%.

Limitations

The surveys aim to show global and regional perceptions of AI in financial services. However, the research has limitations that should be noted when interpreting the findings of this study:

- **Voluntary participation**: The surveys were distributed through partner networks and participation was voluntary, which may introduce self-selection bias and limited sample sizes for some questions. Efforts were made to ensure representative samples in all three surveys, in terms of geographical dispersion, income groups and type of regulatory bodies surveyed.
- **Self-reported data**: All responses reflected the perceptions and self-assessments of respondents rather than independently verified metrics. Reported levels of AI maturity, spend and workforce preparedness should be interpreted as perceived rather than objectively measured values. However, as aforementioned, respondents were officials with direct knowledge of their organisation's AI strategy, adoption, and governance practices.
- **Sample composition**: While the sample covered 151 jurisdictions, the distribution was not uniform. Some regions and sub-sectors had smaller representation, which limits the granularity of country, regional or sectoral-level analysis. Findings for sub-groups with small sample sizes should be interpreted with appropriate caution. There is also a difference in the sample sizes of each survey in terms of geographical distribution which limits cross-survey comparison in some instances. The sample includes a relatively higher share of smaller financial institutions and participation in the survey was voluntary. As such, the results may reflect some degree of self-selection bias, with institutions more engaged in AI potentially more likely to respond, while less mature firms

may be underrepresented. In addition, aggregate results may mask differences across firm sizes and types, particularly for larger or systemically important institutions. Where relevant, the analysis seeks to highlight these differences, however, the findings should be interpreted with these limitations in mind.

- **Cross-sectional design**: The study captured a single point in time. Causal relationships between variables (for example, between AI spend and profitability) cannot be established from this data. Observed associations should be treated as correlations rather than causal findings. Where possible, associations with past research efforts will be identified, with the limitation that both the nature of AI and the questions posed then and now are not identical.
- **Question comparability**: While standardised wording was used where possible across the three surveys, not all questions were identical. Some questions were tailored to the specific context of industry, vendor or regulatory respondents, which may limit direct comparability on certain themes. This was an intentional choice to bring a nuance to the survey in terms of the industry, AI vendors and regulators.

Analysis approach

The analytical framework employed a combination of descriptive statistics, binary segmentation and cross-survey comparison to identify patterns across the dataset.

Cross-survey mapping

A structured mapping framework was developed to enable cross-survey comparison. Questions addressing common themes (for example, AI technologies used, risk perception and accountability preferences) were mapped across the three survey instruments using a standardised question wording reference. This mapping supported the All-Surveys Aggregation analysis, where responses from all three stakeholder groups were combined and compared on equivalent questions.

Exploratory data analysis platform

A purpose-built data analysis dashboard was developed using Streamlit (an open-source Python-based data analysis framework) to support the exploratory analysis, validation and visualisation of the survey data. The dashboard served as the primary analytical tool for the research team throughout the study.

Cross tabulation analysis

Twelve binary segmentations were applied to the survey respondent datasets. These cuts of the data provided sizable sample sizes, enabling systematic comparison of responses across key organisational characteristics. Further details on this can be found in the appendix. These segments included:

- Institution type (fintech versus traditional financial institution)
- Economic development (advanced versus emerging and developing markets)
- Geographic regional base
- AI maturity (transforming and scaling versus exploring and piloting)
- AI spend (above versus below USD 100,000 in the most recent financial year)
- Firm size (above versus below 50 Full Time Employed (FTE) in the most recent financial year)
- Revenue (above versus below USD 10 million in the most recent financial year)
- Profitability impact (increased versus negative or no change)
- Productivity impact (positive versus negative or no change)
- Workforce preparedness (moderate/highly prepared versus limited/not prepared)
- Value measurement difficulty (challenging versus easy/neutral)
- AI offering type (tailored financial services versus general-purpose)

Survey demographics

Respondents by economic development

A key goal of this study is to offer comparative analyses not only by region, but also by economic development. This allows this study to further explore intrinsic characteristics of AI adoption and its effects across the different groups of stakeholders surveyed by the income level of the countries they operate.

Classification in Advanced Economies (AEs) and emerging market economies was undertaken using World Bank income classifications (World Bank, 2025),[2] with those classified as high income being placed in the advanced economy category, and economies of low income, lower-middle income and upper-middle income being classified as emerging markets and developing Economies (EMDEs).

Figure A.3: Sample demographics AEs versus EMDEs

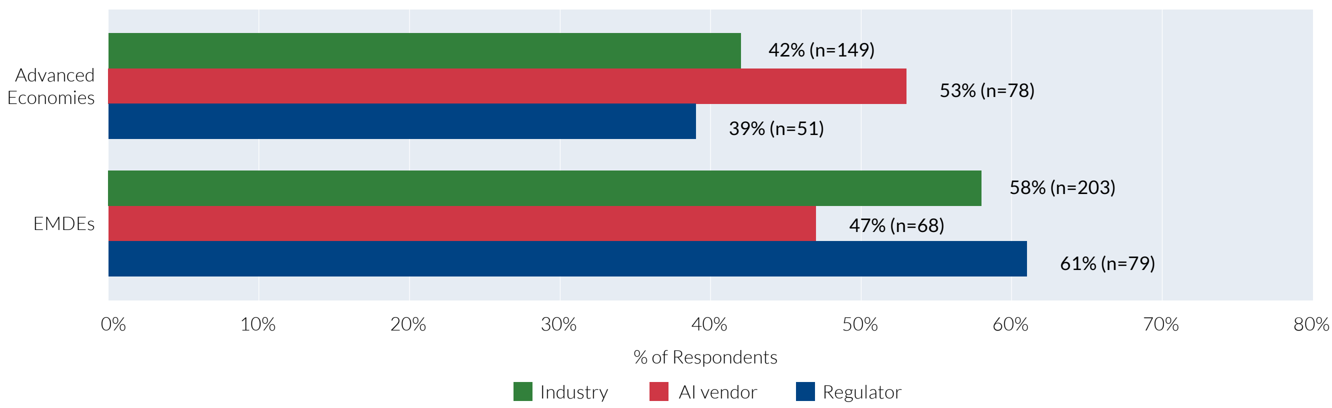

Overall, 55% of respondent organisations are based in EMDEs, compared with 45% in AEs, which provides a useful data cut for analysis in this study.

Region of operation

Table A.1: Distribution of operating firms by region with the number of data points

Region	Industry %	Industry	AI Vendor %	AI Vendor	Regulator %	Regulator
Europe	33%	115	45%	66	27%	35
Asia Pacific	36%	125	39%	57	12%	16
Latin America & Caribbean	29%	101	30%	44	25%	32
Middle East & North Africa	16%	56	27%	39	19%	25
North America	16%	56	38%	56	2%	2
Sub-Saharan Africa	14%	50	22%	32	15%	20
Total responses	–	352	–	146	–	130

The regional distribution of the industry survey is robust, with at least 50 firms in each region. APAC, Europe and LAC make up the top three with 29% to 36% of firms, while SSA had the lowest proportion with 14% of firms operating in the region. For the vendor survey, there were at least 32 respondents in each region.

Industry vertical

Figure A.4: Financial Sector Composition of the Sample (Multiple choice question)

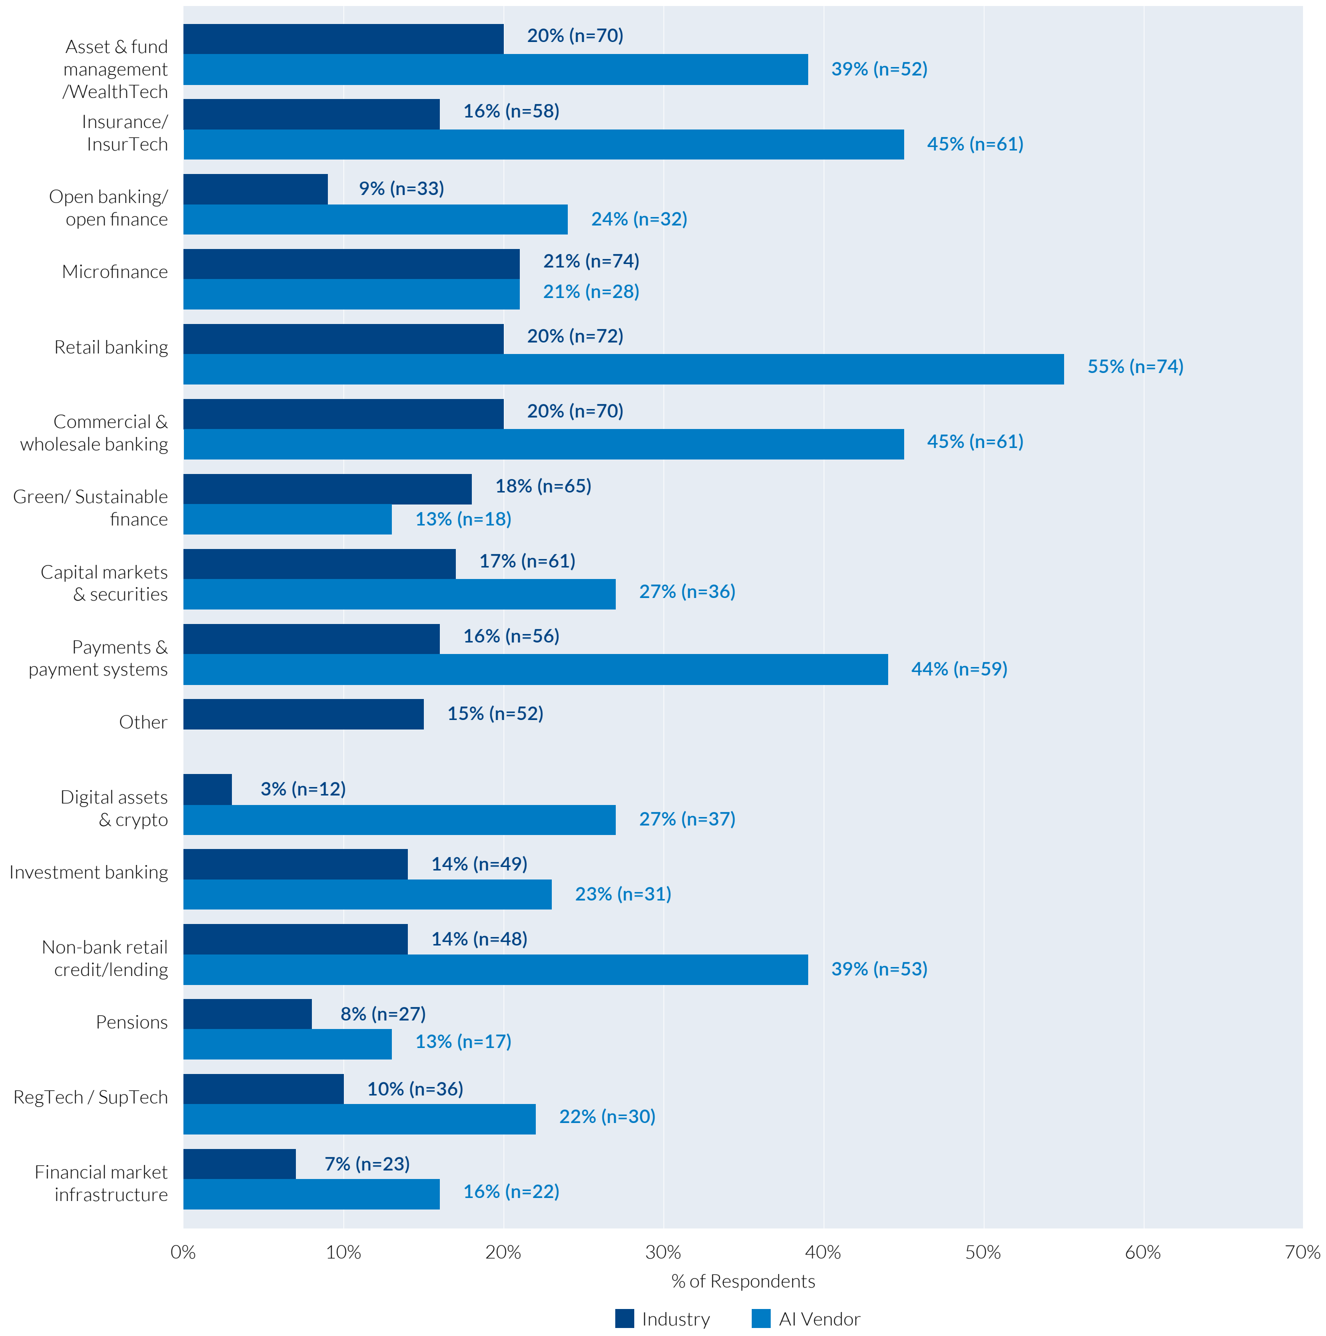

Industry participants selected their primary and secondary financial activities in which they operate, while AI vendor survey participants selected the financial sector they serve. The sectors are based on the CCAF alternative finance classification.[3] Overall, the data set is composed of at least 20 firms engaging in each financial industry sector.

Number of Full Time Employees (FTEs)

Figure A.5: Sample Firm Size: ≤50 versus >50 Employees

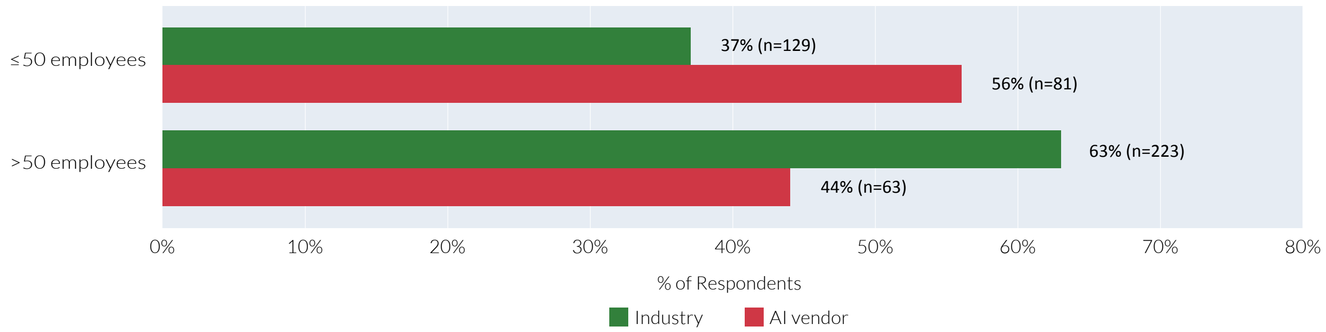

Fintechs have over 80% of firms with fewer than 50 or 51-500 employees, whereas most traditional FIs have 51-500 employees and 501-5,000 employees.

With regards to AI vendors, 80% of survey respondents have up to 500 employees, with most having fewer than 50 employees.

Table A.2: Distribution of responses by the number of full-time employees (FTE)

FTE Range	Industry %	Industry n	AI Vendor %	AI Vendor n
less than 50	37%	129	56%	81
51-500	31%	110	24%	35
501-5,000	18%	64	10%	15
5,001-50,000	9%	30	5%	7
more than 50,000	5%	19	4%	6
>50 employees (combined)	63%	223	44%	63
Total (responded)	100%	352	99%	144

Regulator demographics

Regulator institution type

Central banks, prudential regulators and securities market regulators responded to this survey alongside some ministries of finance and data and competition regulators. It is often the case that prudential regulators have both securities market oversight and banking oversight and thus the categories below should not be considered as mutually exclusive.

As AI can straddle many different regulatory mandates, there was a smaller sample of respondents that were in charge of specific issues relating to AI regulation, such as data and information regulators, as well as some government ministries who within their own jurisdiction might have oversight over the broader AI regulatory framework.

Figure A.6: Regulator Type by Economic Development

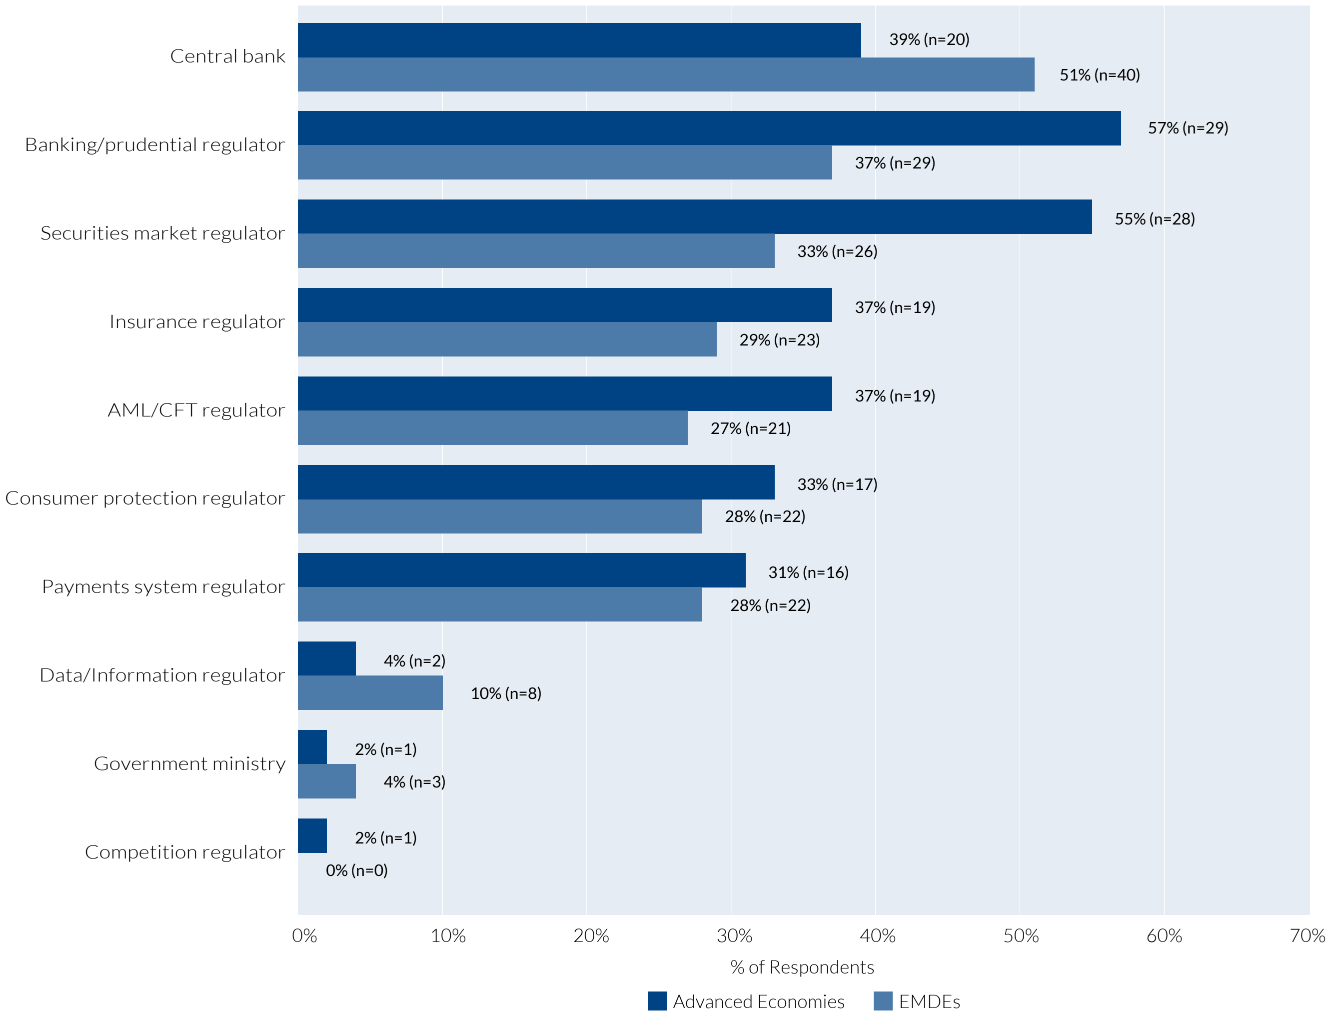

Chapter 1

Overview of Adoption of AI in Financial Services

Across the 352 industry respondents and 130 regulators surveyed, the data reveals a global financial ecosystem in mid-transition: broadly engaged with AI yet heavily concentrated in the implementation stages that precede institutional transformation.

Key highlights

- **The execution gap**: 81% of industry respondents are adopting AI at some level, however, only 14% view their deployment as transformational to their organisational strategy and competitive advantage – signalling a deep execution gap between early-stage experimentation and institution-wide AI integration.
- **The economic divide**: Economic context is one the sharpest dividing lines: advanced economies (AEs) are nearly twice as likely to have reached the Transforming stage (19%) than their emerging markets developing economy (EMDE) peers (9%), reflecting structural gaps in data infrastructure, talent and resourcing.
- **Fintechs are ahead**: Fintechs are moving faster than traditional FIs. 87% of fintechs are now Piloting AI or in more advanced stages of deployment (up from 76.6% in 2020) and are three times more likely to have reached organisation transformation (19% of fintechs versus 6% within traditional FIs).
- **The regulator innovation delta**: Only 10% of EMDE regulators report reaching the Transforming or Scaling stages of AI adoption compared with 33% AE regulators. Tighter budgets, limited capacity and scarcer AI talent are likely constraining regulatory modernisation with AI.
- **Evolving technology adoption**: Classical Machine Learning (ML) is widely adopted (76% fintechs, 73% traditional FIs, 53% regulators). However, newer technologies are rapidly scaling given lower technical barriers to adoption compared with more traditional ML methods. Traditional FIs cite slightly higher adoption in GenAI than fintechs (75% versus 69%), while fintechs report higher adoption of agentic AI (57% versus 45% within traditional FIs).

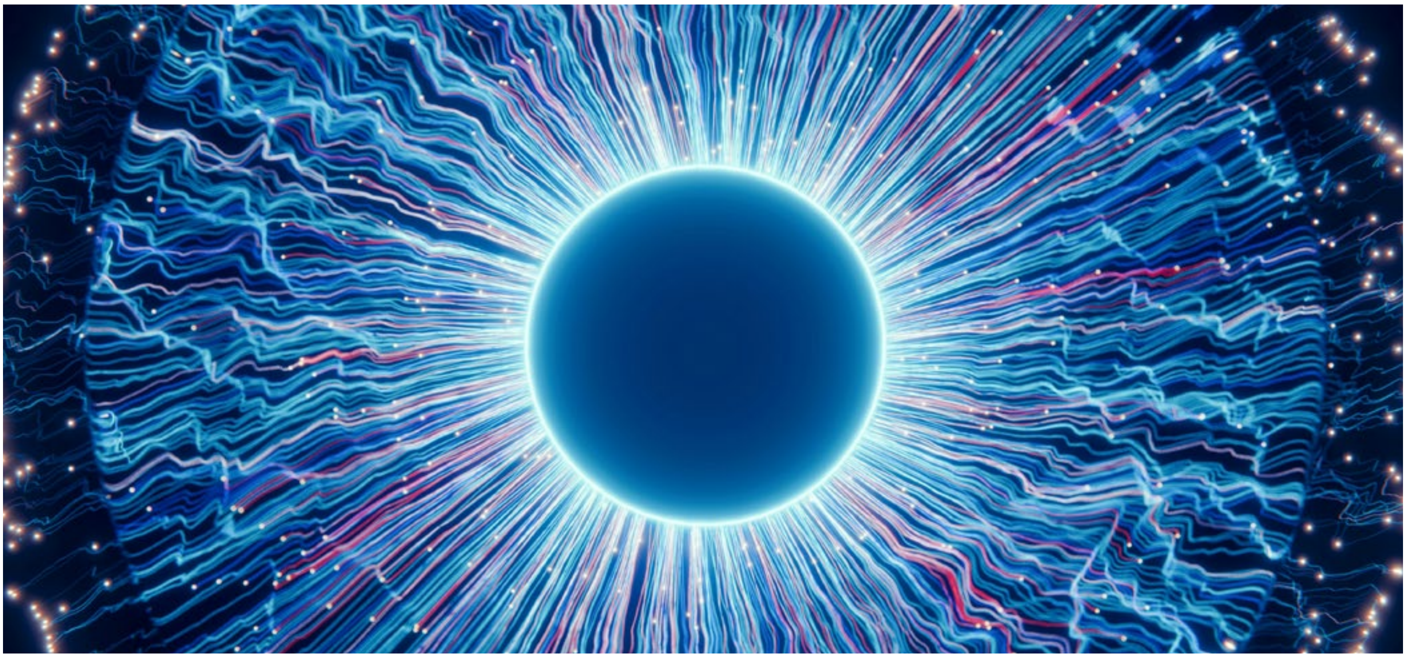

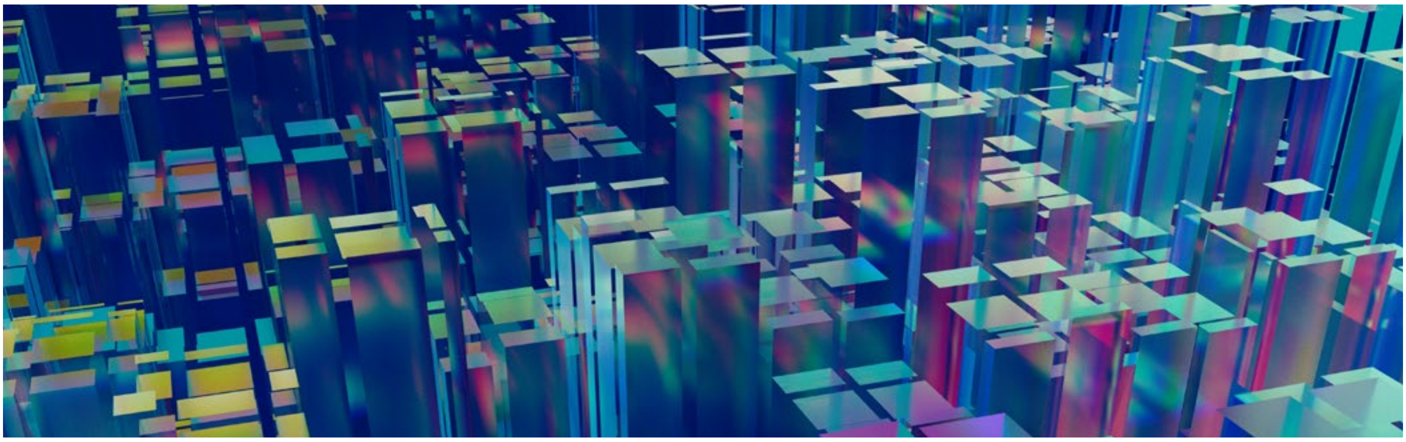

Industry versus regulators

Comparing the maturity of AI adoption between industry (fintech and traditional FIs) and regulator respondents shows marked differences. There is much greater concentration of more advanced AI use within the financial industry which are twice as likely as regulators to report being at either Scaling or Transforming (40% vs 20%) including a sevenfold gap in those at the Transforming stage alone (14% versus 2%).

Figure 1.0: AI adoption maturity: industry (n=352) versus regulators (n=130)

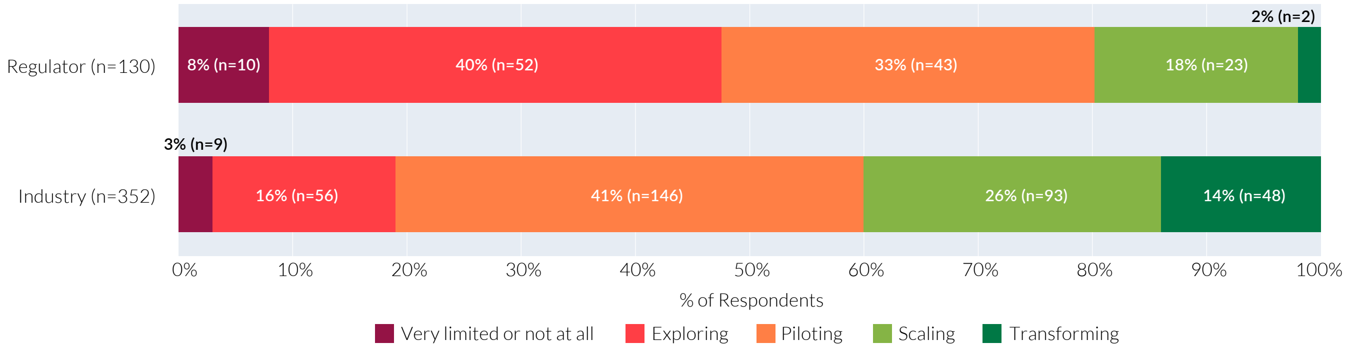

In contrast, regulators remain clustered within earlier stages of AI adoption maturity. They are more likely to report being in the Exploring stage (40% versus 16% for industry) while both groups report relatively low shares of limited or no adoption. These differences are likely driven by institutional context. Firms operating in competitive markets face strong incentives to turn AI into operational and commercial value quickly. Regulatory bodies typically operate under different institutional incentives, procurement constraints, and risk tolerances thereby naturally adopting a more cautious posture.

However, this gap is also compounded by resource constraints. Limited access to internal technical talent, data infrastructure, and restricted investment availability all serve to slow regulator adoption. This maturity gap is further explored in later chapters which shows it to be a structural condition. Regulators face capacity constraints around building technical expertise, in risk governance priorities, and in supervisory infrastructure deficits that regulators themselves identify as their most pressing constraint.

Fintech versus traditional financial institutions

Figure 1.1: AI adoption maturity by firm type – fintechs (n=203) versus traditional FIs (n=149)

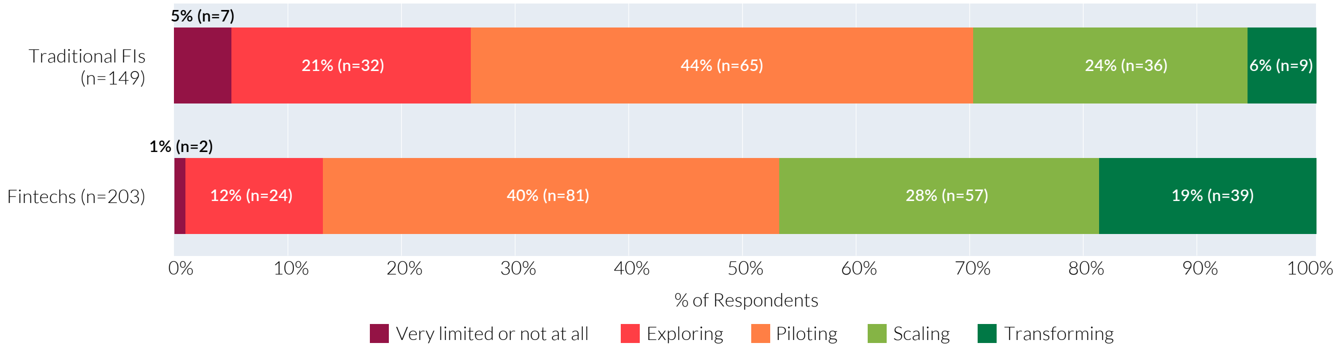

Within industry, the clearest differences between fintechs and traditional FIs appears at the most mature stages of AI adoption. Fintechs are more than three times more likely as traditional FIs to have reached the 'Transforming' stage (19% versus 6%). Traditional FIs show higher shares of Exploring (21%) and Piloting (44%), indicating slower progression through the AI adoption maturity curve.

This pattern is consistent with the view of fintechs, as digital-first, more agile adopters of new technologies. Established, traditional FIs typically face greater organisational inertia, legacy complexity and more demanding integration and security requirements that complicate the path to scaling deployment.

Economic development and AI adoption in industry

Figure 1.2: Industry AI adoption maturity by economic development – advanced economies (n=149) versus EMDEs (n=203)

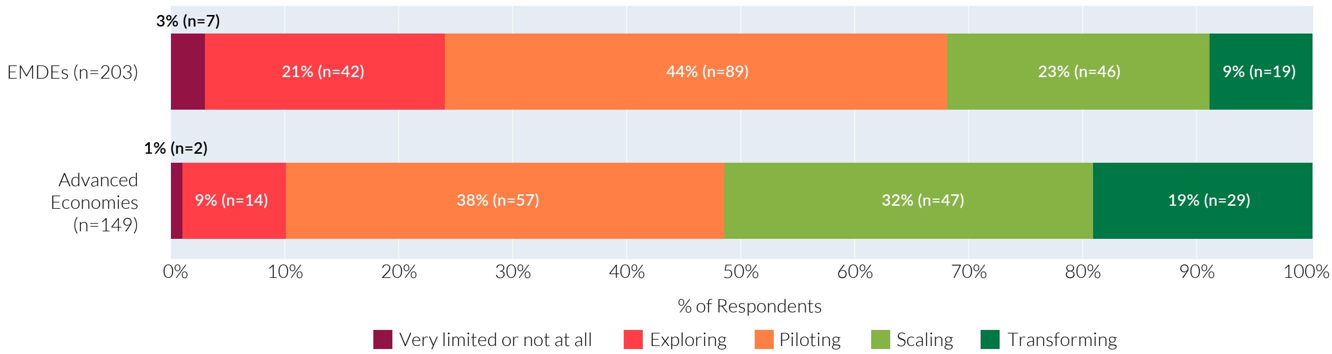

Macroeconomic context correlates with rates of AI adoption. In AEs, over 50% of surveyed firms reported advanced levels of AI adoption maturity. In emerging markets, this drops to less than a third. Furthermore, around a quarter of EMDE financial institutions are in the earliest phase of exploration or limited AI adoption. Emerging markets likely face more obstacles in the critical enablers of transformative AI adoption, such as access to AI talent, data quality, and technological infrastructure.

Economic development and adoption maturity in regulators

This economic divide is mirrored across the regulatory landscape too. Regulators in EMDEs report lower maturity than their advanced economy peers. Around one-third of AE regulators are already scaling AI for their operations, while more than half of EMDE regulators remain in exploratory stages (52%) or very limited (11%), while 27% are Piloting solutions. Only a very small number have reached more advanced stages of maturity.

Figure 1.3: Regulator AI adoption maturity by economic development – advanced economies (n=51) versus EMDEs (n=79)

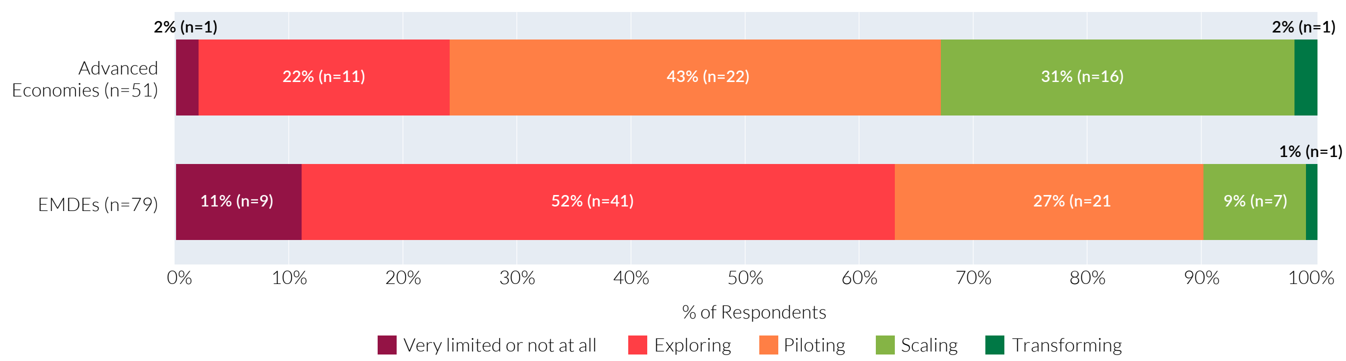

Regulators in EMDEs often face tighter budgets that limit investment in AI, weaker access to scalable technical infrastructure and challenges in competing for high-cost specialist AI talent. These factors likely slow the transition from early experimentation to embedded operational use.

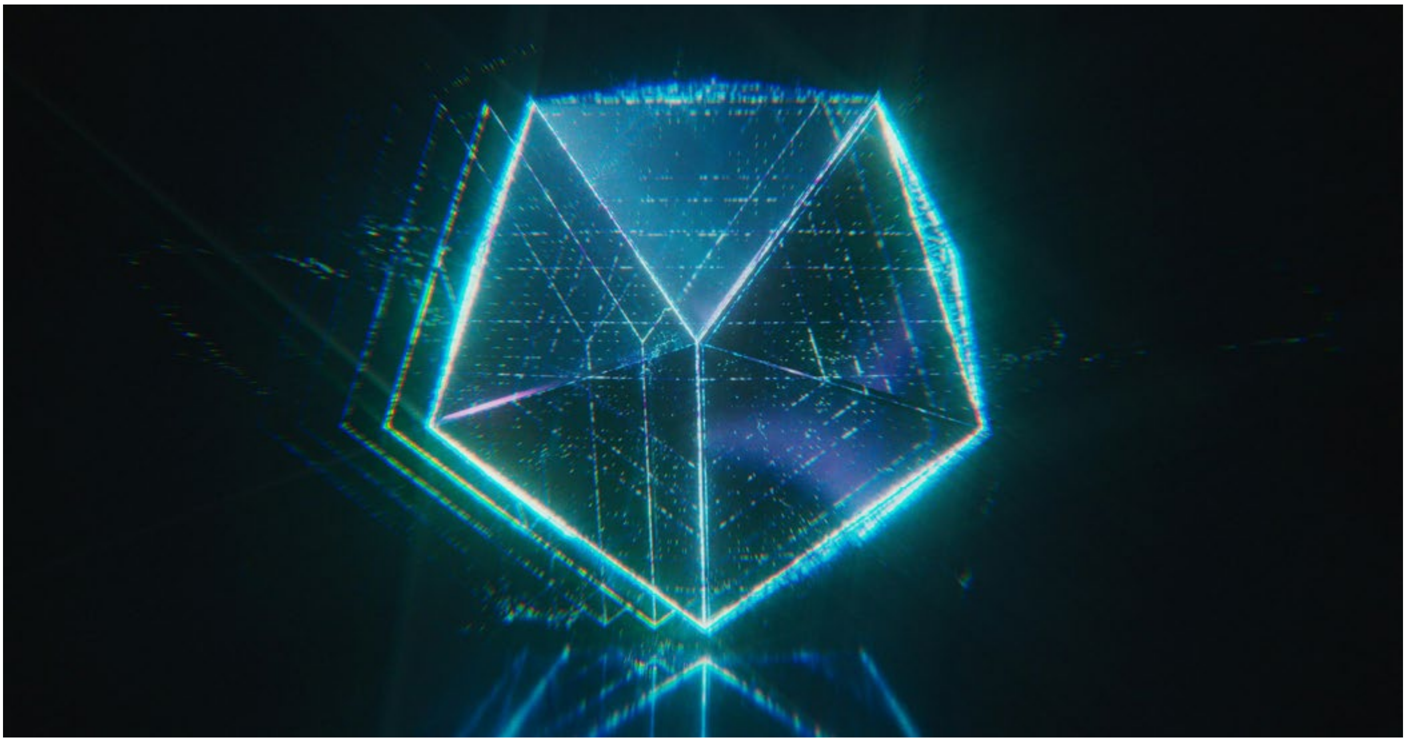

Adoption Rate of AI Category by Stakeholder Group

Beyond overall AI adoption maturity, tracking the adoption of specific AI categories – Classical ML, Generative AI (GenAI) and agentic AI – reveals an uneven state of play across stakeholder groups.

Figure 1.4: Active AI adoption by AI type and stakeholder group – % active adoption covers Piloting, Scaling or Transforming

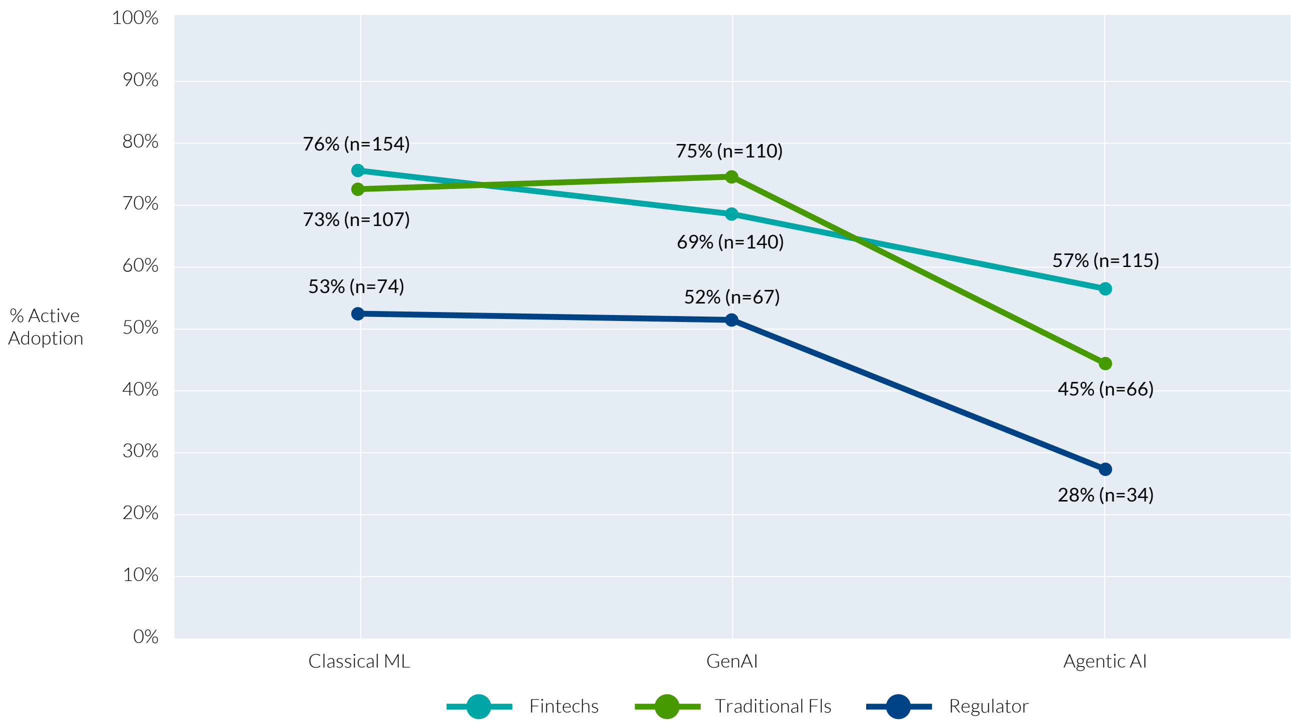

Classical ML is well established within the industry, with three-quarters of respondents citing active adoption, while just over half of regulators are using these techniques. Classical machine learning systems learn statistical patterns from labelled historical examples: during training, a model is exposed to past data, such as transactions, applications and customer records alongside known outcomes, and adjusts its internal parameters until it can reproduce those outcomes reliably. The family, as defined here, spans a wide range of architectures, from logistic regression and gradient-boosted trees to deep learning, which comprises multi-layer neural networks that learn representations directly from raw inputs, such as transaction sequences or document images, without requiring manual feature engineering. Common financial services applications of this ML paradigm include fraud scoring, credit underwriting, anti-money laundering anomaly detection, and time-series forecasting for liquidity and market risk.

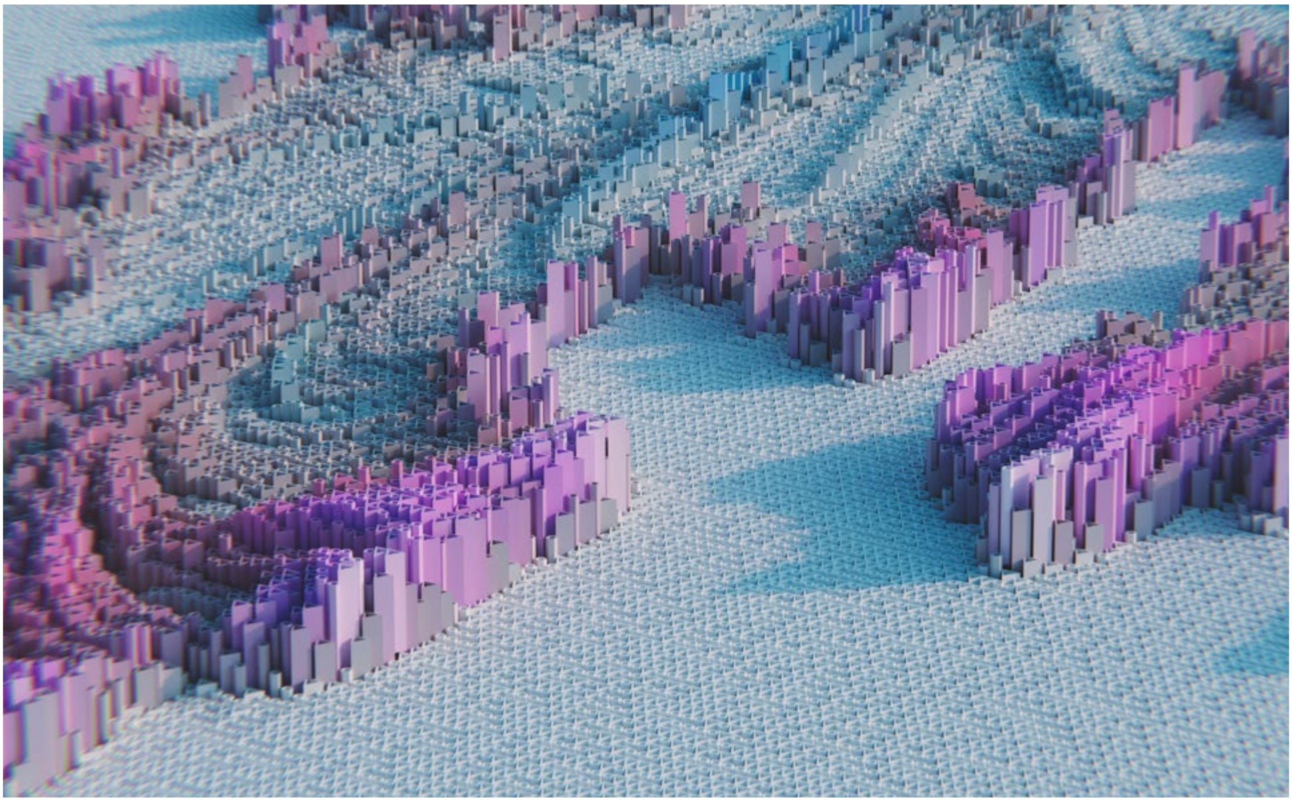

GenAI, having gained traction since the launch of ChatGPT in late 2022, has already reached adoption parity with Classical ML methods. Notably, traditional FIs have slightly higher levels of adoption than fintechs (75% versus 69%) and more than half of regulators also show a meaningful level of uptake. GenAI, as defined here, encompasses foundation models, which refer to large ML models trained once on a vast corpus of text, code, or other data, and then adapted to many downstream generation tasks. Training such architectures from scratch is highly computationally intensive and is undertaken by a small number of specialised firms, meaning that many financial institutions, AI vendors, and public sector bodies that deploy generative AI do so by drawing on models they did not build themselves. This separation between model development and model deployment characterises the current GenAI supply chain and may also explain the rapid uptake visible in the survey responses, as packaged AI solutions from foundation model vendors require little to no AI engineering or technical capabilities compared with Classical ML pipelines.

Agentic AI refers to systems that pursue objectives through autonomous, multi-step sequences of actions. An agentic system is given a goal and works towards it: decomposing the objective into sub-tasks, selecting and calling tools (e.g., foundation models, databases, APIs), observing intermediate results, and revising its approach accordingly. Currently, 57% of fintechs, 45% of traditional FIs, and 28% of regulators cite some level of adoption at the piloting stage or beyond.

Agentic AI is a priority area for supervisory attention because it introduces autonomous decision-making characteristics, which create new challenges and risks relating to issues including, but not limited to, model risk management, consumer protection and concerns around liability and accountability. Later chapters explore these issues in more detail, examining the risk and accountability questions that agentic AI raises and documenting the supervisory infrastructure available to address them.

AI maturity across data groupings

To help understand the landscape of AI adoption, the survey data was segmented across several key variables enabling comparison across cohorts. Clear correlations emerge with firm type (fintechs), size (over 50 employees), geographic base (AEs) and most critically, AI investment.

Figure 1.5: Industry AI adoption maturity by comparison grouping

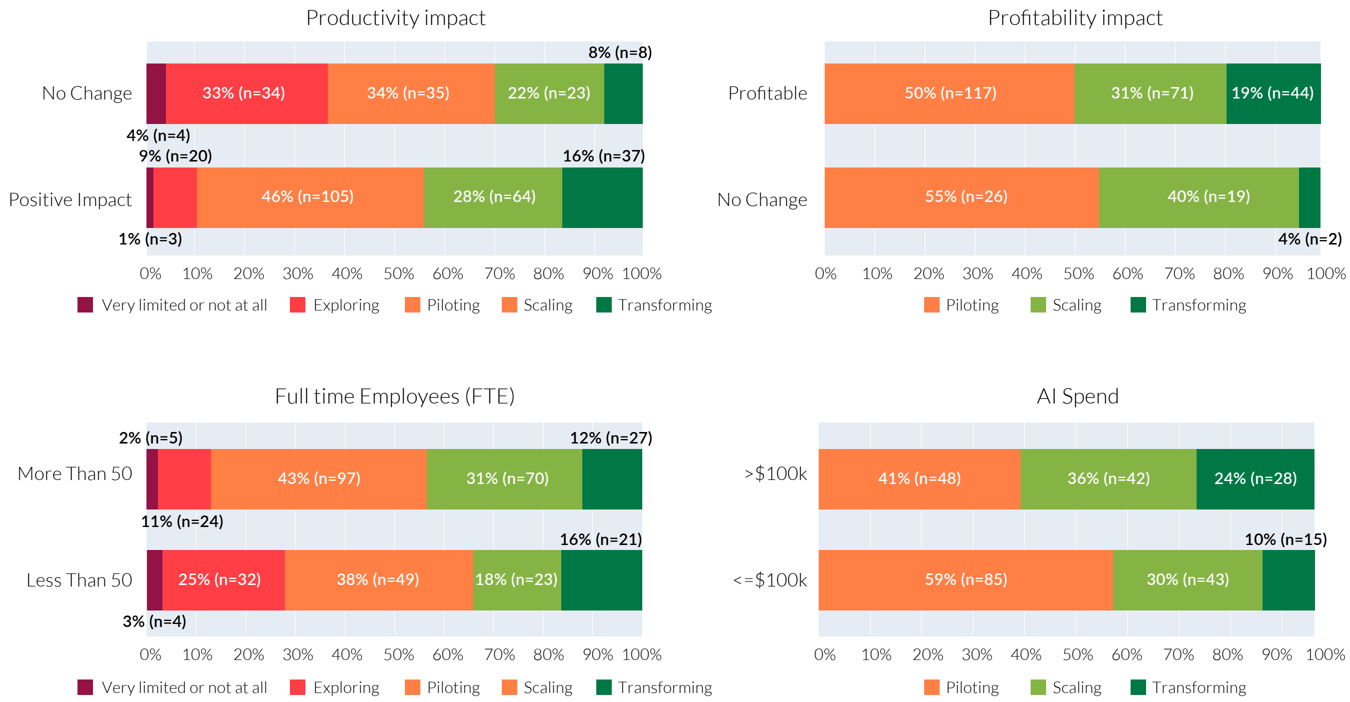

Investment as a catalyst for AI adoption: AI adoption maturity is strongly linked to positive business outcomes. Among firms spending above USD 100,000 per year on AI, 24% have reached the Transforming level of adoption maturity and 36% are Scaling, with 60% of firms in advanced-stage adoption. Higher spending likely supports the data and technical talent foundations needed to transition beyond Pilot use cases and into wider organisational deployment. However, it could also point to more digitally capable firms being better positioned both to invest and to scale AI across business functions.

The AI maturity and profitability link: Higher AI maturity is related to positive profitability outcomes. Exactly half of respondents reported that AI has increased profitability are also in the most advanced stages of adoption.

However, this data must be interpreted critically. While financial returns are more likely to emerge once AI is more deeply embedded within the organisation, this finding may be linked to self-reporting bias. Organisations – and especially the individuals championing these initiatives that were likely to be the people completing this survey – that have invested heavily in AI, may have strong psychological and strategic incentives to report positive financial outcomes.

Agentic AI adoption: Agentic AI, defined as autonomous agents capable of planning, tool use and iterative reasoning, has emerged as one of the clearest frontier technologies in financial services. Industry-wide adoption is accelerating rapidly. External reports, such as NVIDIA's 2026 State of AI in Financial Services, found that 42% of financial institutions are already using or assessing agentic AI, with 21% having deployed AI agents into production.[4]

52% of financial industry respondents are piloting or are at more advanced deployment stages with agentic AI, confirming movement well beyond basic exploration. The question is no longer whether institutions are considering it, but where adoption is concentrated and what differentiates the institutions moving fastest.

The figure below presents agentic AI adoption across six paired comparisons based on the data cuts outlined in the research methodology, each isolating a specific dimension of the data.

Figure 1.6: Industry agentic AI adoption maturity by comparison grouping

AI maturity correlates: Financial institutions with more mature levels of AI adoption report greater use of agentic AI at 34% (transforming or scaling), compared with less mature institutions (17% at transforming or scaling). This 17-percentage-point gap suggests that AI agents may not be an entry-level technology but rather a secondary stage of implementation that relies on pre-existing data infrastructure, internal resources, and governance frameworks.

The profitability link: 65% of firms that cite increased levels of profitability from AI are Piloting or deploying agentic AI (39% Deployed, 26% Piloting), compared with just 18% (deployed) reporting no change to profit.

Fintechs, geography and scale: Fintechs, characterised as early tech adopters, are outpacing their traditional FI peers in piloting agentic AI (33% versus 24%) and full deployment (25% vs 23%). Interestingly, financial institutions based in EMDEs had slightly higher levels of agentic AI adoption than AEs (55% versus 51%). This may be due to sampling bias with a higher proportion of respondent fintech firms operating in EMDEs. Notably, larger organisations (≥50 FTE) show higher adoption rates (57%) compared with smaller firms (47%), potentially indicating that the resource-intensive nature of developing and managing autonomous agentic systems favour institutions with greater resource depth.

The regulatory blind spot: A significant institutional divergence between the private sector and regulatory bodies is evident in the adoption of AI agents. Regulators have far lower levels of agentic AI adoption compared with industry (28% overall) compared with the industry average of 52%. This 24 pp gap points to a potential emergent supervisory risk area: an increasingly autonomous financial services industry outpacing regulatory adoption speed of agentic oversight tools, which could create unpredictable and volatile market outcomes.

Partner perspectives:
The IMF – Macroeconomic impacts of AI adoption

By Miguel Segoviano, IMF

Insights and Key Issues: The IMF notes CCAF survey collects data with informative heterogeneity across respondents—including traditional financial institutions, FinTech firms, AI vendors, and regulators across multiple countries—in terms of AI adoption timing and intensity, technology choices, deployment status, and impacts on productivity and profitability. It also highlights differences in scalability challenges, perceived risks, market expectations, and data privacy concerns. A key outstanding issue is the limited time dimension of the data collected.

Strategic Alignment: The IMF is advancing an institution-wide effort to assess the macroeconomic impact of AI adoption. In this context, the IMF Statistics Department (STA) is developing data frameworks to identify the most relevant variables and build indicators to support this broader IMF institutional effort. Against this backdrop, the development of CCAF's database could offer STA a valuable opportunity to leverage these frameworks to assess the relevance of CCAF's data.

Future Directions: CCAF survey data may feed into proposed STA's data frameworks to test variables that may be relevant for the IMF's assessment of the macroeconomic impact of AI adoption. Some ideas by STA include using the CCAF database to develop and test measures of financial industry AI adoption, and asymmetric AI adoption, quantifying gaps between traditional financial institutions and fintech firms. These measures could serve as valuable inputs to assess how such differences affect productivity, profitability, risk-taking, and implications for financial inclusion and financial health.

In summary, this chapter has shown that AI adoption in financial services is split into two distinct paths. A significant minority of institutions are operating at scale or in active transformation, while the majority remain at exploration and piloting stages. That gap runs along consistent lines: between fintechs and traditional FIs, between advanced and emerging market contexts and between the financial sector and the regulatory apparatus that oversees it, raising questions on whether the adoption speed is outpacing institutional readiness to mitigate or control correlated risks. The next chapter provides detail on the specific application areas underpinning adoption, providing insights on the divide between using AI for automation and cost-cutting as opposed to actively generating new revenue potential.

Chapter 2

Industry AI Adoption

The AI applications deployed by financial institutions follow a consistent pattern: they are heavily concentrated in functions where the business case is established, and the risk of error is bounded. This chapter examines that deployment map across front office, back office and risk and compliance functions, with attention to how patterns vary by firm type, macroeconomic development context, and overall AI maturity stage.

Key highlights

- **Software development and long-established fraud detection methods**: Currently the most mature and widely deployed AI use cases globally.
- **Fintech versus traditional FIs**: AI is more widely used to reshape products and platforms in fintechs, while mostly used to augment established expertise in traditional FIs.
- **Back-office gap**: More mature AI adopters trend towards applications within internal and back-office operations (HR, legal and data management).
- **Geographic revenue generation divide**: Front office value-creating use cases show the largest divide between advanced economies (AEs) and emerging markets and developing economies (EMDEs).
- **Future deployment expectations**: Firms expect the strongest growth over the next three years in AI deployment across data, legal operations and HR use cases.

Financial services activities typically flow through three functional areas:

1. **Front office**: Where client interaction and commercial decisions take place.
2. **Back office**: For execution, operational and analysis-based activities.
3. **Risk and compliance**: Continuous oversight by risk and compliance functions.

This structure provides a simple framework for analysing AI use cases across the financial industry. Across all these functions, the highest levels of AI adoption are concentrated in operational, back-office activities (software development, data visualisation, process automation and data management), followed by risk-related applications (fraud detection) and structured front office tasks (customer support).

Crucially, except for software development and fraud detection, which show relatively higher rates of full deployment, all other listed use cases remain predominantly in Piloting or Development stages.

At the less mature end of the spectrum, environmental and climate risk applications represent the least adopted use case in the sample, (9.1% piloting, 3.4% deployed). This pattern is consistent with external research that shows AI adoption in climate-related financial applications remains at an early stage due to regulatory uncertainty, significant data and infrastructure requirements, and ongoing concerns around AI model reliability (RAND, 2025).[5]

The landscape of AI use cases across the financial services industry

Figure 2.0: Industry AI adoption maturity across use cases

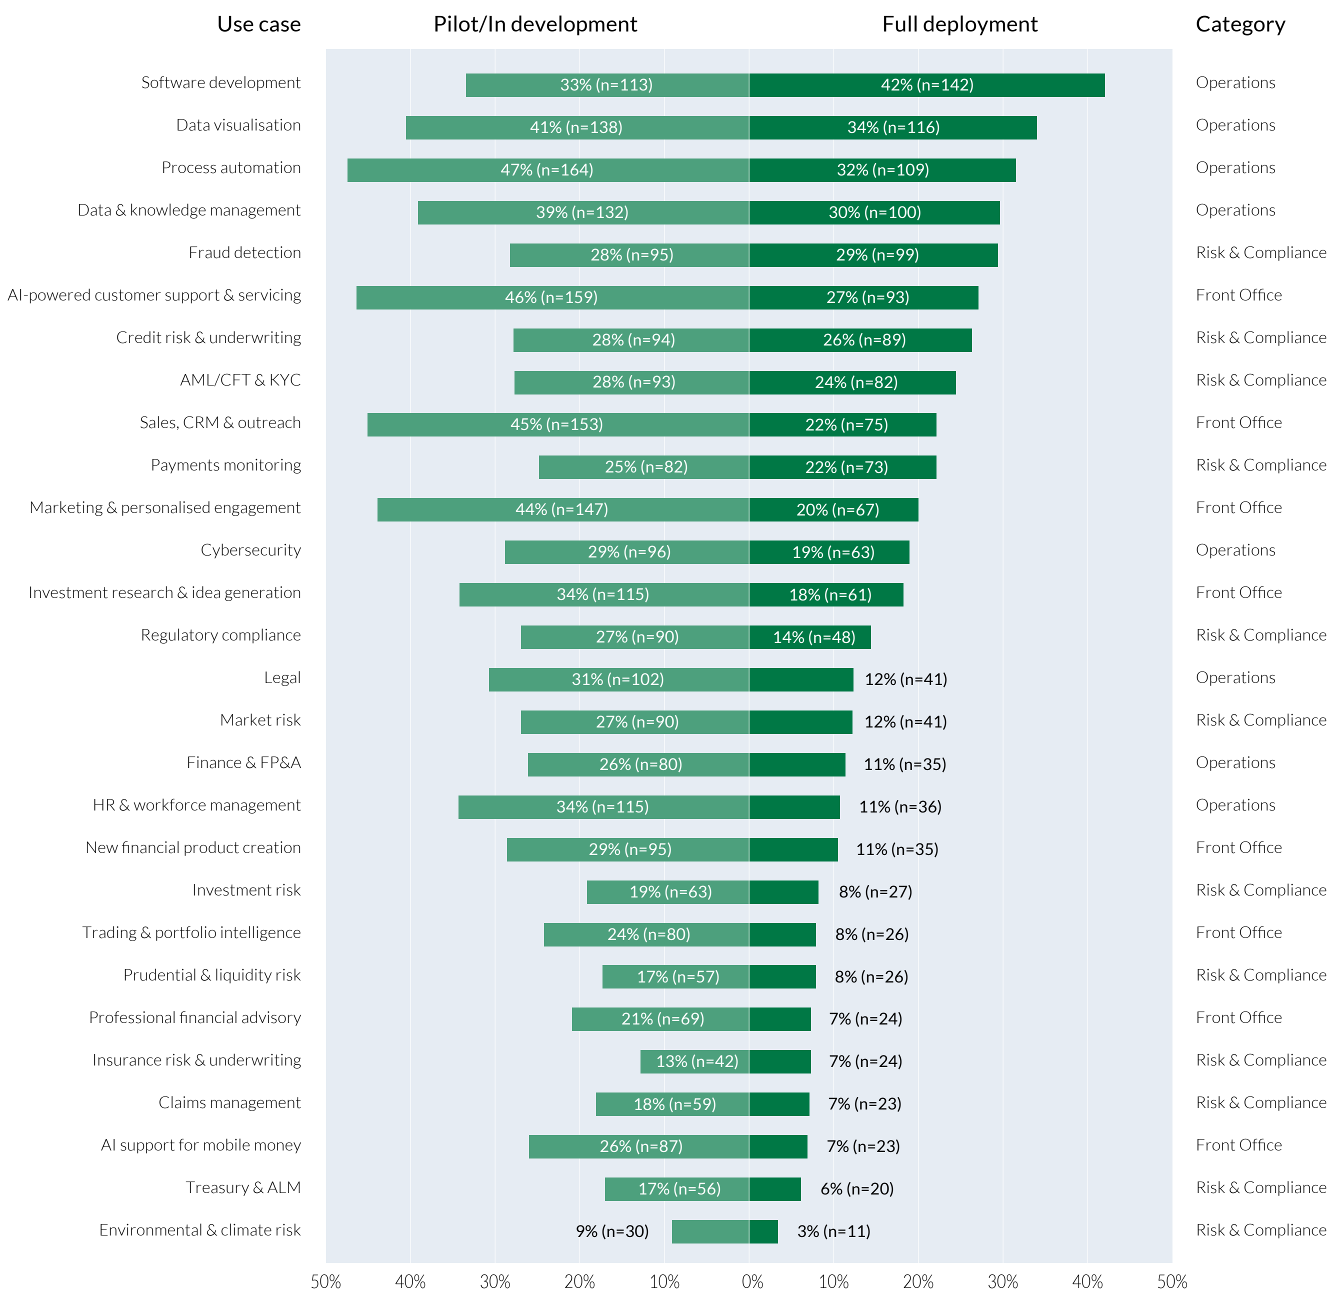

Analysis by function

Front office: Within the front office, AI-powered customer support (73%), sales, CRM and outreach (67%), and marketing and personalisation (64%) are the most widely adopted front office use cases. These applications primarily support client relationship management and customer acquisition. Current adoption remains lower in revenue-generating front office activities such as trading and advisory services and new product creation.

Back office: In back office and operations, AI adoption is highest in process automation (79%), data visualisation (75%) and software development (75%). By contrast, HR, legal operations, and finance, including financial planning and analysis (FP&A), show the lowest level of adoption. Firms reporting AI use in these lower-adoption back-office functions tended to report higher profitability from AI applications as detailed in the chart below.

Risk and compliance: A similar pattern is observed in risk and compliance. While fraud detection (57%), credit risk and underwriting (54%), and AML/CFT and KYC (52%) are the most widely adopted use cases, firms reporting AI adoption in less prevalent areas such as treasury, insurance risk, claims management and climate risk reported higher profitability outcomes. This trend shows that higher reported profitability is more commonly observed among firms applying AI in more specialised use cases.

Correlations with profitability: An interesting pattern emerges regarding firms reporting increases in profitability from AI. While adoption of customer support and process automation were widely reported, firms that deployed less common, highly specialised applications – within new product creation, trading and portfolio intelligence, treasury management and FP&A – also report higher overall increases in profitability from their organisation's overall use of AI. While this does not demonstrate a causal relationship, it suggests that firms that are able to embed AI applications into more complex, domain-specific workflows also extract higher financial premium than those firms using AI for more generic automation tasks.

Figure 2.1: Use case adoption rate versus increased profitability

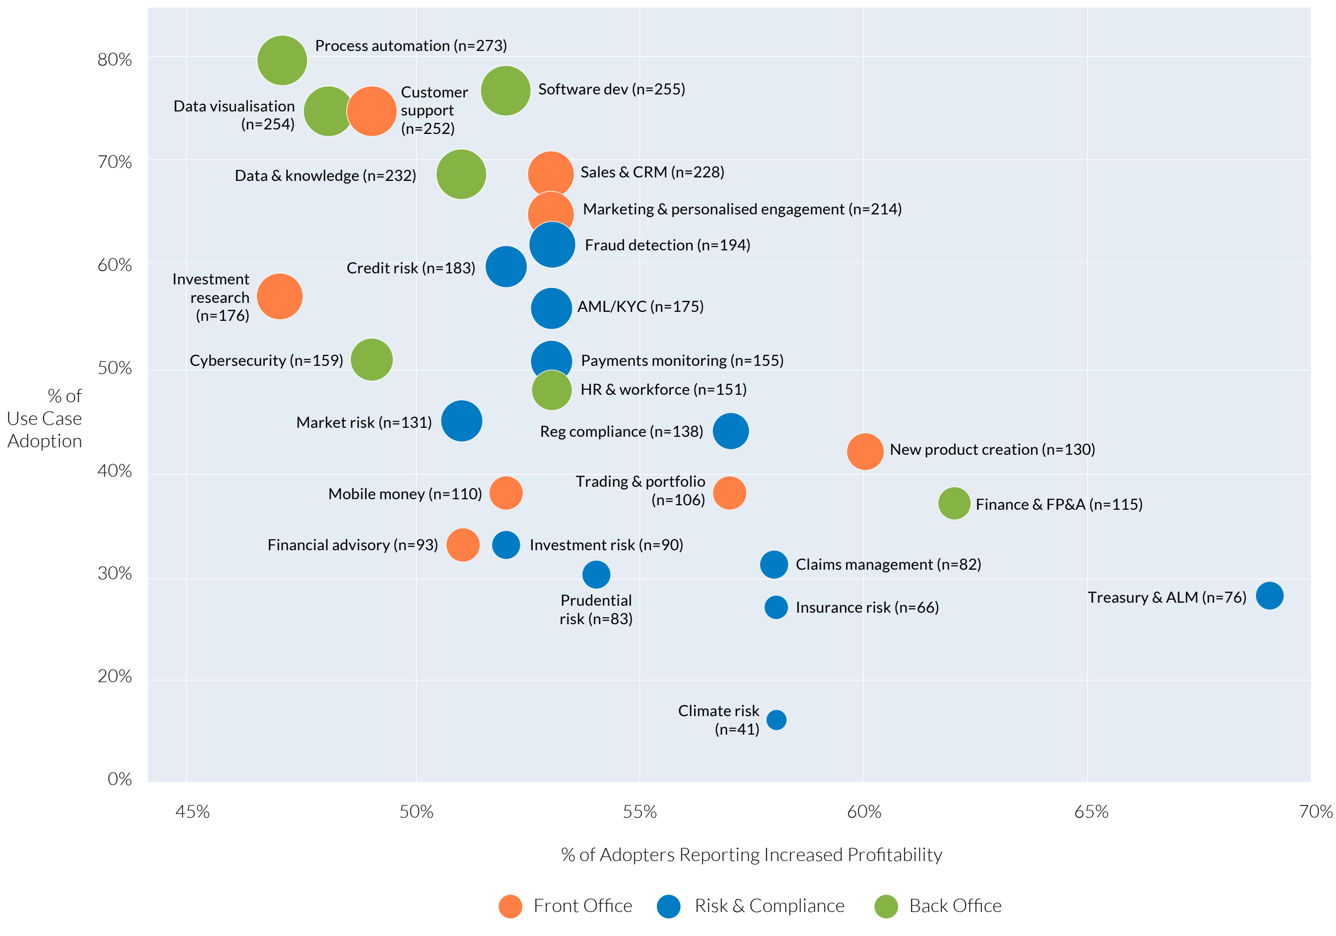

Fintech versus traditional FIs

Fintechs and traditional FIs show similarities in AI adoption, with software development (84% for fintechs; 72% for traditional FIs) and process automation (82% for fintechs; 78% for traditional FIs) ranking among the most widely adopted use cases for both groups.

The divergence becomes evident beyond these common foundations, as fintechs have greater rates of AI deployment within customer delivery, product innovation, and decision-making activities such as mobile money, payment monitoring, and new product creation.

Traditional FIs report modest leads in trading and portfolio intelligence and professional advisory services.

Figure 2.2: AI use case adoption gaps – Fintechs versus traditional FIs

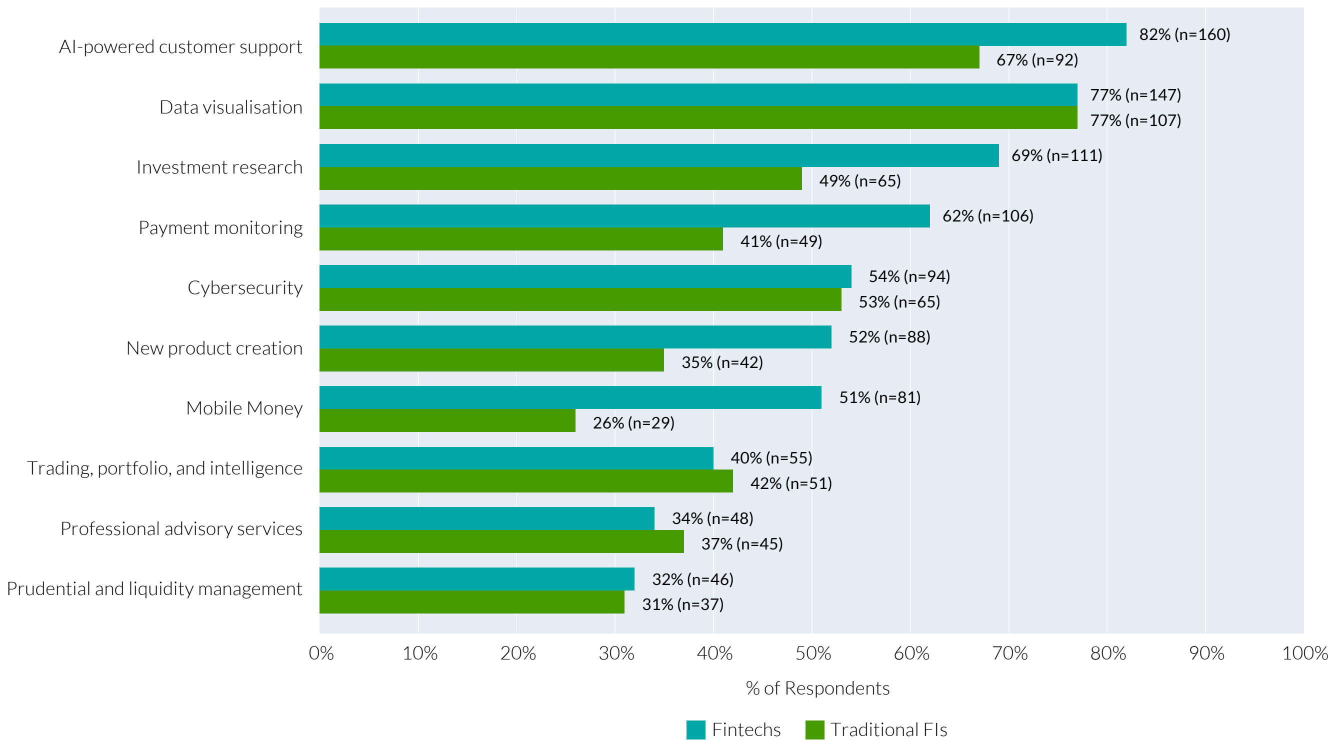

Economic Development

The widest adoption gaps between AEs and EMDEs can be seen predominantly in front office value-creating use cases, with four of the five largest differences relating to investment research, new product creation, advisory services and trading-related intelligence.

Figure 2.3: Use case adoption by economic development

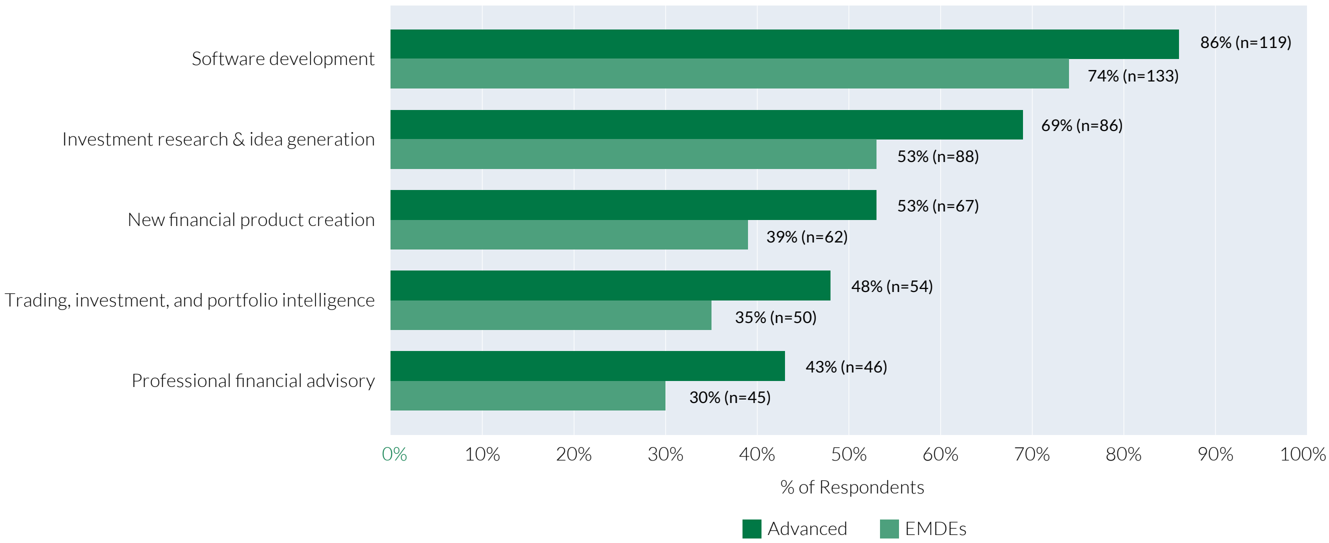

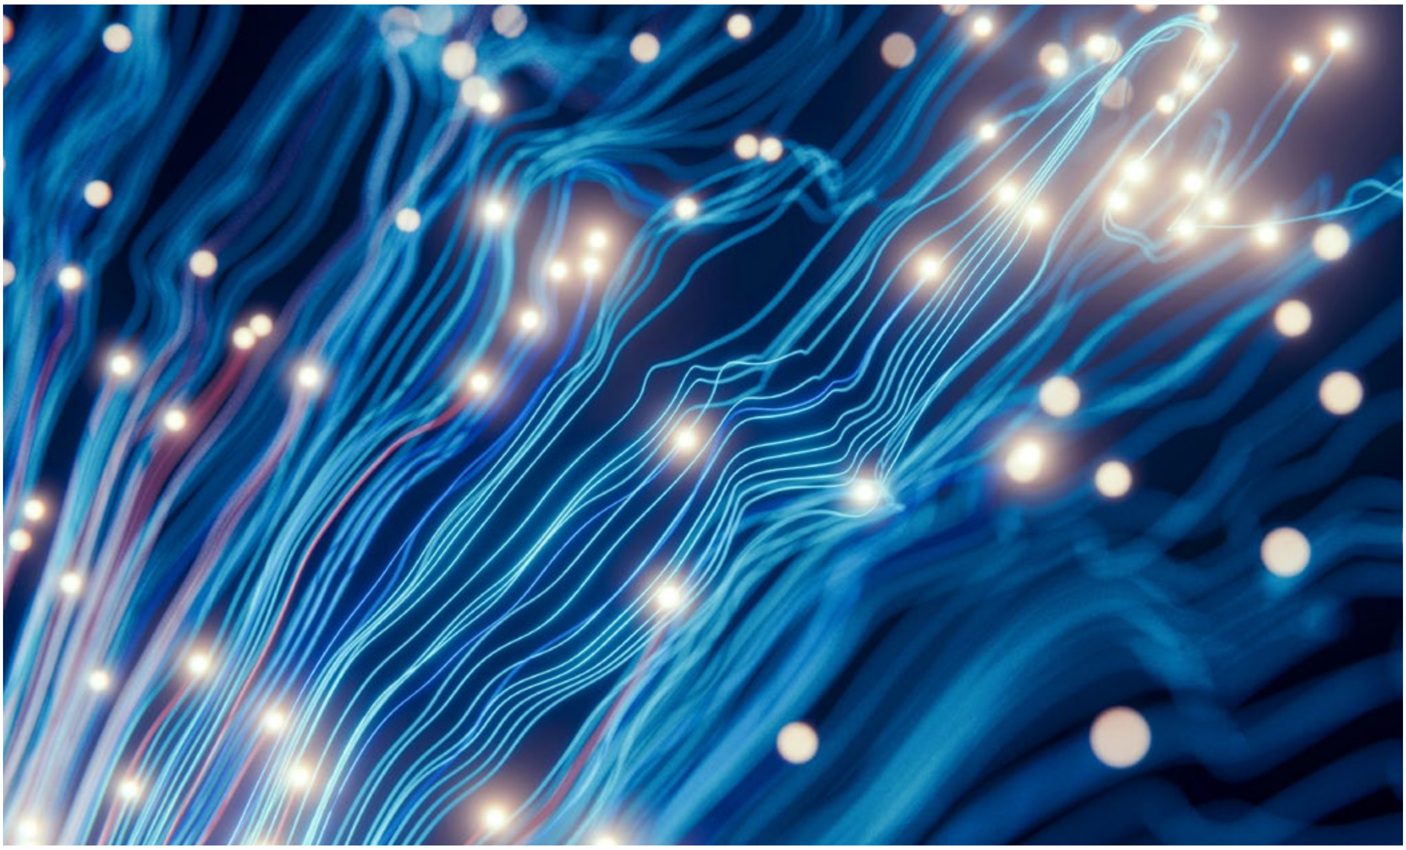

Maturity

The study further explores differences in use case adoption between AI-mature (trasnforming or scaling their institutional use of AI) and less mature (exploring or piloting their institutional use of AI) respondents. Nearly all AI-mature firms report using AI for software development and for data management and visualisation. The largest adoption gaps are mainly seen in internal operations and back-office functions, which account for four of the ten widest differences observed. This indicates that the main distinction between more and less mature firms lies in how deeply AI is embedded in operational activities.

Figure 2.4: AI use case adoption gaps by AI maturity

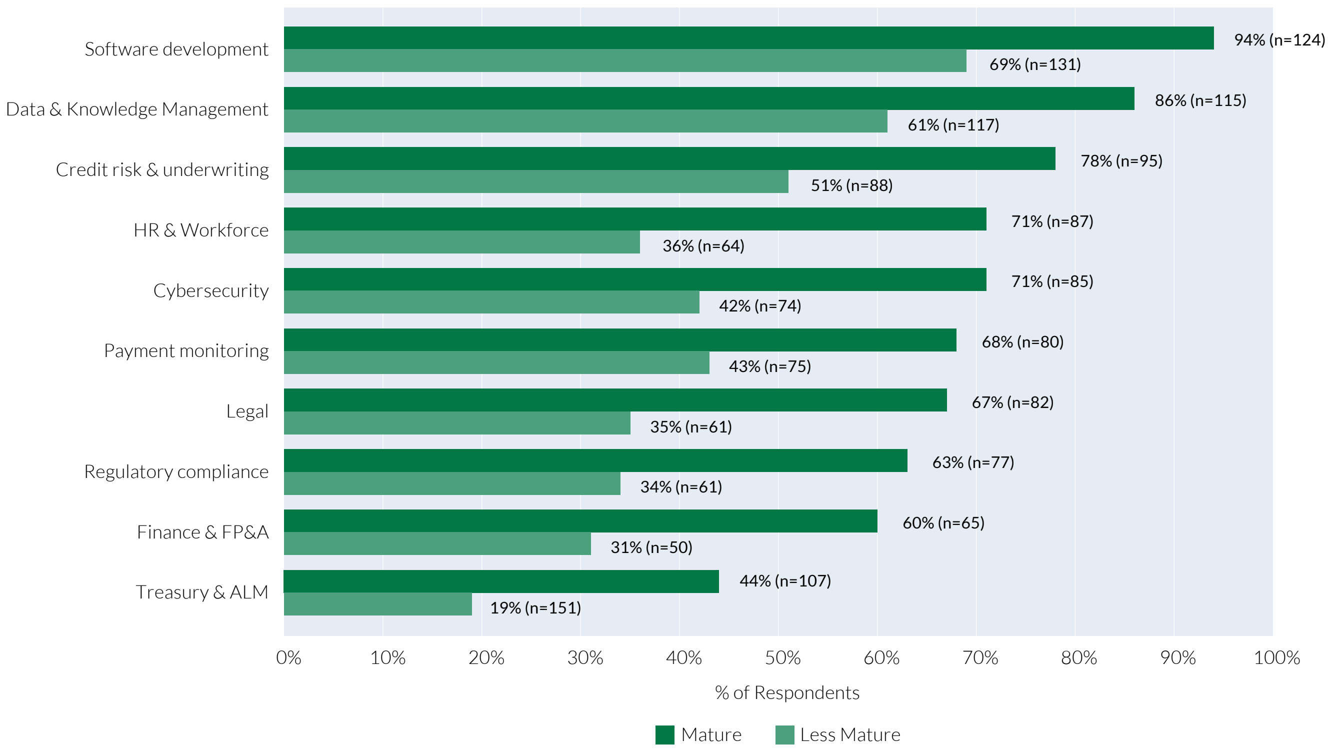

HR and workforce show the single largest gap between the two groups. AI use in HR is classified as a high-risk application under the EU AI Act (2024) and is commonly associated with heightened regulatory scrutiny, concerns around bias, and challenges integrating AI into legacy systems. These characteristics could make adoption more feasible for AI-mature firms with more advanced governance frameworks, controls and implementation experience than for less mature firms.[6]

Three-year deployment expectations

The survey also examined which AI use cases are expected to see greater deployment over the next three years. While some firms have already fully deployed these applications, many others continue to report them at earlier pilot or developmental stages.

Respondents most frequently point to data and knowledge management, legal, and HR and workforce applications as areas where deployment is expected to expand. These expectations may be guided by recent advances in large language models and related generative AI capabilities, particularly in search, summarisation, extraction and drafting. Such capabilities align well with data, legal and HR functions, which depend heavily on unstructured text, document-intensive workflows, and proprietary internal information.

Figure 2.5: AI use case – pilot/development versus expected to be deployed in three years

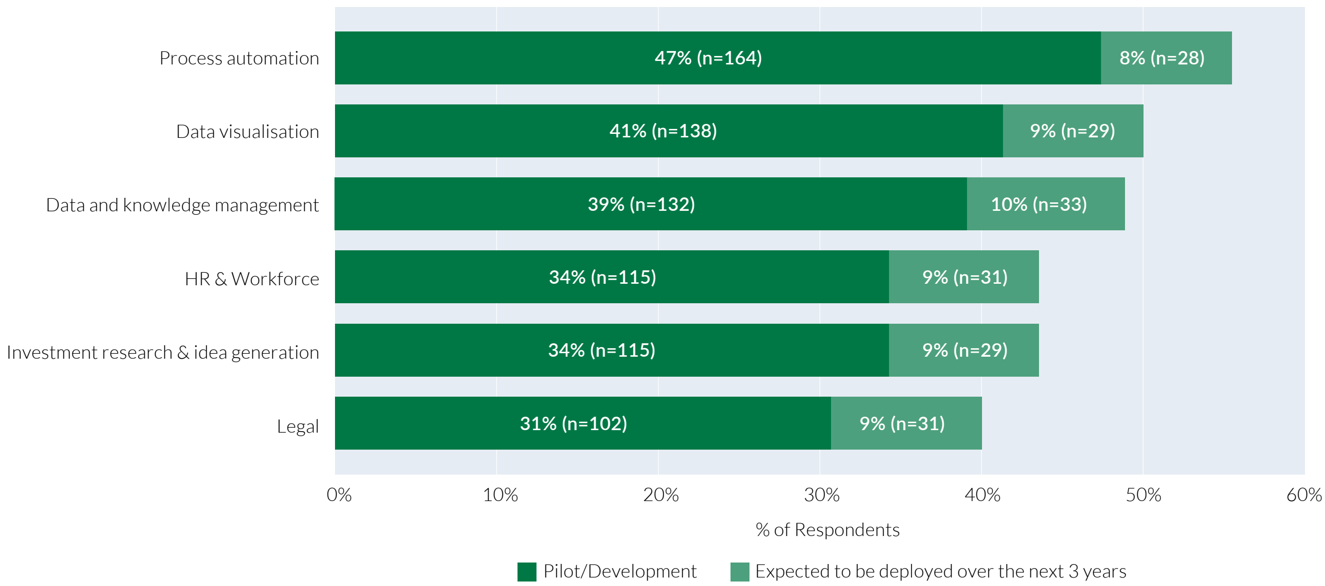

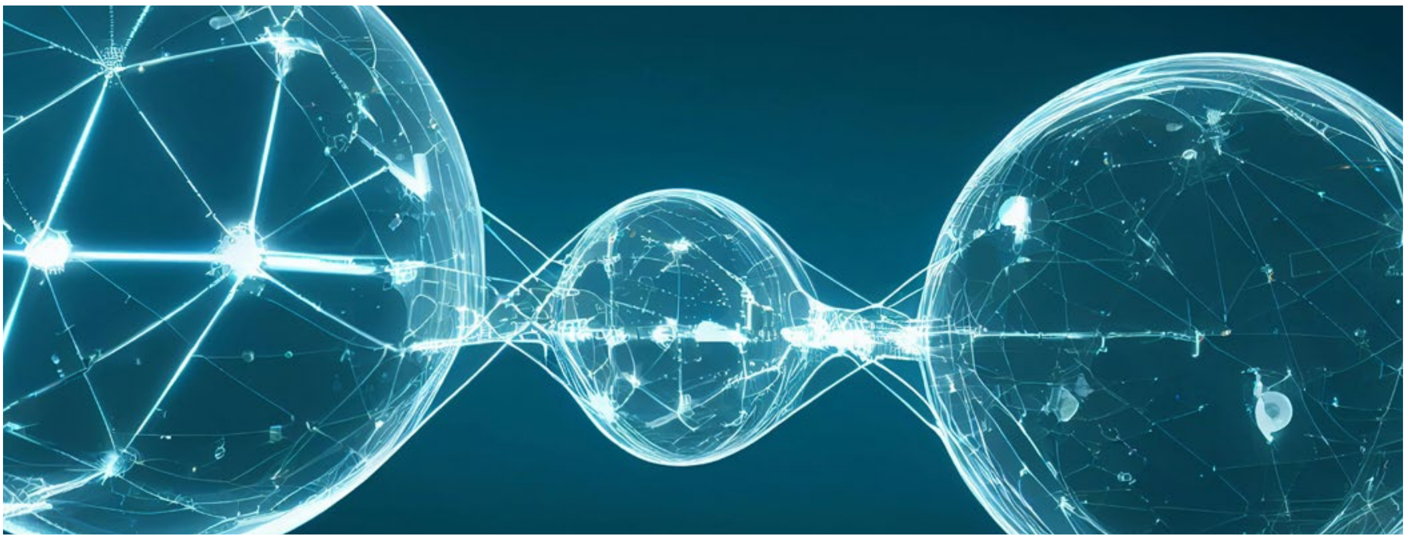

Partner perspectives: The AMF – AI adoption and implications in the Arab world

By Nouran Youssef, Karim Mouaffak

AI is now embedded in financial services across the Arab region, and its adoption is no longer confined to isolated experimentation. Financial institutions are actively embedding AI into operational processes, risk management functions, and customer-facing activities. This progression is taking place within systems that remain structurally diverse, both in terms of institutional capacity and technological readiness.

Digital Infrastructure, Data, and Financial Inclusion

Digital infrastructure differentiates AI deployment across the region. AI systems depend fundamentally on data, while the availability of high-quality and structured data is closely linked to the maturity of digital public infrastructure.

In several Arab countries, significant progress has been made in areas such as digital identity, real-time payment systems, and open banking frameworks. These developments provide a foundation for more advanced AI applications, particularly in credit assessment, fraud detection, and customised financial services.

In other countries of the region, digital identity systems remain fragmented, interoperability is limited, and data ecosystems are still developing. This variation directly affects the ability of financial institutions to deploy AI effectively. Where data is incomplete or inconsistent, model performance is constrained, and risks related to bias and exclusion become more pronounced.

This is particularly relevant when considering the role of AI in financial inclusion, where the potential benefits are significant. For instance, AI-driven models can extend access to finance by incorporating alternative data sources and reducing reliance on traditional credit histories. This is especially important for SMEs and informal sectors, which are often underserved by conventional financial systems. The global data confirm that these segments are a primary focus, with SMEs representing around 70% of target users and retail consumers approximately 64%.

However, these benefits are conditional. The same systems that enable access can also reinforce exclusion if underlying data is biased or incomplete. If algorithmic decisions are not properly governed, they may replicate existing inequalities in new forms. Therefore, ensuring that AI contributes positively to inclusion, requires technological capability and robust governance frameworks, including mechanisms for transparency, accountability, and fairness.

Structural Heterogeneity across the Region

The Arab financial system is heterogeneous. The region includes economies that are actively investing in AI and digital infrastructure on a scale, alongside others where foundational capabilities are still being developed. This divergence is consistent with global trends, where approximately 51% of institutions in advanced economies report scaling or transforming through AI, compared to around 32% in emerging markets.

In practical terms, this means that priorities differ significantly across the region. In more advanced jurisdictions, the focus is on scaling AI capabilities, developing sovereign infrastructure, and positioning the financial sector within broader technological strategies. In others, the priority remains on building foundational digital infrastructure, strengthening regulatory frameworks, and developing institutional capacity.

This divergence has direct implications for policy design. A single regulatory model applied uniformly across the region would fail to account for these differences. A more effective approach lies in combining regional coordination with national flexibility. Jurisdictions with more advanced capabilities can contribute to the development of standards and regulatory practices, while others focus on building the foundations required for gradual integration.

The trajectory of AI in the Arab financial system will depend on how effectively technological adoption, supervisory capacity, and infrastructure development are aligned. Where alignment is achieved, AI can support financial stability, inclusion, and resilience.

In summary, the deployment map documented in this chapter is also a map of institutional priorities. The data shows that financial institutions have invested most heavily in AI where the business case is established and where the consequences of failure are bounded, such as back office automation and well-understood risk functions.

The frontier applications where competitive differentiation is most likely to emerge, including agentic deployment, new financial product creation and customer-facing personalisation, remain the territory of more advanced institutions. Fintechs consistently reach further into that territory than their traditional counterparts. The following chapter dives deeper into the technological lifecycle by examining the infrastructure used to develop and deploy AI in the surveyed institutions.

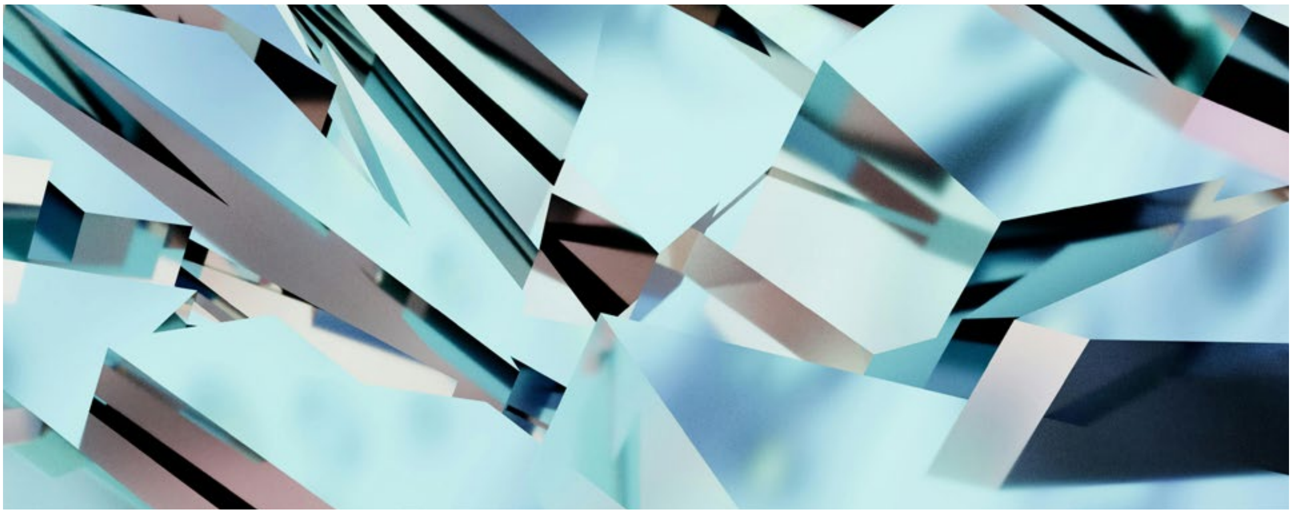

Chapter 3

Technology and Investment Dynamics

The adoption of AI in financial services is shaped not only by strategic intent but by the practical realities of budgets, procurement timelines, workforce capacities and technical infrastructure.

The 2020 CCAF-WEF AI Report documented an industry where experimentation was widespread, but scaled deployment remained limited. Five years on, the technological foundations have transformed paradigms. With global spending on AI in banking alone expected to reach USD 97 billion by 2027 (World Economic Forum),[7] this chapter examines the core technological drivers of that investment: how firms are spending, how quickly they are deploying, the severe constraints of the talent market and growing concentration risks around critical cloud and foundation model providers.

An AI investment divide

The distribution of annual AI expenditure across industry respondents reveals a marked divergence between fintechs and traditional FIs. Of the 352 industry respondents, 234 reported their annual AI expenditure in USD (149 fintechs and 85 traditional FIs).

Figure 3.0: Annual AI spend by firm type (in USD) – fintechs (n=149) versus traditional FIs (n=85)

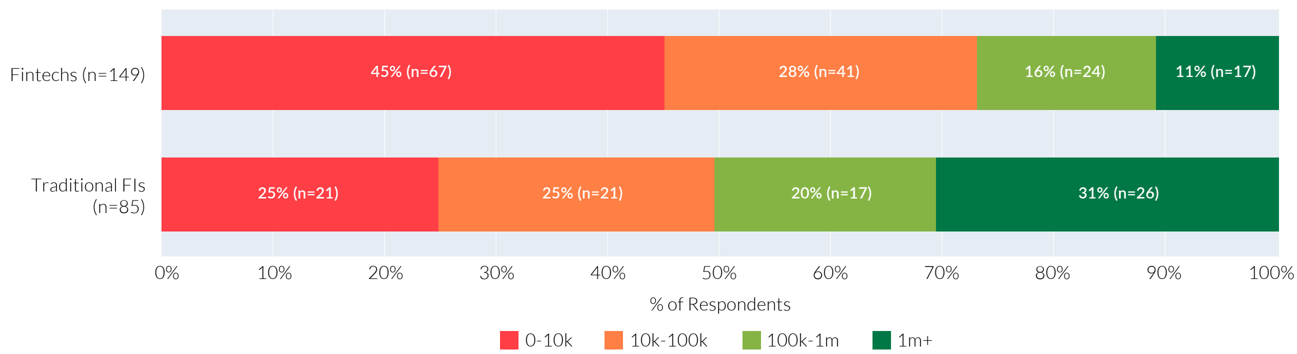

45% of fintechs reported annual AI spend under USD 10,000, and over 70% spent less than USD 100,000 annually. Despite these lean budgets, fintechs maintain the highest rates of reported AI adoption maturity (as documented earlier). This suggests the prevalent pairing of efficient, cloud-based architectures with easier-to-adopt, off-the-shelf GenAI and API-driven foundation models.

Meanwhile, traditional FIs surveyed report larger capital investment, with 31% spending over USD 1 million on AI annually. This makes sense as large AI vendors typically charge higher fees for higher numbers of users. Larger firms also face the realities of integrating AI within long established legacy technology estates, and the need to navigate more complex security requirements and funding large scale change management programmes which necessitate higher capital investment.

Speed to deployment

Two-thirds of AI vendors (66%) reported deployment timelines within their clients of six months or less. Nearly a third (33%) of vendors execute within three months from initial procurement to full integration within clients.

Figure 3.1: AI procurement-to-deployment time

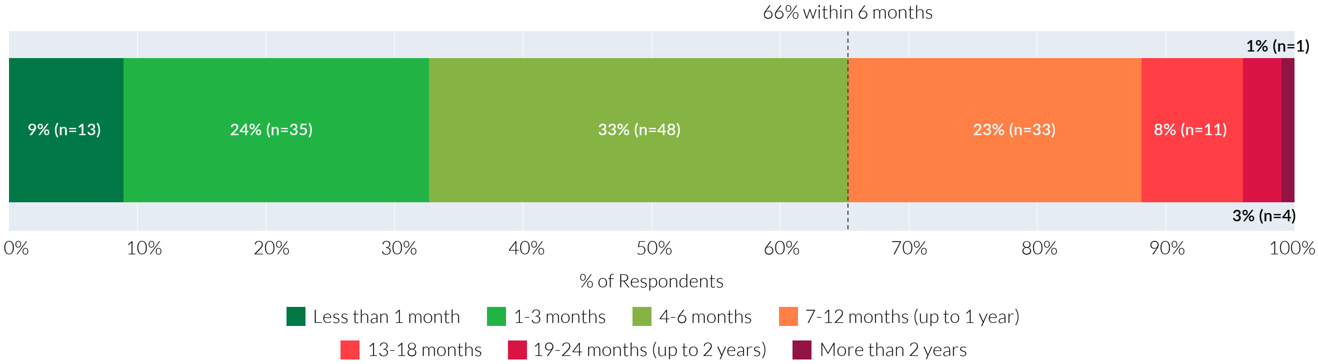

This suggests a notable acceleration since the 2020 CCAF-WEF AI report, which cited deployment friction and legacy integration as primary adoption blockers. Eleven percent of vendors report timelines exceeding one year, predominantly serving the insurance sector, or providing specialised, tailored AI solutions in highly regulated sectors.

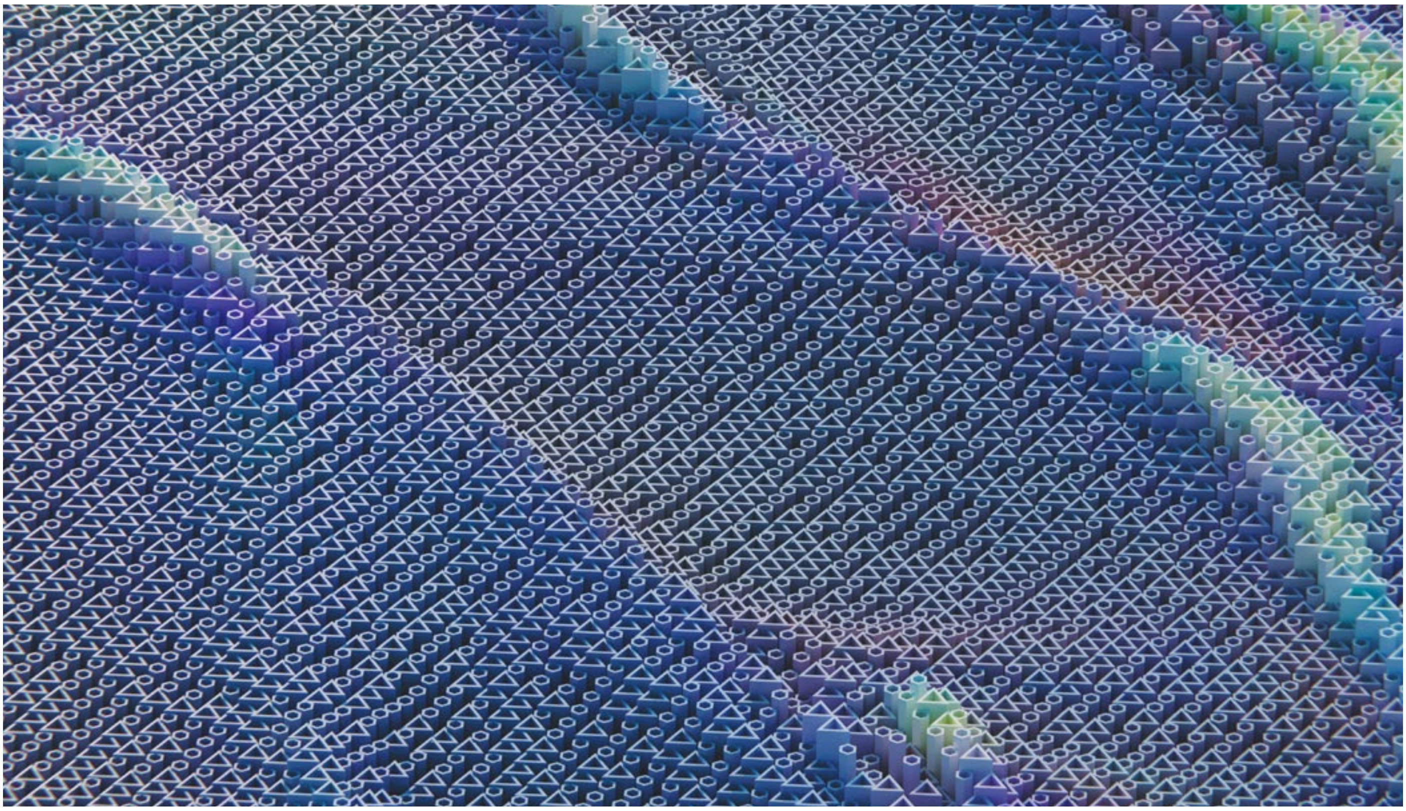

Dedicated AI technical teams

The scale of dedicated AI teams within financial institutions is largely determined by organisational size. 86% of firms with fewer than 50 employees reported five or fewer dedicated AI staff. 38% operated with just one dedicated AI staff member.

By contrast, 68% of relatively larger firms (50+ employees) have dedicated AI teams of six or more, underscoring how AI team capacity remains resource intensive. This is consistent with external reports that suggest that major AI investments act as barriers to implementation and can heighten the existing disparity between large and small firms.[8]

Figure 3.2: Number of dedicated AI staff by firm size (FTE)

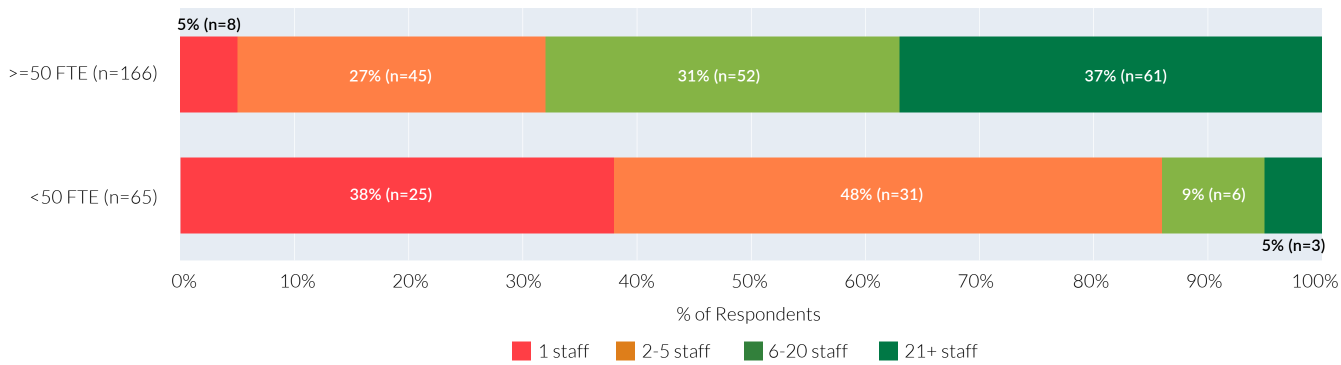

Workforce preparedness for AI

Access to talent directly impacts organisational readiness for AI. The 2020 CCAF-WEF AI report identified workforce skills as a critical barrier to AI deployment in financial services. Five years on, this finding persists. Only 10% of respondents describe their workforce as highly prepared. This capacity gap is even more acute for regulators, where 68% cited limited internal expertise as their top constraint in overseeing AI. This preparedness gap reflects the talent barriers which will be further examined in the next chapter, where access to AI expertise remains among the top two hurdles to adoption, unchanged from the 2020 study.

Critically, workforce preparedness is a leading indicator for positive financial returns from AI. Institutions that reported positive AI profitability outcomes were nearly four times more likely to describe their workforce as highly prepared (23%) compared with firms reporting no change in profitability (6%). This strongly indicates that extracting value from AI depends on both human capital investment as well as technology procurement and adoption.

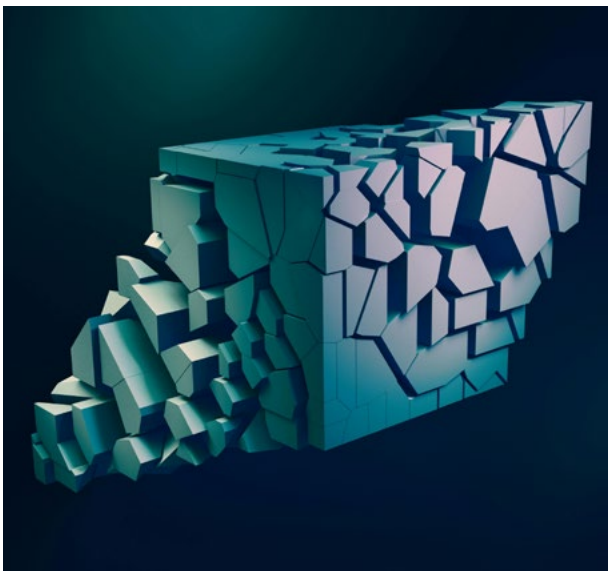

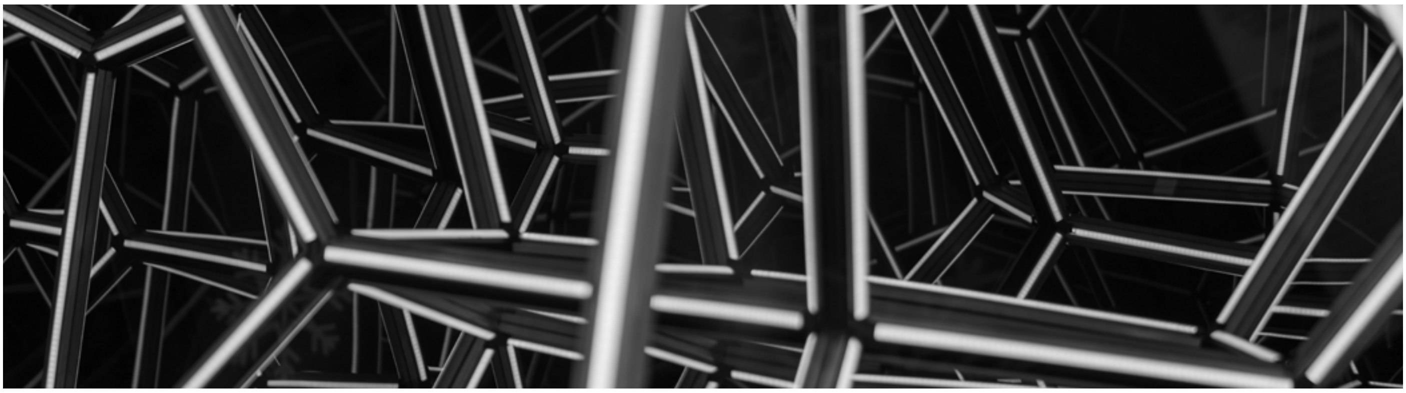

The figure below presents self-reported workforce preparedness across six binary variables for all industry respondents (both traditional FIs and fintechs) that responded to this question. It reveals a consistent pattern in which institutions that had invested more heavily in AI, or had already realised measurable returns from it, also reported higher levels of workforce readiness

Figure 3.3: Industry workforce preparedness for AI by comparison grouping

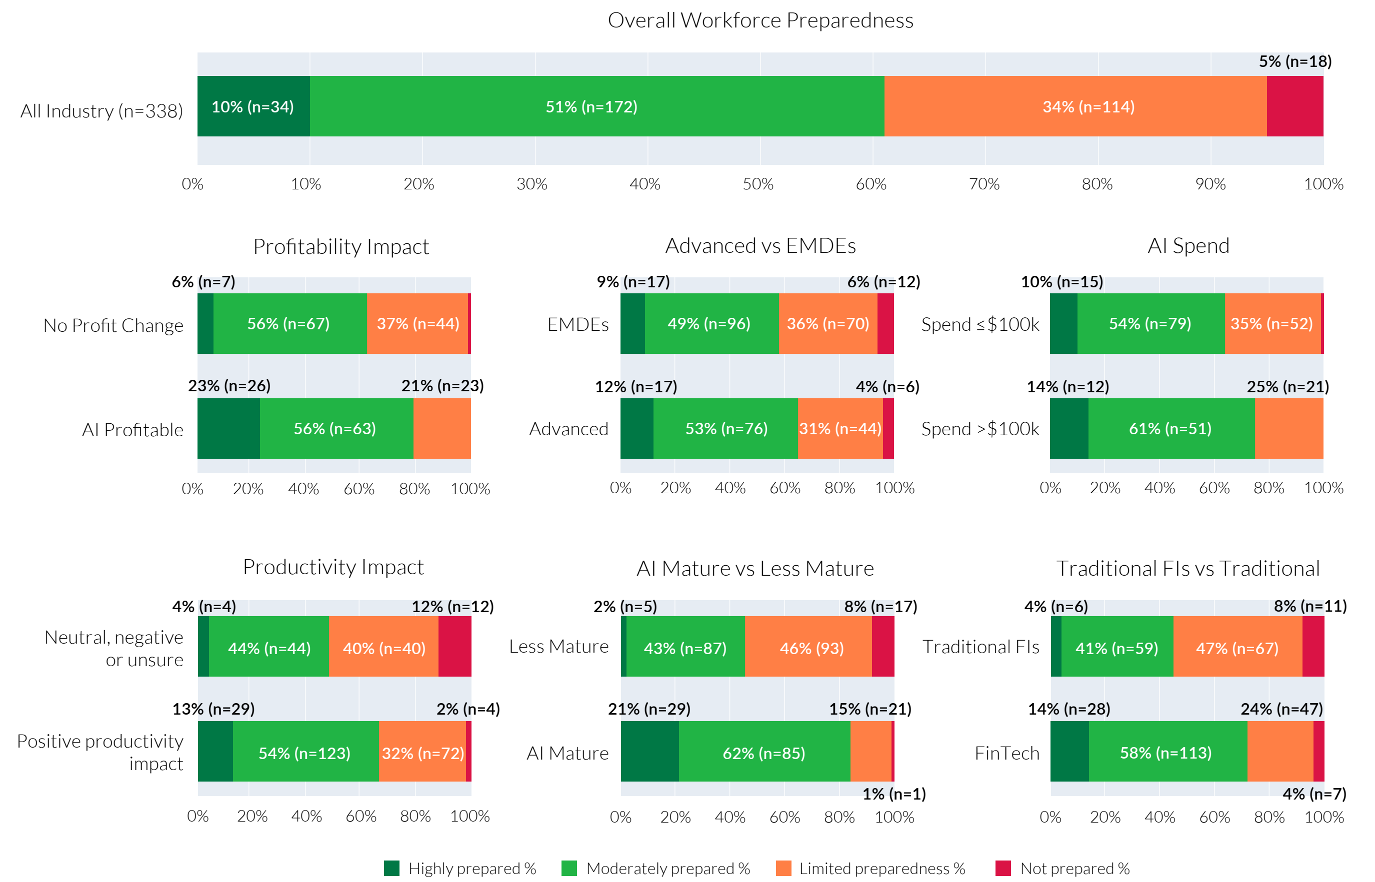

Fintechs reported higher preparedness than traditional FIs (14% versus 4% highly prepared), and advanced economies (AEs) edged ahead of emerging markets and developing economies (EMDEs), though not significantly (12% versus 9%). Higher AI spend also correlated with greater preparedness, with 75% of institutions spending more than USD 100,000 on AI infrastructure reporting high or moderate preparedness compared to 64% of those spending less.

AI tech adoption gap: fintech versus traditional FIs

The adoption gap between fintechs and traditional FIs varies significantly depending on the AI technologies used. The figure below presents the different levels of adoption across 10 AI technologies surveyed, ranked by the magnitude of the gap in the level of technology adoption.

Figure 3.4: AI technology adoption gap by firm type – fintechs (n=203) versus traditional FIs (n=149)

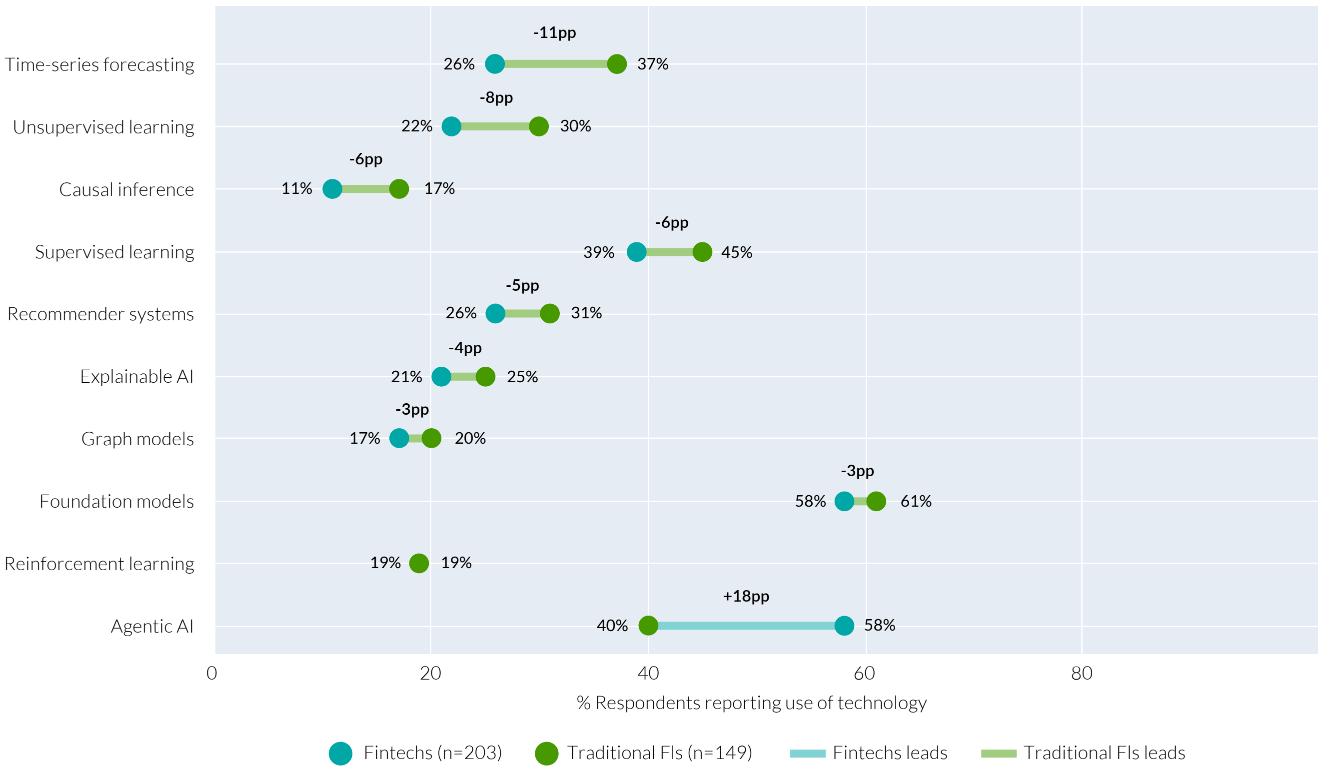

Interestingly, traditional FIs have much higher rates of adoption in two areas: time series forecasting (37% versus 26%, an 11pp lead) and unsupervised learning (30% versus 22%, 8pp), both Classical ML methods.

Foundation models show one of the narrowest gaps (3pp, 61% versus 58%), indicating that access to external foundation model providers is broadly comparable across institution types regardless of digital maturity.

The adoption-maturity paradox

The relationship between how widely a technology is adopted and how deeply its adopters have matured in their overall use of AI reveals an important paradox. The chart below plots each AI technology's adoption rate against the percentage of its adopters who have reached the Scaling or Transforming stage of deployment, with bubble size representing the number of adopters.

The trend line shows an inverse relationship: the most widely adopted AI technologies tend to have the lowest concentration of mature users. GenAI, foundation models and agentic AI are widely adopted (despite being some of the newest technologies) but by firms with lower AI maturity on average. This may be due to these technologies being easier to experiment with given lower barriers to entry.

Figure 3.5: AI technology adoption rate versus AI adoption maturity concentration in industry (Scaling/Transforming)

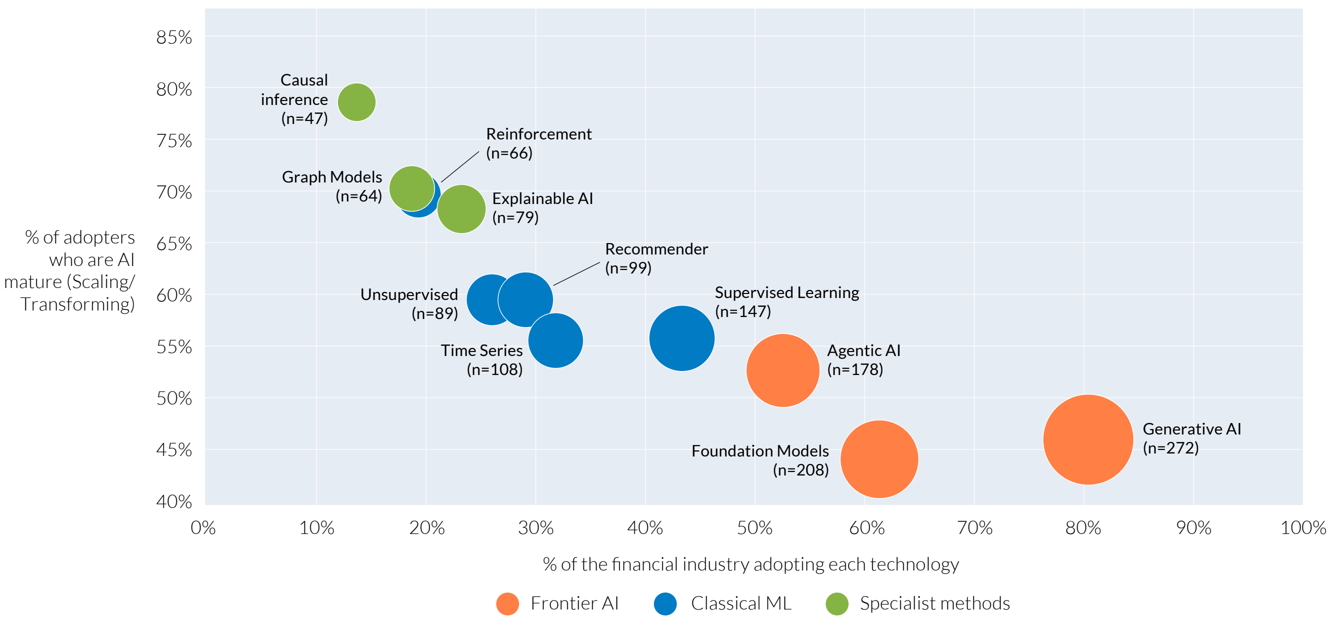

GenAI and foundation models are used by most respondents, yet a comparatively smaller share of these adopters have progressed beyond piloting. Conversely, more complex, specialist methods (for example, causal inference and graph models) have fewer adopters, but the firms using these are largely in the Scaling or Transforming stages.

This suggests a two-track financial sector: widespread, lower-barrier experimentation with newer Frontier AI and GenAI technologies, in parallel with deep, more sophisticated deployment of Classical ML and specialist methods used by a smaller cohort of firms that are seeing increased levels of positive business outcomes from the use of AI.

Value and technology: Spend versus profitability

Analysing technology by investment and profitability correlations adds to this paradox. The relationship between AI investment and profitability outcomes vary significantly by the types of AI technologies used by surveyed institutions. The figure below maps each technology by the percentage of its adopters reporting increased profitability against the percentage spending more than USD 100,000 on AI, with bubble sizes representing the number of adopters. The quadrant structure reveals distinct clusters by technology type.

Figure 3.6: AI spend and profitability by technology

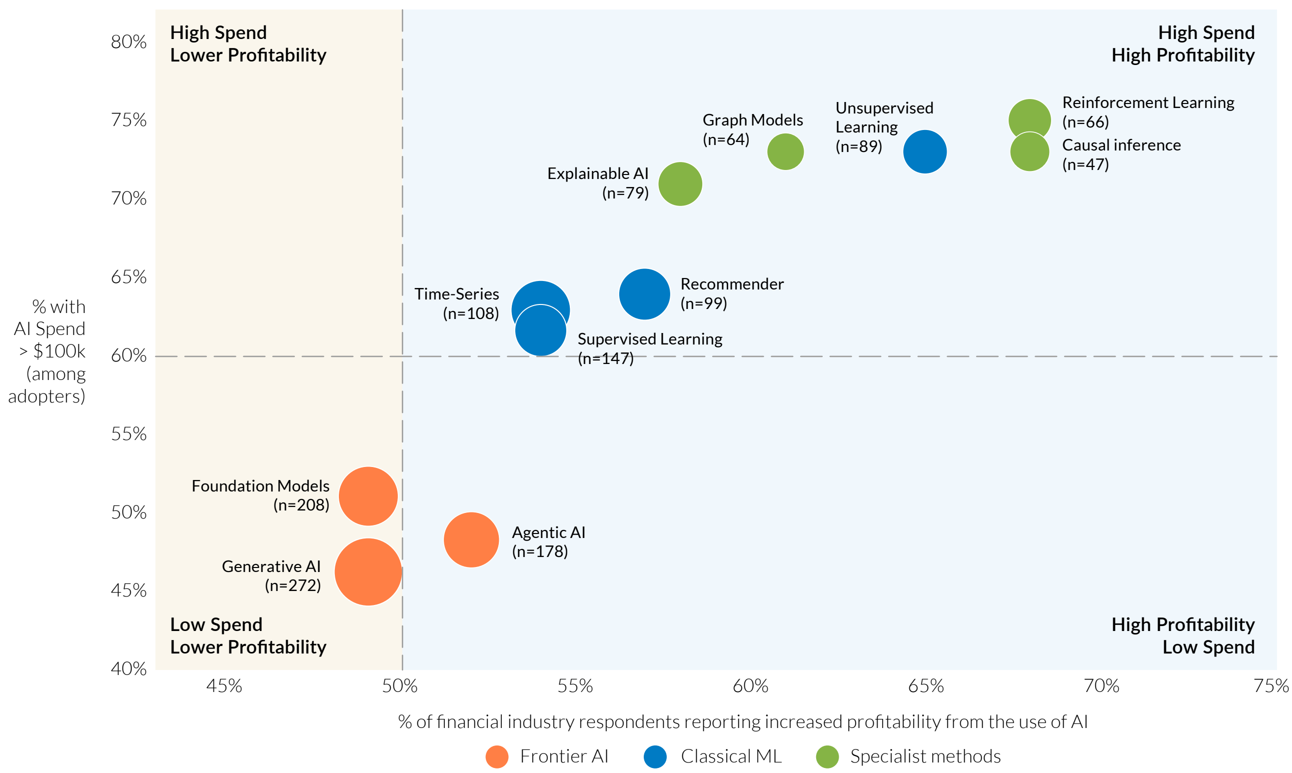

Frontier AI technologies (GenAI, foundation models, agentic AI): These technologies cluster in the lower left quadrant, characterised by generally lower spend and a lower proportion of firms reporting increased profitability among adopters. These are more accessible technologies (lower average spend) but currently correlate with a lower proportion of firms reporting increased profitability. GenAI, the most widely adopted technology (n=272), shows 49% adopters reporting increased profitability and just 46% reporting higher spend (>USD 100,000).

Specialist AI methods: Specialist tools like reinforcement learning and graph models dominate the upper-right quadrant. These likely require high levels of investment (over 70% of adopted spend >USD 100,000 annually) but correlate with more firms that see increased profitability from AI.

While GenAI democratises access to AI capabilities by lowering barriers to entry, the high financial returns in the current markets seem to be accruing to firms who are also able to deploy more bespoke, capital-intensive specialist AI technologies.

Chapter 4

AI Vendor Trends

A total of 146 AI vendors completed the survey.

The AI vendor ecosystem: generalists versus domain specialists

Figure 4.0: AI product offering by vendor type

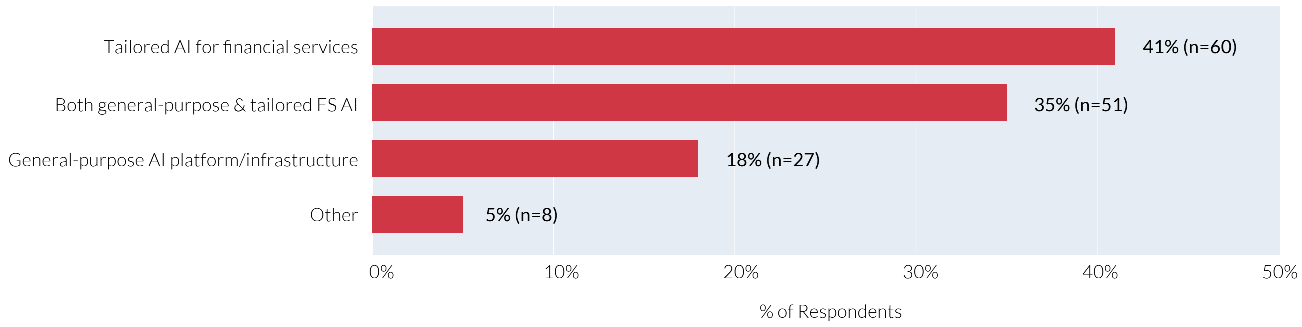

The AI vendor landscape comprises Financial Service specialists (41% of respondents) and general-purpose platforms (18%) with the remainder offering hybrid solutions.

Core tasks AI products perform

Figure 4.1: AI capabilities of financial services-specific versus general-purpose vendors

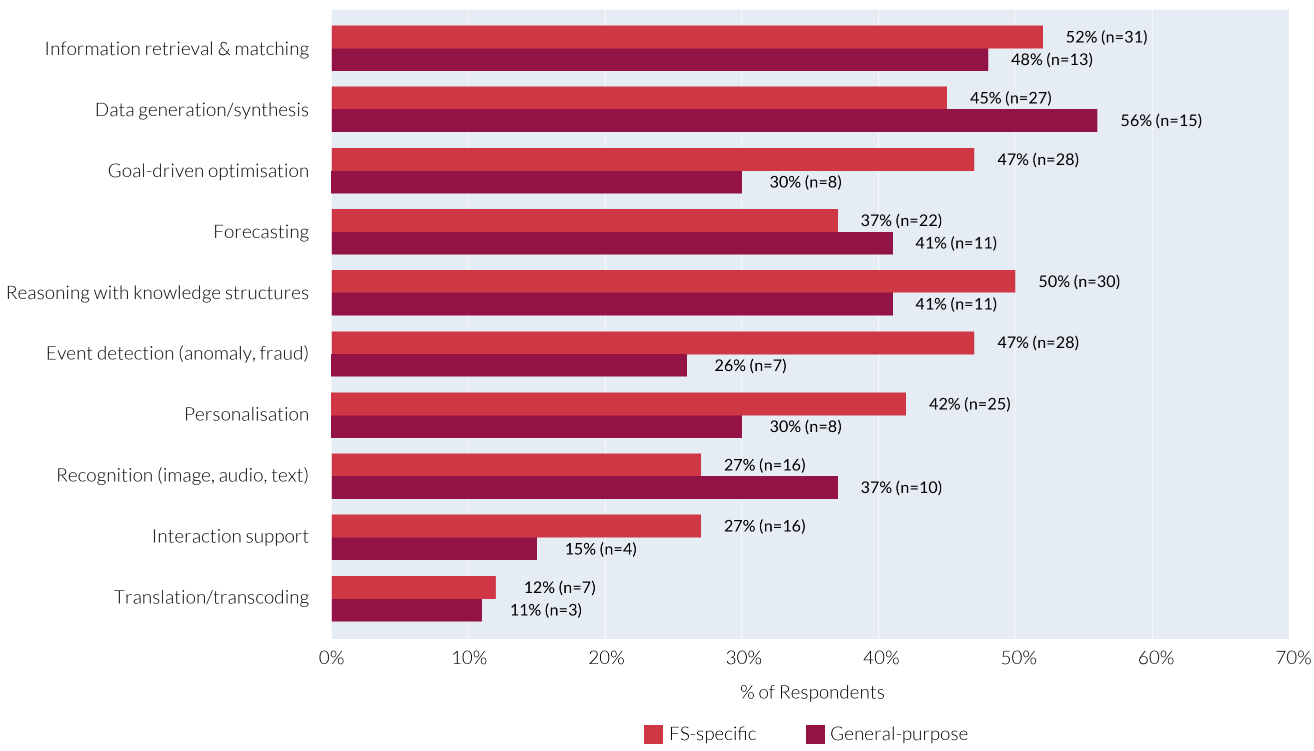

Surveyed AI vendors offered a range of capabilities defined by the OECD AI System Classification System.[9] Both groups of vendors most commonly offer information retrieval and matching and data generation/synthesis, likely reflecting the growing prominence of GenAI-era Retrieval-augmented generation (RAG), document processing and synthetic data capabilities.

The most notable divergence between general-purpose tooling and financial industry specific tooling appears in event detection (for use in, for example, anomaly or fraud detection): 47% of financial services vendors versus only 26% of general-purpose vendors.

Conversely, general-purpose AI vendors lead on data generation/synthesis (56% versus 46%) and recognition tasks (37% versus 27%), reflecting their strength in multi-modal capabilities. Qualitative responses from vendors suggest these capabilities translate into financial services applications such as KYC identity document verification, intelligent document processing for onboarding and claims, and transaction pattern recognition for AML, areas where general-purpose models are increasingly embedded into compliance and operational workflows.

AI vendor levels of autonomy

When classified by the most advanced level of autonomy that their AI products/services enable, vendors' AI products are relatively evenly distributed across the four tiers: from less autonomous human support (21%) and human-in-the-loop (29%) to more autonomous, human-on-the-loop (26%), and fully autonomous (24%) AI vendor systems.

The largest single group (29%) operates at the human-in-the-loop level, where AI acts but requires human approval. Notably, half of all vendors (50%) have progressed beyond pure support advisory to human-on-the-loop or fully autonomous operations.

Figure 4.2: Levels of AI action autonomy by vendor product offering

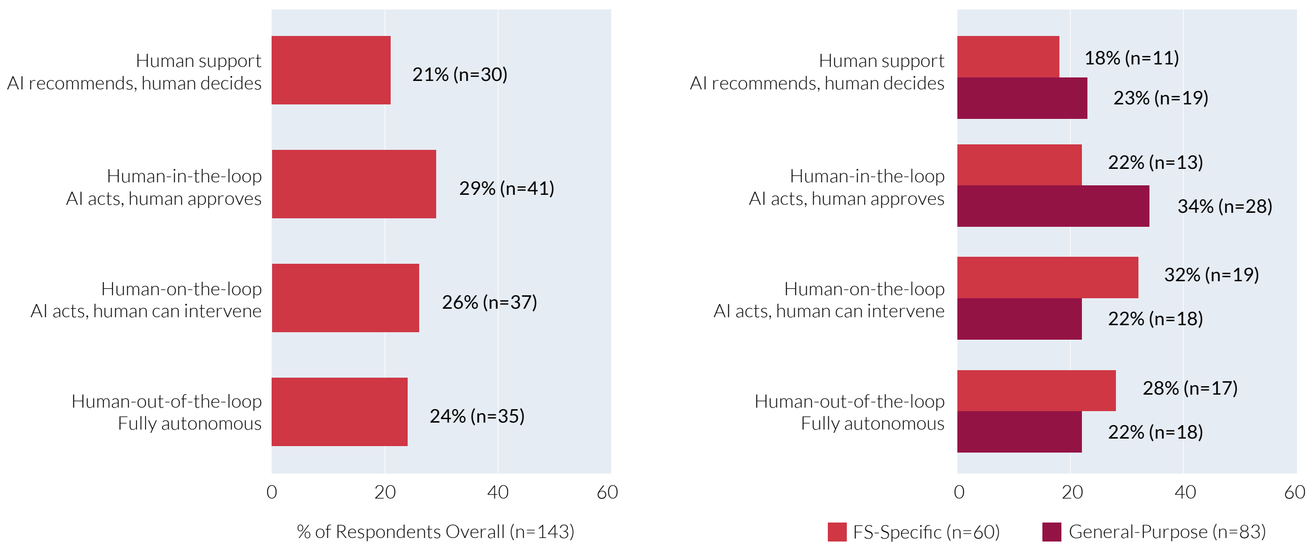

Levels of Action Autonomy: Financial services-specific tools tend to embed more complex regulatory and business process logic into their solutions. They also tend towards offering more autonomous systems.

28% of financial services-specific products operate fully autonomously (human-out-of-the-loop), compared with 22% of general-purpose tools, which predominantly default to human-in-the-loop (34% versus 22% of financial services-specific). This may reflect the fact that workflow-specific vendors embed enough domain logic (regulatory rules, risk thresholds, compliance guardrails) to be trusted with higher autonomy in well-defined, repeatable tasks. Examples of such tasks include transaction monitoring and alert triage. Conversely, general-purpose platforms remain as 'advisory by design' since they lack the built-in regulatory context to operate independently in high stakes financial services environments.

The vendors that enable autonomous operations (28% financial services) are largely the same ones performing event detection and goal-driven optimisation; tasks where real-time, high-frequency decision-making demands reduced human latency and manual review and approval.

The survey's findings around increasing autonomy of AI systems, correlate with Anthropic's recent Economic Index data (2026)[10] which show directive task delegation jumping from 27% to 39% in eight months. Here, the data in this study show that 60% of financial services-specific vendors are already offering more autonomous human-on-the-loop or fully autonomous workflows versus 44% for general-purpose tooling.

Figure 4.3: Highest autonomy level across AI use cases by vendor type

FS-Specific Vendors (n=60)

	Human support (advisory)	Human-in-the-loop	Human-on-the-loop	Fully autonomous
Event detection (fraud)	10%	8%	20%	8%
Goal-driven optimisation	10%	8%	15%	13%
Reasoning & knowledge	8%	10%	13%	18%
Info retrieval & matching	8%	15%	13%	15%
Forecasting	10%	3%	13%	10%
Personalisation	5%	8%	15%	13%
Data generation/synthesis	8%	7%	17%	13%
Recognition	3%	2%	12%	10%
Interaction support	5%	3%	8%	10%
Translation/transcoding	2%	2%	2%	7%

General Purpose Vendors (n=83)

	Human support (advisory)	Human-in-the-loop	Human-on-the-loop	Fully autonomous
Event detection (fraud)	5%	12%	11%	12%
Goal-driven optimisation	8%	13%	11%	16%
Reasoning & knowledge	7%	12%	11%	11%
Info retrieval & matching	5%	27%	10%	14%
Forecasting	10%	17%	11%	16%
Personalisation	10%	14%	12%	10%
Data generation/synthesis	13%	18%	11%	14%
Recognition	4%	17%	6%	10%
Interaction support	5%	11%	7%	11%
Translation/transcoding	5%	10%	6%	10%

Highest Autonomy Level Offered

0%-4% 5%-9% 10%-14% 15%-19% 20%-24% 25%-30%

% of vendors

Cloud and foundation model providers

Foundation models

A foundation model is a large-scale, general-purpose artificial intelligence system trained on vast amounts of diverse data. Rather than being designed for a single specific function, a foundation model possesses broad underlying capabilities – such as language comprehension, reasoning, and pattern recognition – and can be adapted or fine-tuned to perform a wide variety of specific, downstream tasks across different industries. The following section analyses some of the data on usage within the financial services sector.

Foundation model concentration: A substantial 69% (n=423 of 615) of the institutions that responded to this question use OpenAI's foundation model. This represents a 19 per cent lead over their nearest competitor (Google at 57%). This indicative market composition will have no doubt shifted in the months since the completion of this survey, given how quickly the market is shifting. This level of single-provider dominance could create critical supply-chain, pricing and operational resilience vulnerabilities – echoing recent warnings from the Financial Stability Board (FSB) in its report on the financial stability implications of AI.[11] The later chapter on risks in this report examines a potentially concerning dimension of this vulnerability: only a fraction of regulators currently collect data on third-party AI dependencies among their supervised institutions, and nearly half have no plans to begin within 2 years.

Figure 4.4: Foundational model providers by stakeholder group – industry (n=342), regulator (n=130), fintechs (n=198), traditional FIs (n=144)

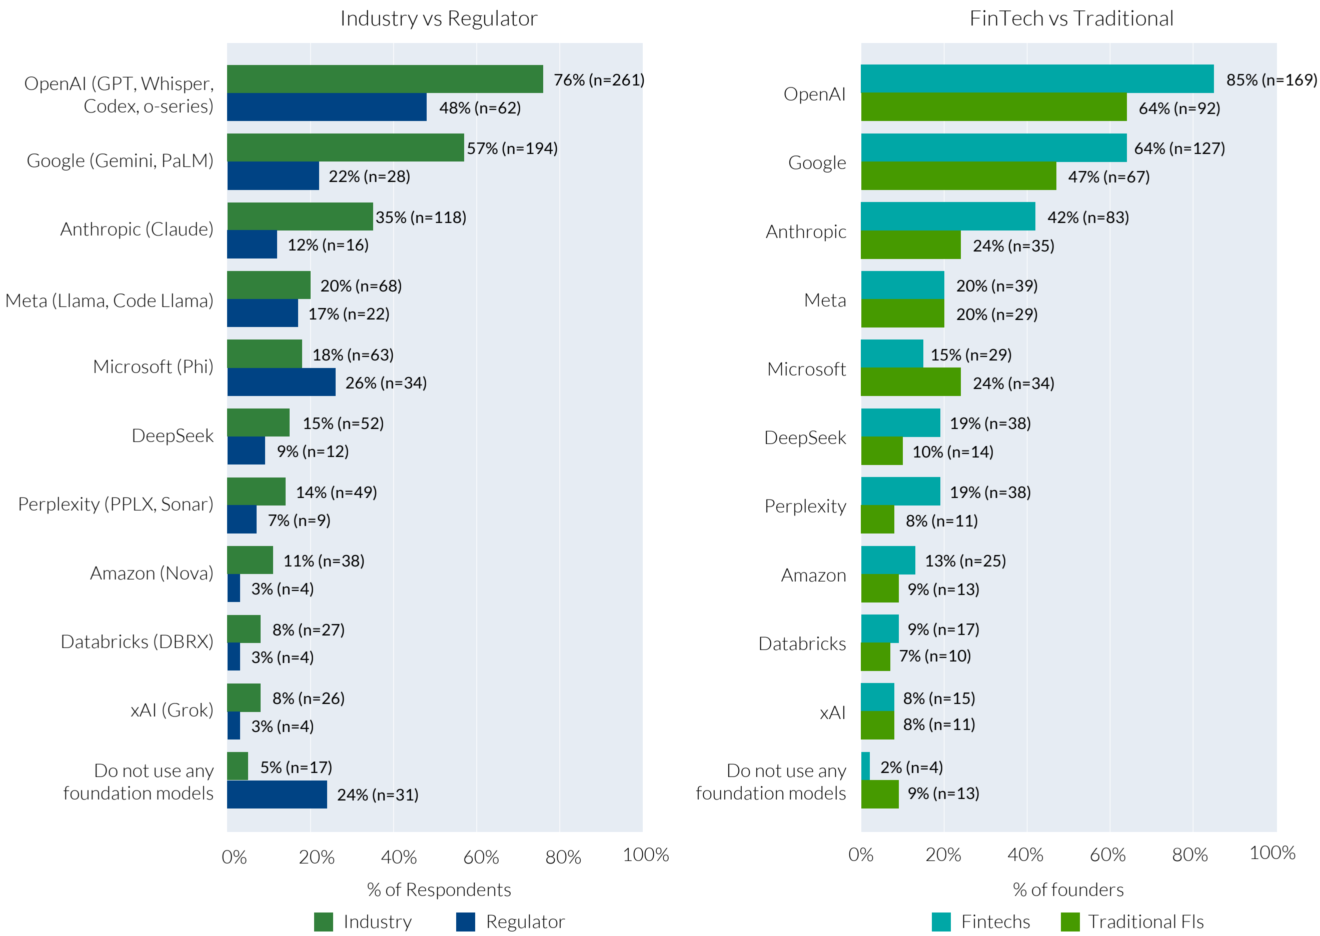

Foundation models – industry versus regulators: In terms of comparison between the financial industry and regulators, the gaps are substantial – 28% for OpenAI, 35% for Google, 23% for Anthropic. Meanwhile, one of the most striking findings from this data is that 24% of regulators use no foundation model at all versus just 5% of industry. This is a significant capability divide that has implications for supervisory technology development and the ability of regulators to keep up with the rapidly evolving AI systems being deployed by the firms they oversee. This divide is further examined in the chapter covering AI use cases by regulators towards the end of this report, showing that 81% of regulators remains at Exploring or Piloting stages in their own AI use.

Foundation models – Advanced Economies (AEs) versus emerging markets and developing economies (EMDEs): OpenAI is used most widely by survey respondents in both AEs (79%) and EMDEs (74%), followed by Google in both groups. However, the adoption gaps between AEs and EMDEs are largest for open-weight and cost-accessible models: Google is used by 62% of EMDE firms versus 50% in AE (+12pp), Meta by 24% versus 15% (+9pp), and DeepSeek by 18% versus 11% (+7pp). Anthropic and Snowflake, by contrast, skew toward AEs.

Figure 4.5: Foundation model providers by economic development – advanced economies (n=144) versus EMDEs (n=197)

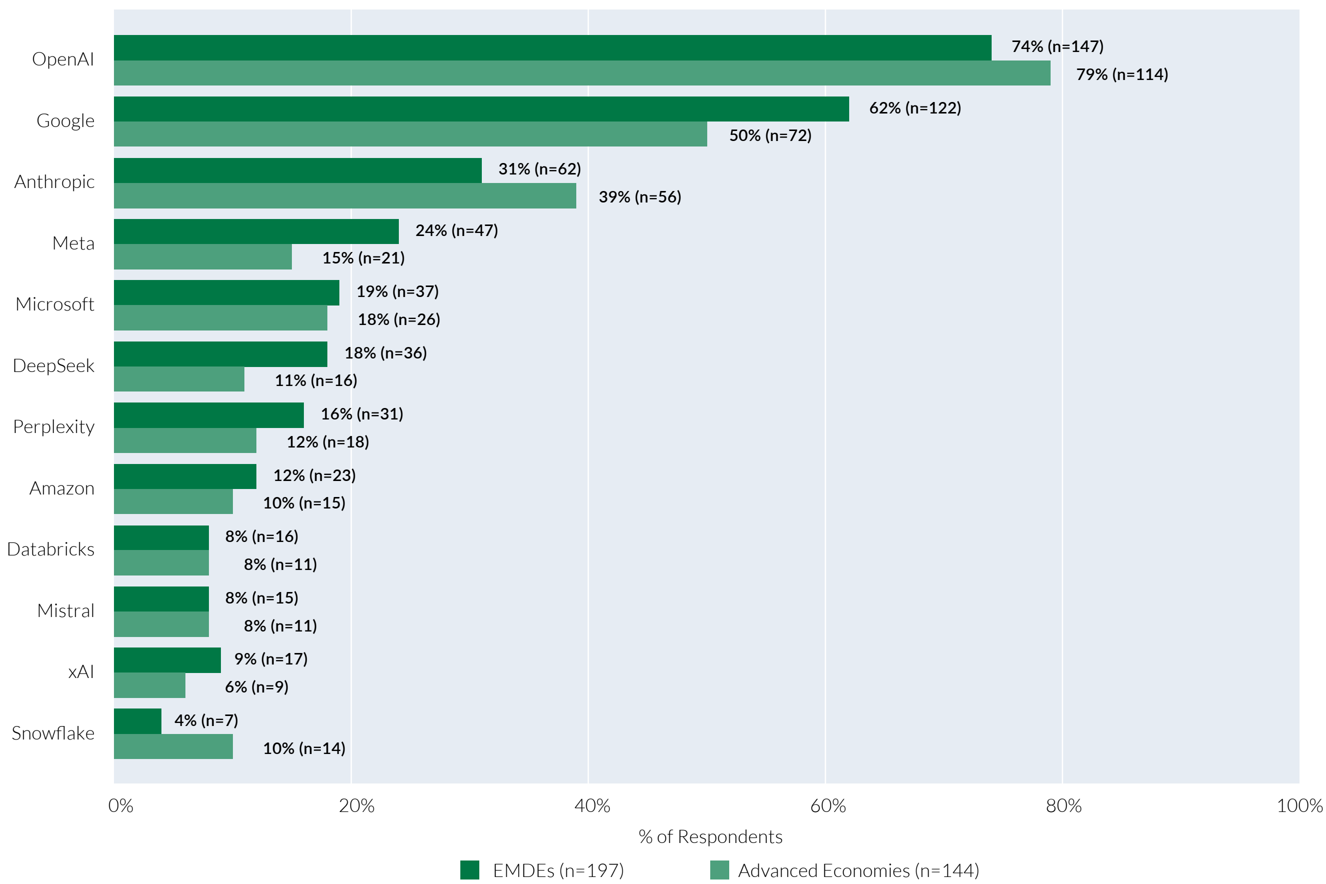

Differences with DeepSeek: Notably, open-weight model DeepSeek has rapidly captured 15% industry adoption, including 19% of fintech firms since launching publicly in late 2024/early 2025. It has leapfrogged models offered by established US players including Amazon (11%).

Figure 4.6: DeepSeek adoption by region (% of region respondents)

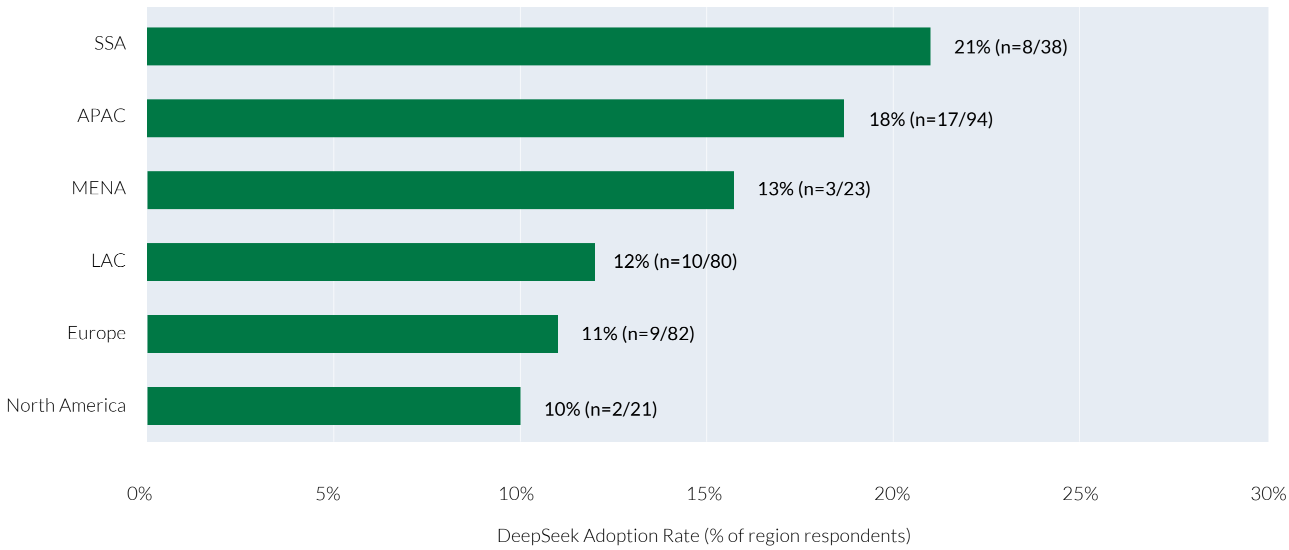

Regionally, DeepSeek adoption is highest in Sub-Saharan Africa (21%) and Asia-Pacific (18%), well above Europe (11%) and North America (10%). DeepSeek offers an open-weight model, (as opposed to the closed proprietary models offered by, for example, OpenAI, Anthropic or Google) which can lower cost barriers for lower-resource institutions. Please note the limited sample sizes.

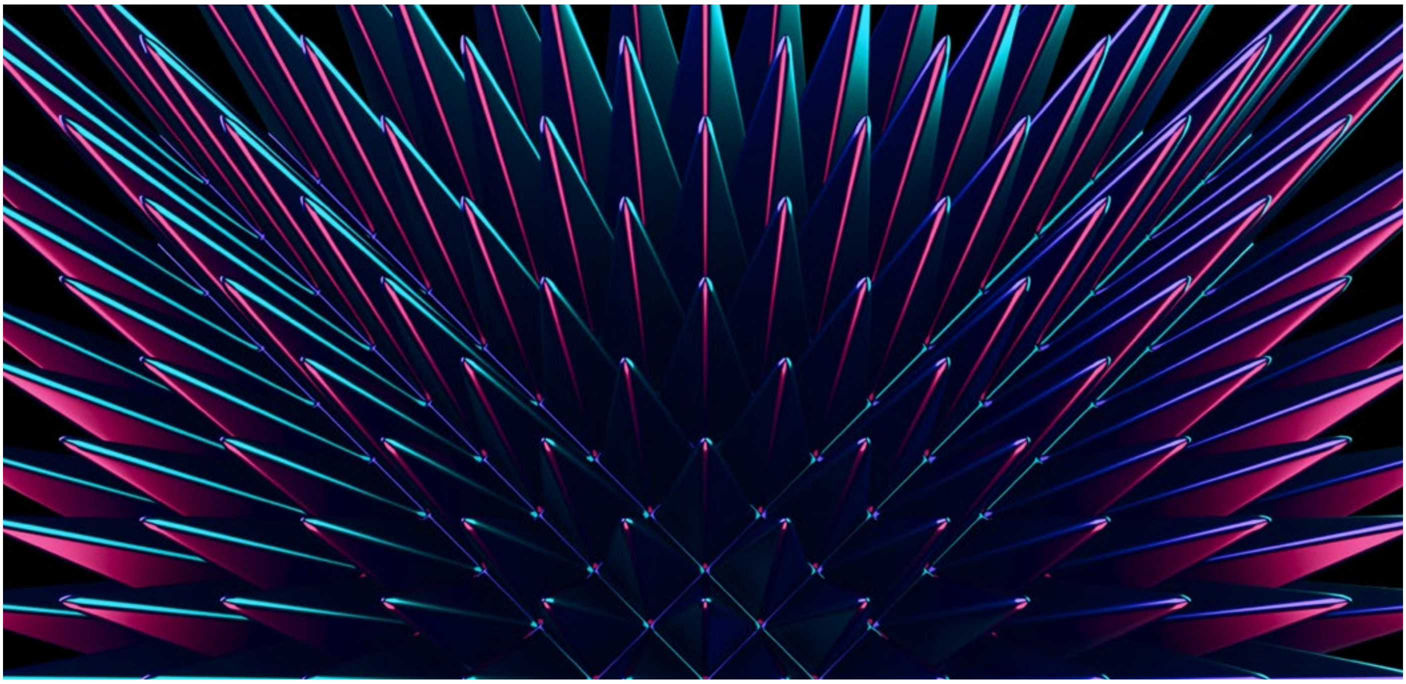

Combining responses across all AI foundation models: Looking at the wider AI model vendor landscape, the tree map below combines all survey responses from across industry, vendors and regulators. The prominence of OpenAI (68.8% of total respondents), Google (46.8%) and Anthropic (32%) is a notable critical third-party risk consideration.

Figure 4.7: Foundational model providers: all survey respondents

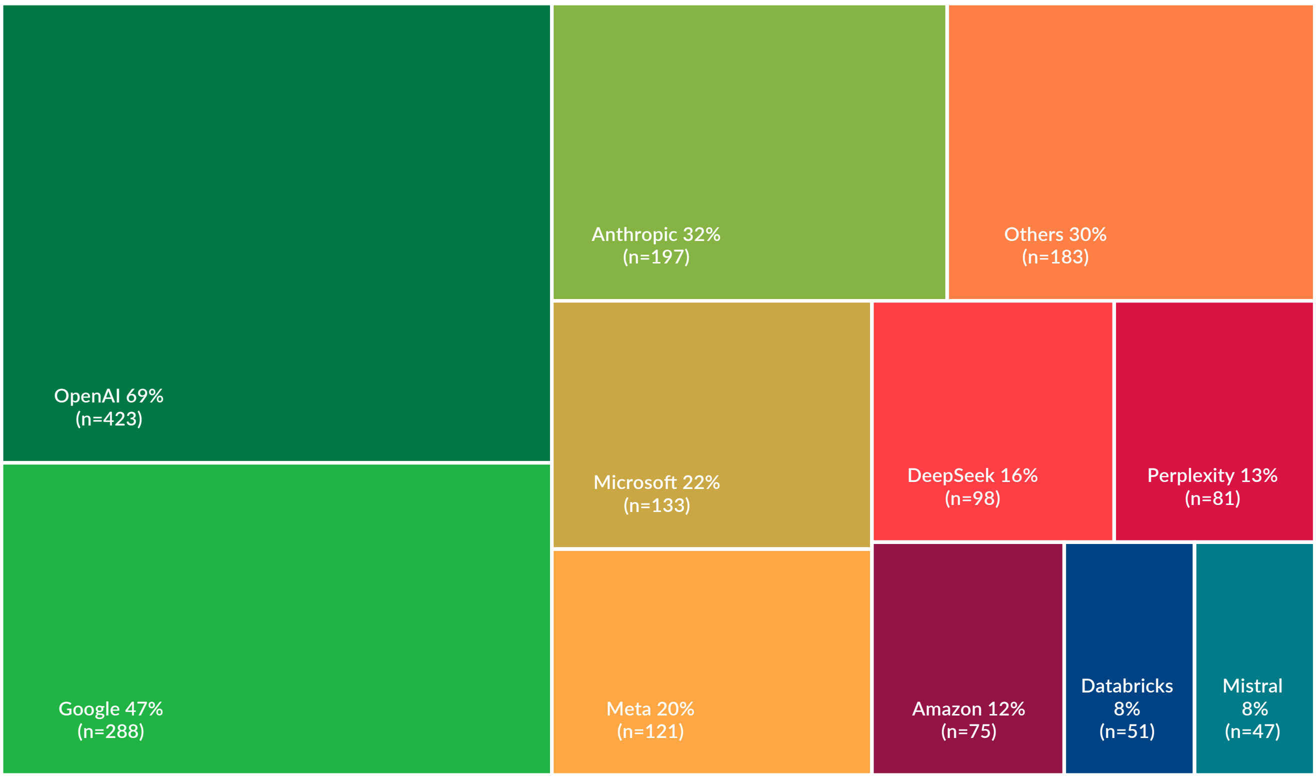

Perspectives on cloud infrastructure

Cloud infrastructure and platform providers supply the data storage, scalable compute capacity and managed environments required for running complex AI workflows. The market is dominated by three major vendors (Amazon, Microsoft, and Google – who are also major AI tech vendors) that provide the foundational physical networks, alongside advanced data platforms that operate on top of these cloud providers to deliver optimised data architectures and AI tooling. This ecosystem also includes firms that choose to run their AI workflows locally, on premises on their own infrastructure.

Industry versus regulator cloud vendors: A noticeable gap exists between regulators and the industry they oversee. AWS has the largest share (46%) of the financial industry based on survey respondents while the most common cloud vendor for regulators is Microsoft (39%).

Figure 4.8: Cloud infrastructure providers by respondent group – industry (n=340), regulator (n=123), fintechs (n=197), traditional FIs (n=143)

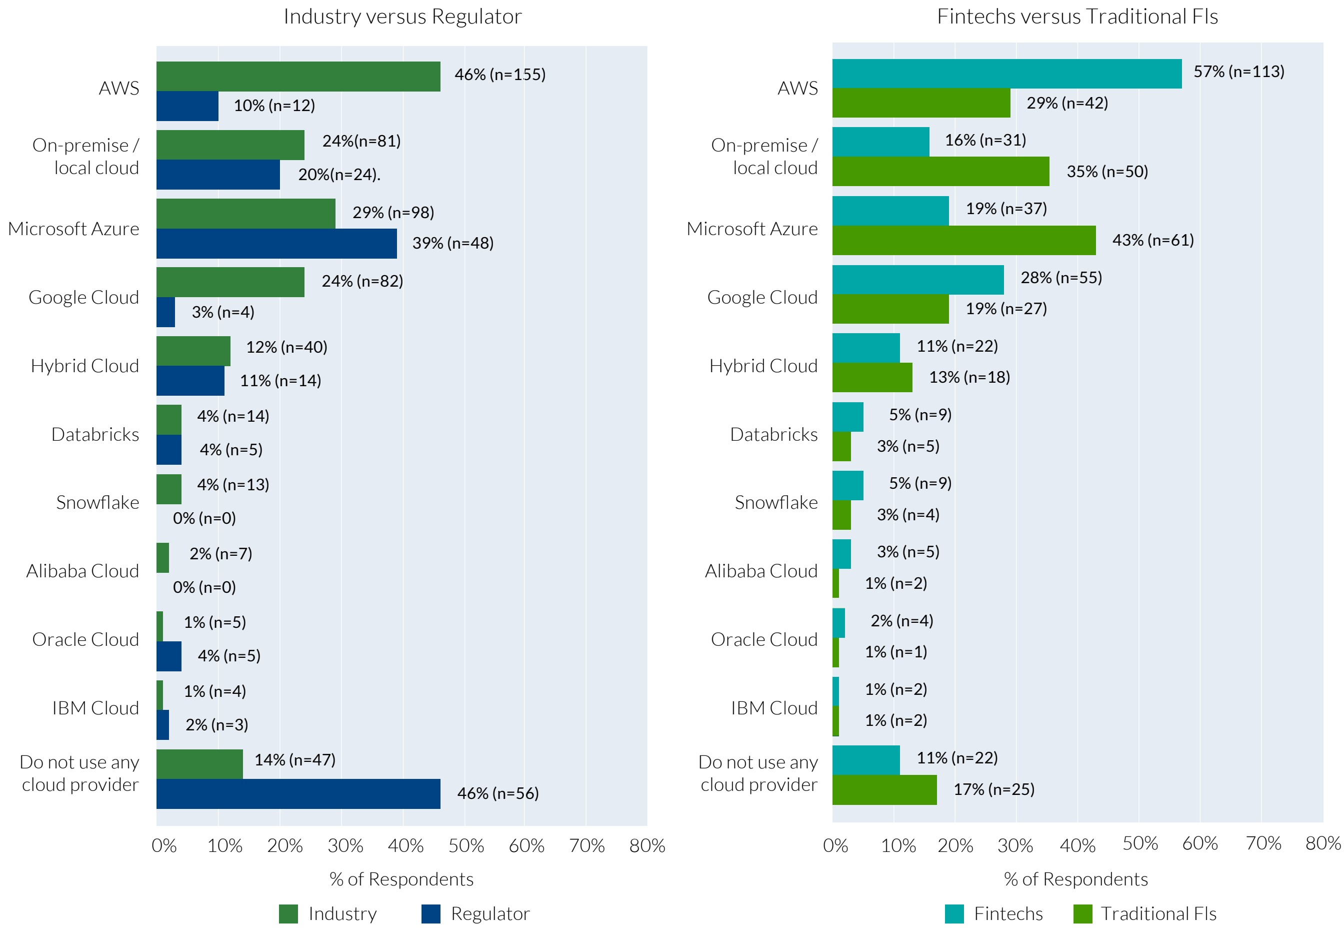

More strikingly, 46% of regulators use no cloud provider at all, the single largest response category, compared with just 14% of industry. This lack of foundational cloud architecture may restrict regulator capacity to monitor massive, cloud-based AI workloads within industry.

Cloud infrastructure in fintech versus traditional FIs: AWS is used by most fintech cloud users (57%) while Azure is used mostly by traditional FIs (43%). Google Cloud is used by around 28% of fintechs that responded to this question, compared with 19% of traditional financial FIs. Interestingly, traditional FIs also rely heavily on on-premises deployment (35% versus 16%) which may indicate a reliance on long-standing legacy infrastructure but could also be driven by cost efficiency and security posture. This infrastructure divergence maps onto the maturity distribution documented in the first chapter, where traditional FIs are 17% behind fintechs at the Scaling or Transforming stages.

Cloud providers: AEs versus EMDEs: AWS leads in both AEs (48%) and EMDEs (44%), but the competitive landscape differs beneath it. Microsoft's Azure is used more widely in AE markets (36% versus 24%, +12pp), while EMDE firms surveyed rely more on on-premises infrastructure (26% versus 21%) and are nearly twice as likely to use no cloud provider at all (17% versus 9%).

Figure 4.9: Cloud providers for AI by economic development – advanced economies (n=143) versus EMDEs (n=197)

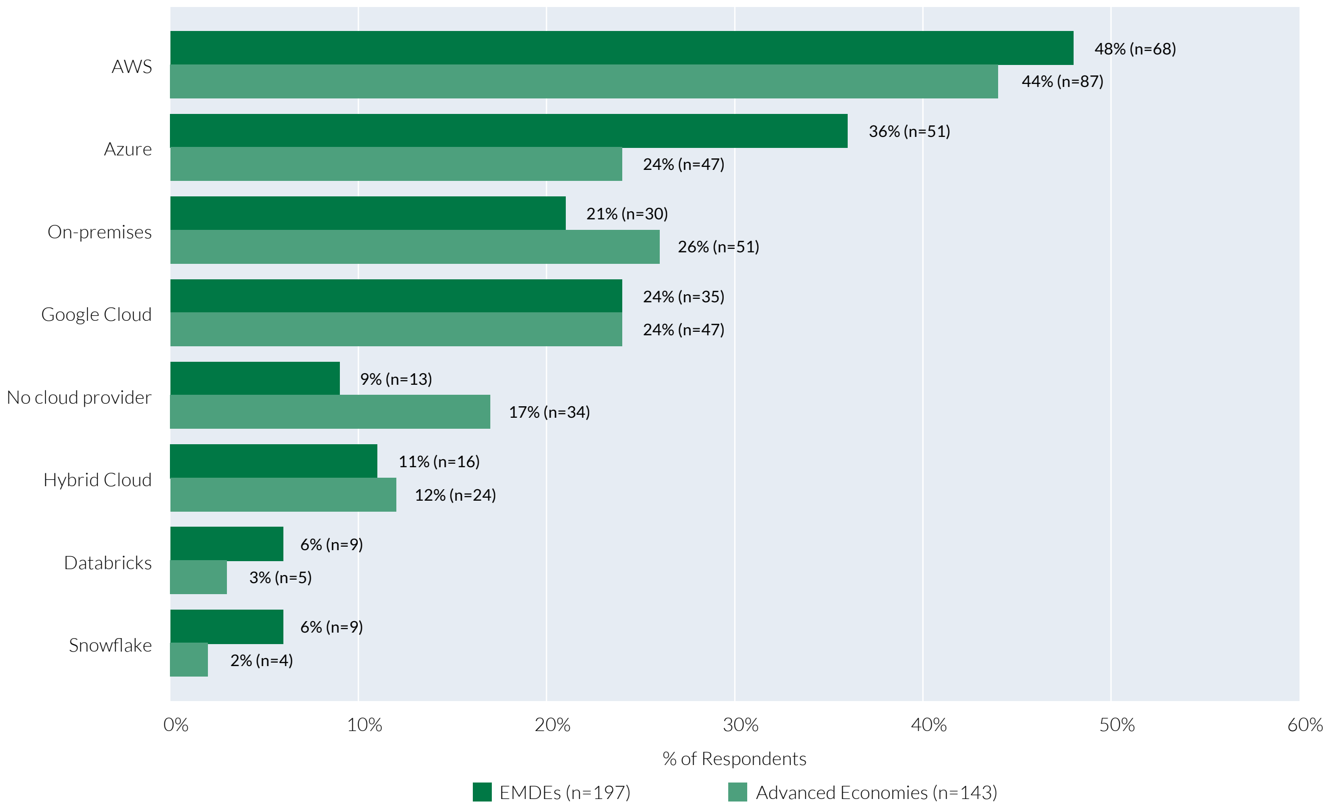

When looking at the total survey responses across all cloud vendors, AWS is used by 41% of the total survey respondents to this question, followed by 33% of respondents for Microsoft Azure. Close to 24% of respondents to this question utilise on premises/local cloud deployment.

Note that survey respondents may have multiple cloud providers and so the total does not add up to 100%. In addition, note that there is a limited sample size for many of the providers which is not reflective of actual market share in the industry, but rather the number of respondents to this question.

Figure 4.10: Cloud infrastructure providers: all survey respondents – industry (n=340), AI vendor (n=142), regulator (n=123) (multi-select)

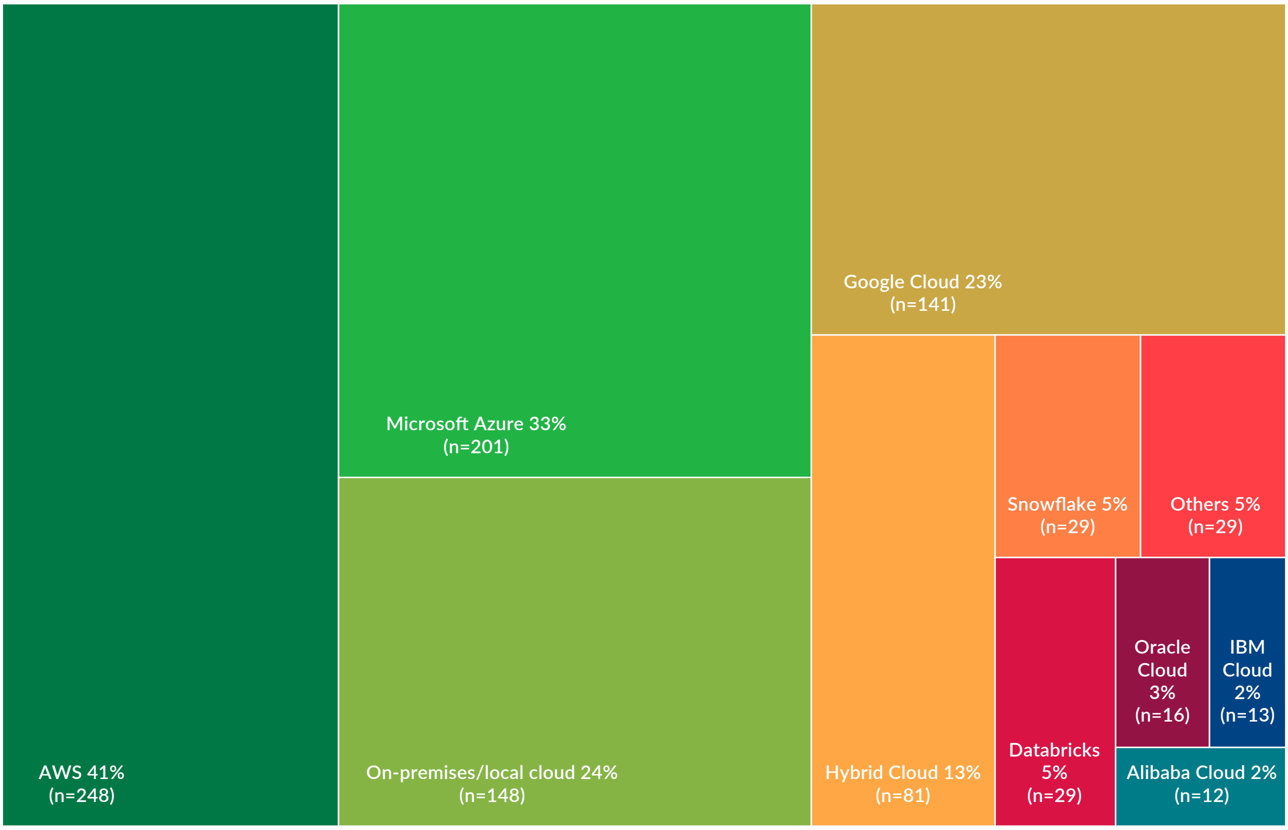

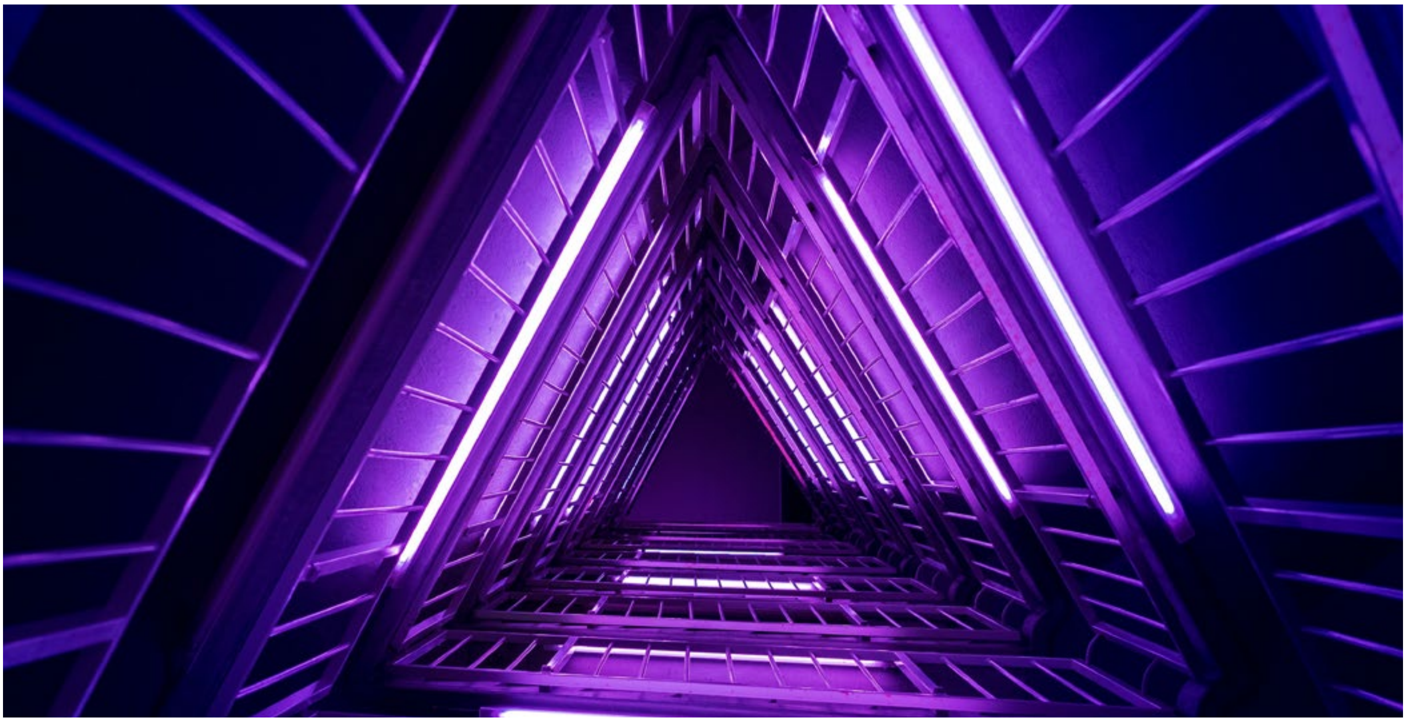

Partner perspectives: BIS – The AI supply and value chains

By Vatsala Shreeti, Kumar Rishabh and Jon Frost[12]

Organisations around the world, within industry, the public sector and beyond, are applying AI tools, at different levels of maturity as this survey highlights.

However, while many organisations are applying AI, far fewer are producing it. This study corroborates findings that a handful of providers dominate the foundation model market. Only five providers account for over 80% of the foundation models used by survey respondents.

Recent BIS work has mapped the AI supply chain that underpins the provision of AI products and services (Gambacorta and Shreeti (2025)).[13] It consists of five interconnected layers:

1. Hardware that includes specialised chips such as graphical processing units (GPUs).
2. Cloud services and related infrastructure.
3. Training data that include vast multimodal datasets from public and proprietary sources.
4. Foundation models.
5. User-facing AI applications for specific use cases.

Each layer has distinct market dynamics, and many are shaped by high fixed costs, economies of scale and scope, and network effects. These economic forces can favour the emergence of large AI firms and reward those that integrate across multiple layers of the supply chain. Indeed, a recent BIS bulletin notes that the global AI supply chain contains a few large firms based in the United States, China, Chinese Taipei, Korea and the Netherlands (Frost et al., (2026)).[14] Among the top 20 global AI firms, the top seven are all US-based and together account for more than twice the market capitalisation of the next 13. Many of these firms are present across most or all layers of the supply chain, playing a dual role of selling products and services externally and building capacity for internal use.

These global AI giants are becoming increasingly important for their respective jurisdictions. By the end of 2025, publicly listed AI giants accounted for 30-40% of the total market capitalisation in the United States, Chinese Taipei, Korea and the Netherlands. Their weight in the real economy is also growing: by the end of 2024, these firms accounted for 25% of aggregate capital expenditure in Korea and 21% in the United States, with their revenue shares also increasing.[15]

Some AI firms are active across all five layers of the AI supply chain.

Figure 4.11: At the top of the AI market, scope and scale go together

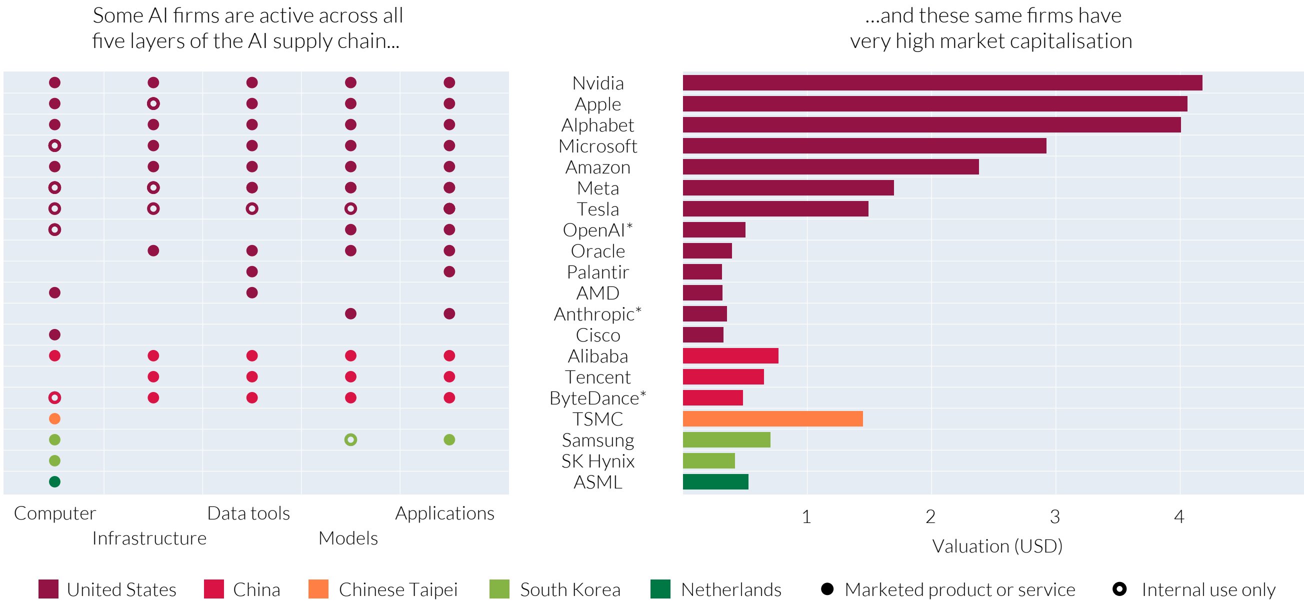

Valuations are in USD and rounded. Jurisdiction in the legend refers to headquarters location. If a firm both sells and uses a capability internally, it is coded as marketed. Valuations are whole firm and are not allocated across layers. * = privately held firm.
Sources: PitchBook (public market capitalisations, private company valuations); authors' coding of supply chain roles from publicly available information on firm products and services.

In addition to understanding the AI supply chain, research is addressing how economies position themselves in the AI value chain. While related, the two are distinct concepts, each reflecting a slightly different lens of analysis.[16] Mapping the AI value chain helps authorities to consider how domestic industries fit into global trade in services. BIS research shows that, for EMDEs, there are opportunities to apply AI to improve growth and productivity.[17] Some EMDEs have found niches in modern service exports, in areas such as data labelling, and domestic uses of AI in areas like the financial sector, healthcare and manufacturing. Some have introduced national AI strategies.

Going forward, further research is needed to better understand the opportunities and assess the impact of AI in practice. Surveys like this one by CCAF are a useful step in this direction.

In summary, the technology data in this chapter reveal that broader deployment does not automatically translate into deeper return. At the same time, the infrastructure on which AI deployment depends (foundation models, cloud platforms) is concentrated in a small number of providers, creating a critical dependency that runs across firm types and geographies. While these dependencies are not specific to the financial services sector, the risk profile may bear more acute implications given that AI-informed decision may have more immediate and far-reaching consequences than in other fields. Correlated exposure to the same providers may also result in systemic events through disruptions which might be contained or absorbed locally in other industries.

Chapter 5

Measuring Value and Adoption Challenges

The Value of AI

AI is delivering clear productivity gains, but value realisation remains uneven and is generally hard to measure for all firms, no matter their size or their sophistication with AI.

AI adoption is increasingly generating measurable improvements across financial services, though the extent of impact varies by function, firm type and organisational maturity. Productivity gains are strongest in technology and operations, while corporate, leadership and risk functions see relatively more modest gains.

The survey findings depict an industry moving from experimentation to selective but meaningful value realisation, with clear gains in operational and technical areas but persistent challenges in measurement, governance and organisational readiness.

AI impact on productivity

AI is broadly perceived to have a positive impact on productivity, with most respondents reporting gains across all functions. This suggests that firms are increasingly well positioned to realise productivity benefits from AI adoption.

The strongest gains were observed in technology, data, and product functions, where 79% of respondents reported positive outcomes. Fintech firms reported greater benefits than traditional firms, with 86% reporting gains compared to 68%. Back office and operations followed closely at 75% overall, with fintech and traditional firms reporting similar results (76% and 72%, respectively), indicating that operational automation has delivered consistent benefits across firm types. In contrast, corporate functions and leadership recorded the lowest positive impact, at 55% overall, with fintech firms at 61% and traditional firms at 48%. Overall, these patterns suggest that fintech firms have advanced more rapidly in deploying AI for customer engagement and technical innovation.

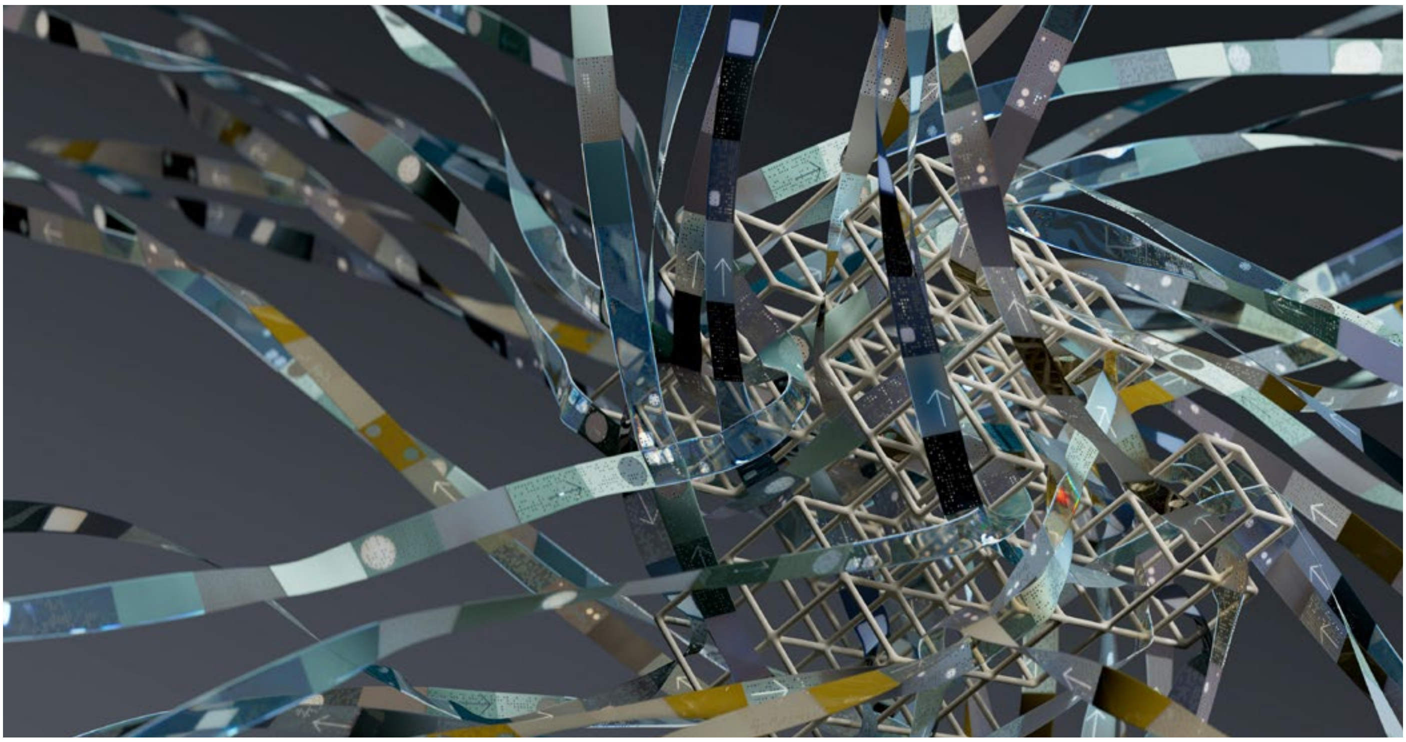

Drivers of productivity

Figure 5.0: AI productivity impact by function and firm type – fintechs (n=203) versus traditional FIs (n=149)

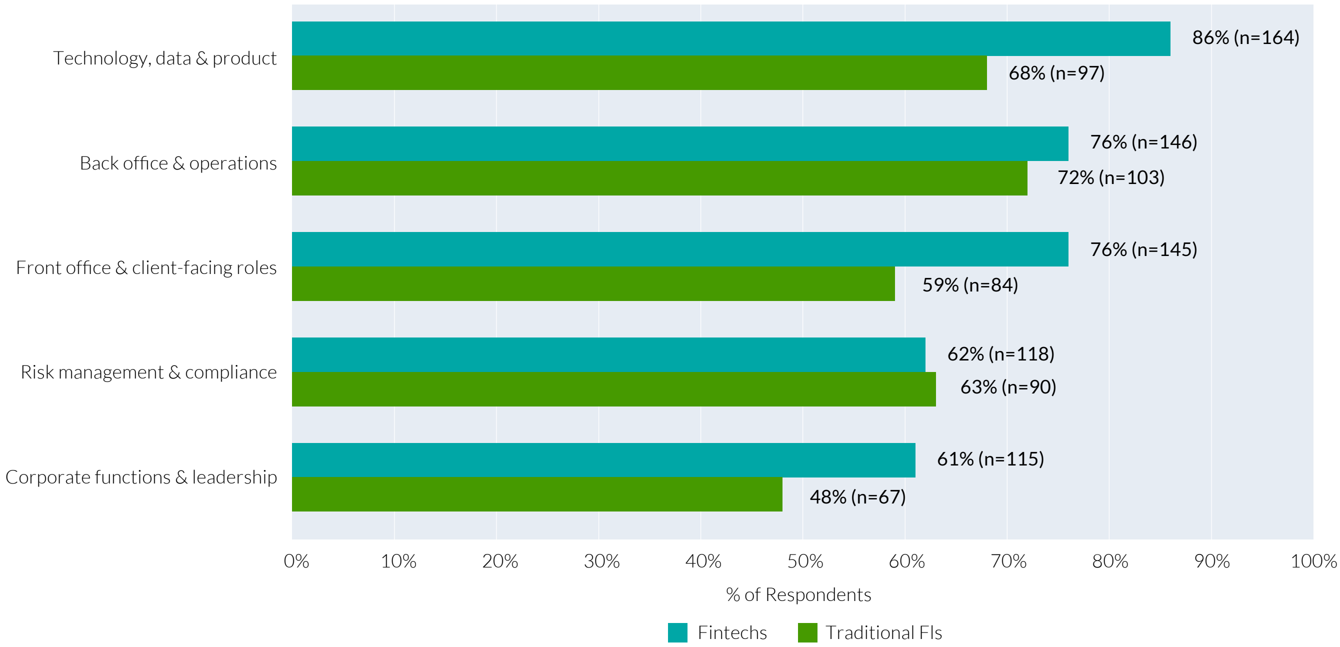

Organisational characteristics also shaped perceptions of AI's impact on productivity. Firms with more mature AI adoption reported higher positive impacts of 74% compared with 60% for those at earlier stages. Also, firms experiencing increased profitability reported stronger perceived productivity gains (80%) compared with those with no impact on profitability (68%), pointing to a reinforcing relationship between financial performance and the value derived from AI.

Geography and investment levels also play a role as firms in advanced economies reported a 73% positive impact, compared with 64% in emerging markets and developing economies. Similarly, firms investing USD 100,000 or more in AI reported slightly higher gains (77%) than those with lower spending (72%).

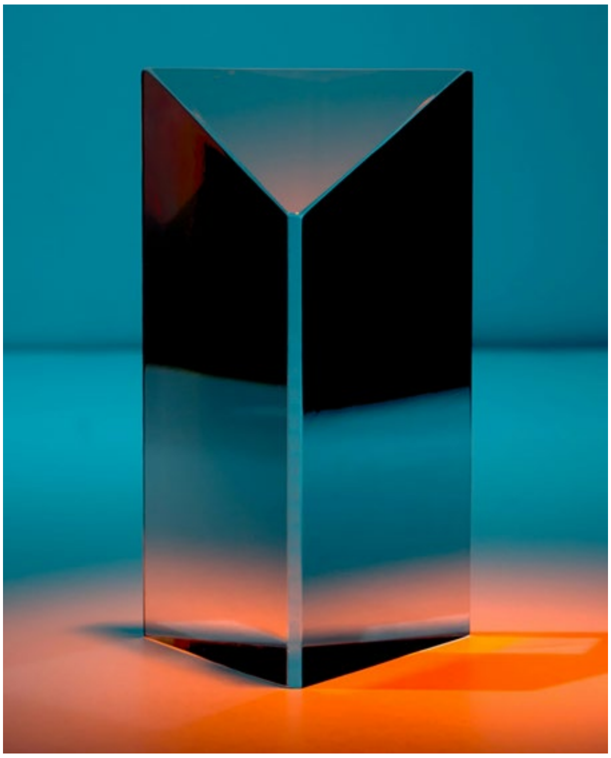

Figure 5.1: Average reported positive AI impact on productivity across comparison groups

AI value proposition: vendor offer versus industry benefits

There is meaningful alignment, as well as divergence, between how AI vendors promote their tools and how financial institutions perceive the resulting value. Both recognise cost reduction and resource optimisation as the primary benefit. In qualitative responses, many firms mentioned drastic reductions in the time required for data entry, document screening, and workflow routing for example, allowing them to manage higher volumes of work without needing to proportionally increase human headcount.

Points of divergence

While vendor and industry agree on the value of efficiency, their secondary priorities differ. Vendors emphasise increased profitability and revenue, and risk management as their primary value propositions. However, industry places a higher premium on speed and latency reduction (53% of industry respondents cited this benefit versus 31% of vendors).

A major gap exists regarding access to expertise. 36% of industry respondents see this as a key value driver while only 6% of vendors see this as their core value proposition to clients. Given that 41% of industry respondents cited access to talent as one of their top barriers to adoption (refer to fig. 5.10), vendors may find a strategic advantage in positioning them as a bridge to securing specialised AI expertise rather than just software providers.

Figure 5.2: Perceived AI value drivers: industry (n=241) versus AI vendors (n=146)

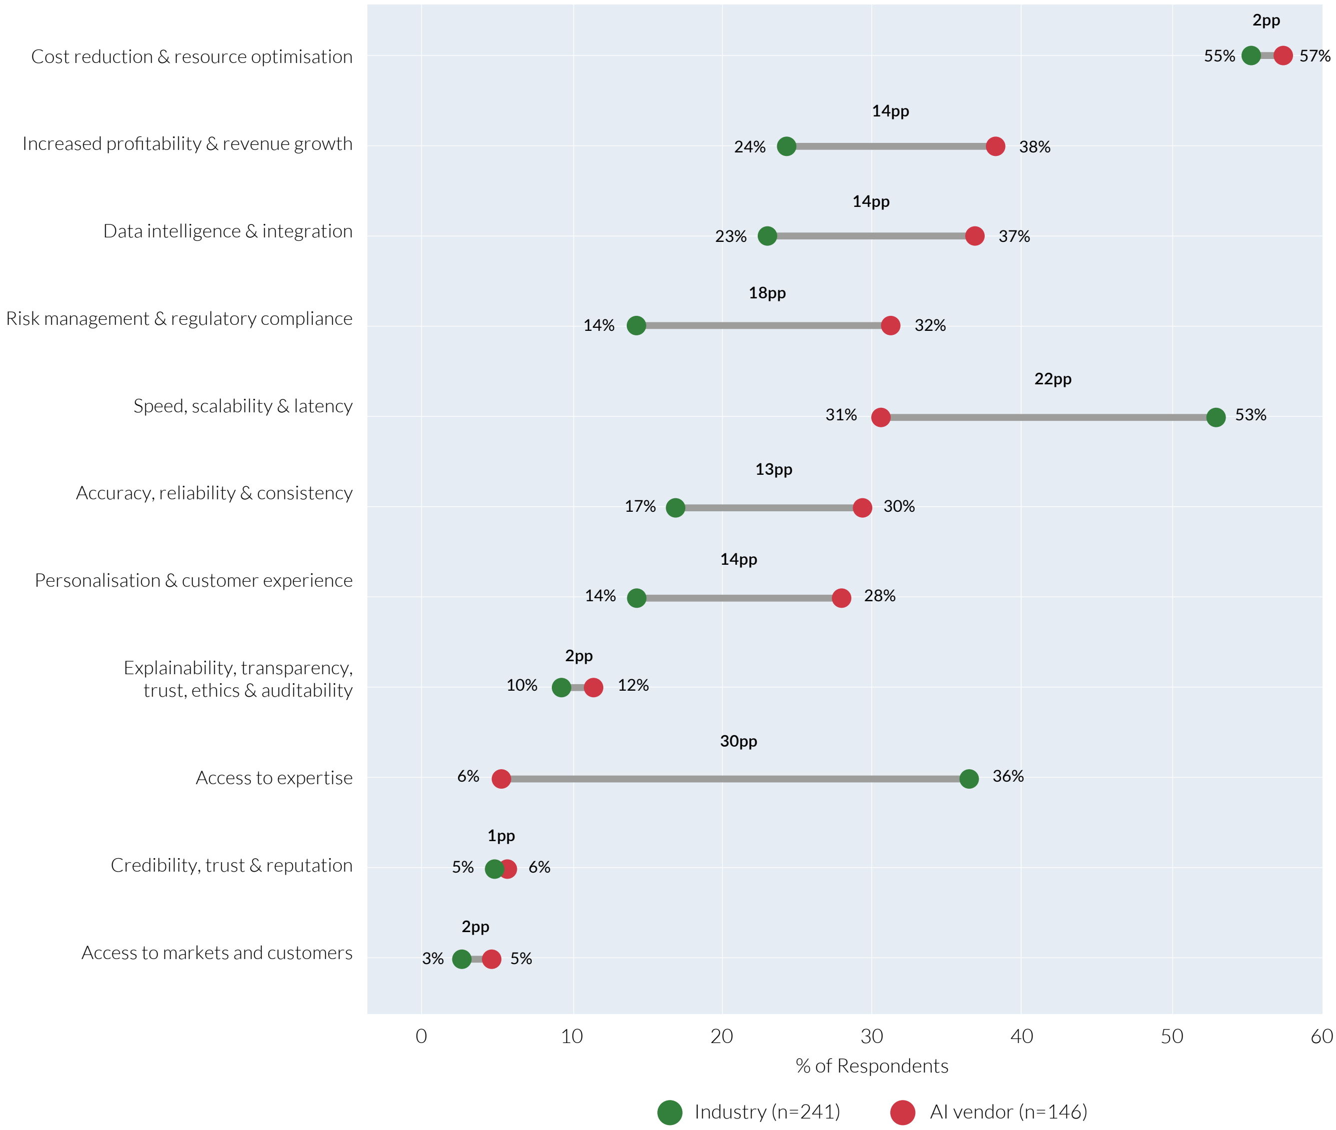

Challenges in quantifying value

While firms broadly perceive a positive impact on productivity, the ability to measure the value of deployed AI products and services remains a significant challenge, regardless of level of AI maturity.

Figure 5.3: Difficulty of measuring AI value by comparison group

Maturity does not solve measurement. Even firms at the Scaling stage of AI adoption reported high difficulty at 60%, suggesting experience alone does not resolve this key challenge.

The measurement difficulty was most pronounced among large firms, with 76% reporting challenges, compared with 49% of small firms. Large firms were also the least likely to report ease (9%, versus 20% among small firms), indicating that greater organisational complexity may play a factor in value of AI measurement issues. However, it is important to note that the sample size for larger firms (n=45) is much smaller. This may reflect legacy IT systems, fragmented data environments, and more rigid governance structures in traditional firms make it harder to integrate and measure AI impact compared with more flexible, data-centric fintech models. [18, 19]

Regulators reported greater difficulty (63%) compared with industry firms (55%) and were less likely to report ease (6% versus 16%).

Overall, the findings highlight that measuring AI value remains a systemic challenge across all firm and organisation types. This may suggest that progress in AI adoption will depend on improvements in data integration, governance and measurement practices.

AI effect on profitability

While productivity gains are perceived across most institutions, approximately 43% of organisations report no impact on overall profitability from AI adoption. However, the downside risk appears low, as only one firm in the sample reported a decrease in profitability attributed to AI adoption. In contrast, nearly 40% reported an increase in profitability.

Figure 5.4: Reported impact of AI on overall profitability – Industry (n=279)

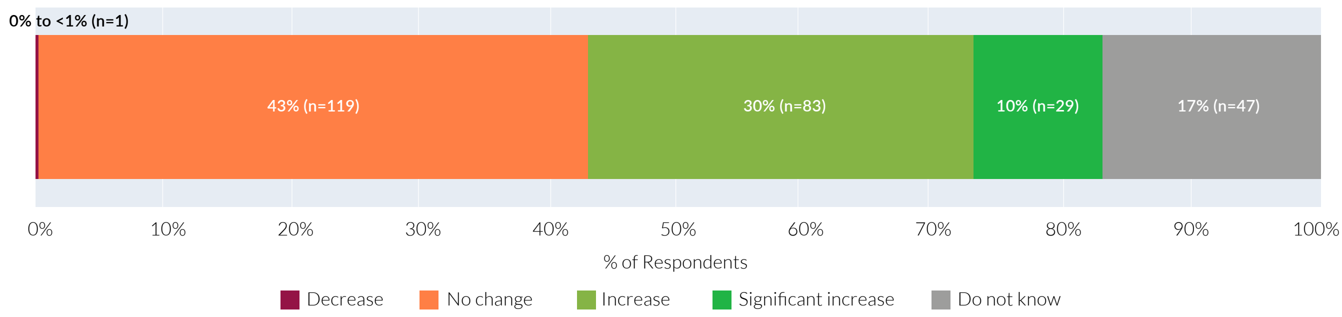

Correlations between AI maturity, spend and profit

Profitability gains are correlated with advanced AI adoption. 64% of more mature adopters of AI reported gains in profitability compared with 33% of less mature firms. 56% of fintech firms recorded increased profitability compared with 34% of traditional FIs. This profitability divergence aligns with the 17% maturity gap in advanced AI adoption between fintechs and traditional FIs identified earlier in this report.

61% of firms that invested over USD 100,000 in the most recent financial year noted increases in profitability compared to 40% of firms spending less than USD 100,000.

Firms with fully in-house or fine-tuned AI models report higher profitability gains (54%) compared with those relying on off-the-shelf or vendor-built solutions (39%). Taken together, this indicates that realising financial value from AI may depend less on adoption alone and more on organisational maturity, technical capability and the level of control over AI development.

While higher investment and more advanced AI maturity are both associated with improved profitability outcomes, these relationships should be interpreted with caution. Correlation alone does not necessarily establish causality, and higher AI spending may reflect underlying strategic commitment rather than act as the direct driver of improved performance. Firms that invested more heavily in AI were likely already more mature along other dimensions such as leadership sponsorship, organisational culture, governance structures or access to specialised AI talent. These conditions may have enabled both greater investment and more effective deployment. As such, spending may be a symptom rather than a cause of higher AI maturity and profitability.

Figure 5.5: Reported AI profitability impact across key comparison groups

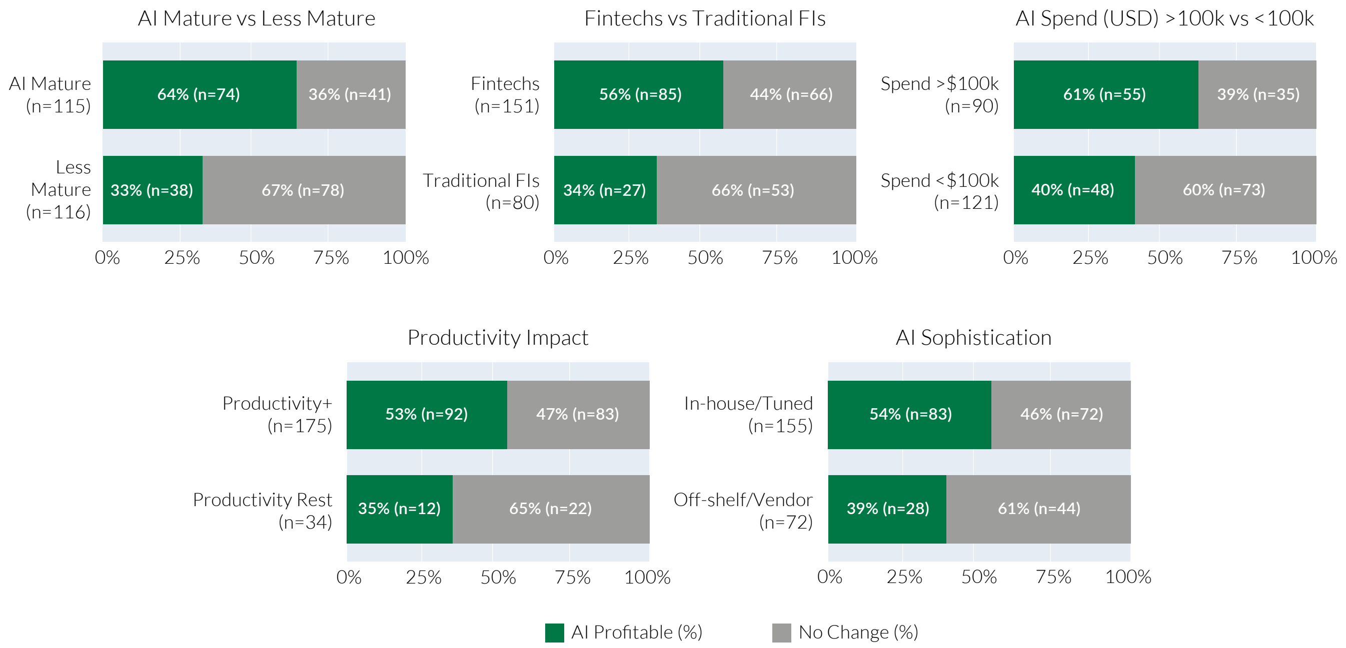

When considering the relationship between reported increases in profitability and size of firms, both in terms of full time employed (FTE) staff, economic development the level of economic development as well as in terms of revenue, there was no discernible impact on profitability between these cohorts based on the respondents to this question. The lack of clear variation across these cohorts may indicate that reported profitability effects depend more on implementation and the translation of AI into operational change. Further research is needed to examine these relationships.

Figure 5.6: Reported AI profitability impact by firm size, economic development, and revenue

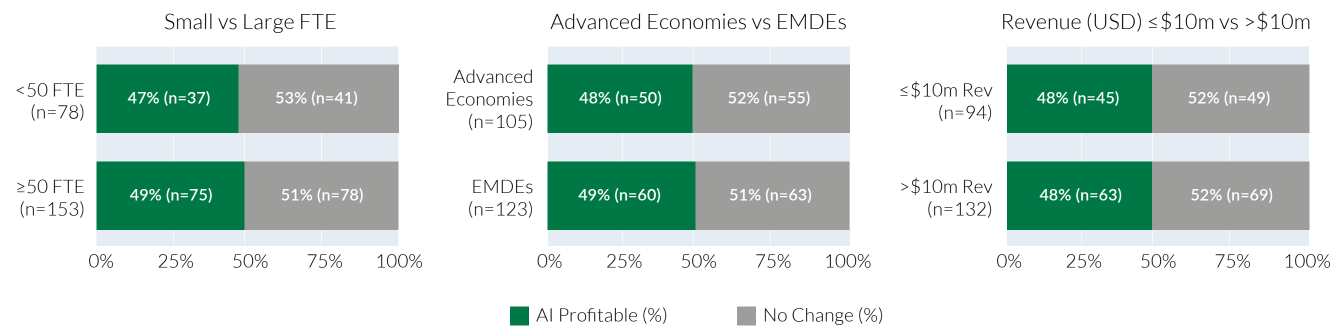

AI impact on the workforce

Turning now to the important area of AI and its impact on jobs within the financial services sector. The results suggest that despite high levels of general maturity and adoption of AI in financial services firms, the actual effect on headcount has remained very limited for the last three years, with 74% of respondents reporting that no significant job losses or gains have been observed due to AI implementation (comprised of 44% noting minimal change and 30% unsure).

Figure 5.7: Reported job impact of AI over the last three years – Industry (n=321)

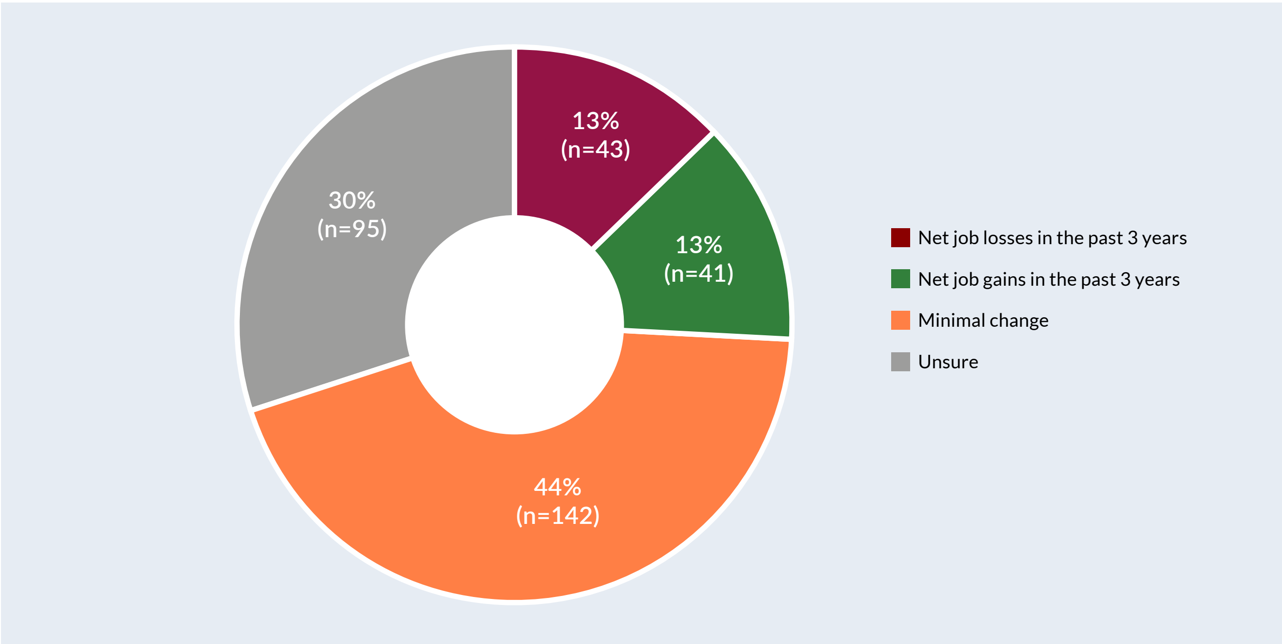

This proportion of responses with no change and unsure may indicate that firms are reorganising their workforce as they grapple with the challenges relating to the adoption and implementation of AI, as discussed in previous chapters of this report.

Expected job impact of AI by 2030

Looking further ahead to 2030, the industry expects structural transformation rather than simple contraction. 25% of firms expect 'Reskilling and Transformation' of the workforce. Combined with the 10% of respondents expecting a net increase, a total of 35% of the industry anticipates a future where job roles are transformed through reskilling or positively impacted by the use of AI. However, a quarter of firms anticipate a net reduction in jobs by 2030.

Figure 5.8: Expected job impact of AI by 2030 – Industry (n=259)

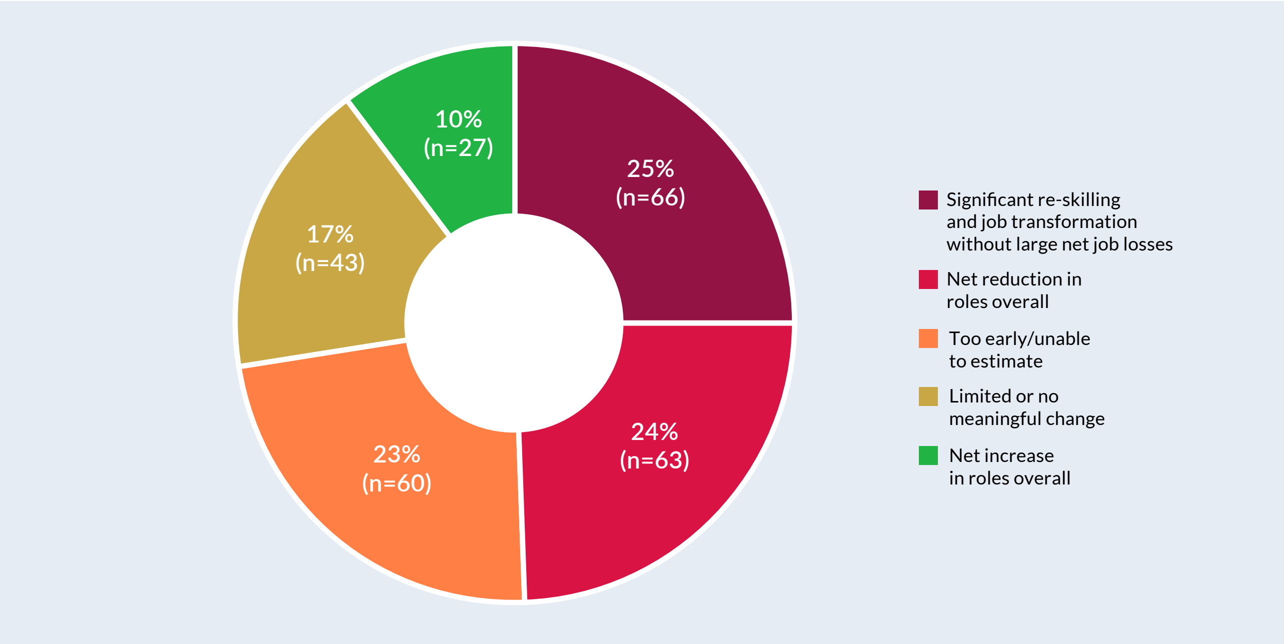

The 2025 WEF Future of Jobs Report,[20] which emphasises that while displacement is a certainty for some jobs and activities, the overarching narrative is of a period of intense volatility where the ability to manage reskilling pathways will be the critical determinant of organisational success.

Expected job impact by 2030: by sub-sector

Employment expectations vary substantially by sector. The Payments sector is the most pessimistic, with 21% of respondents projecting a net reduction. In contrast, the commercial and wholesale banking sector anticipates a net increase of 44% in jobs by 2030.

Figure 5.9: Expected job impact of AI by 2030 across sub-sectors (% of industry respondents)

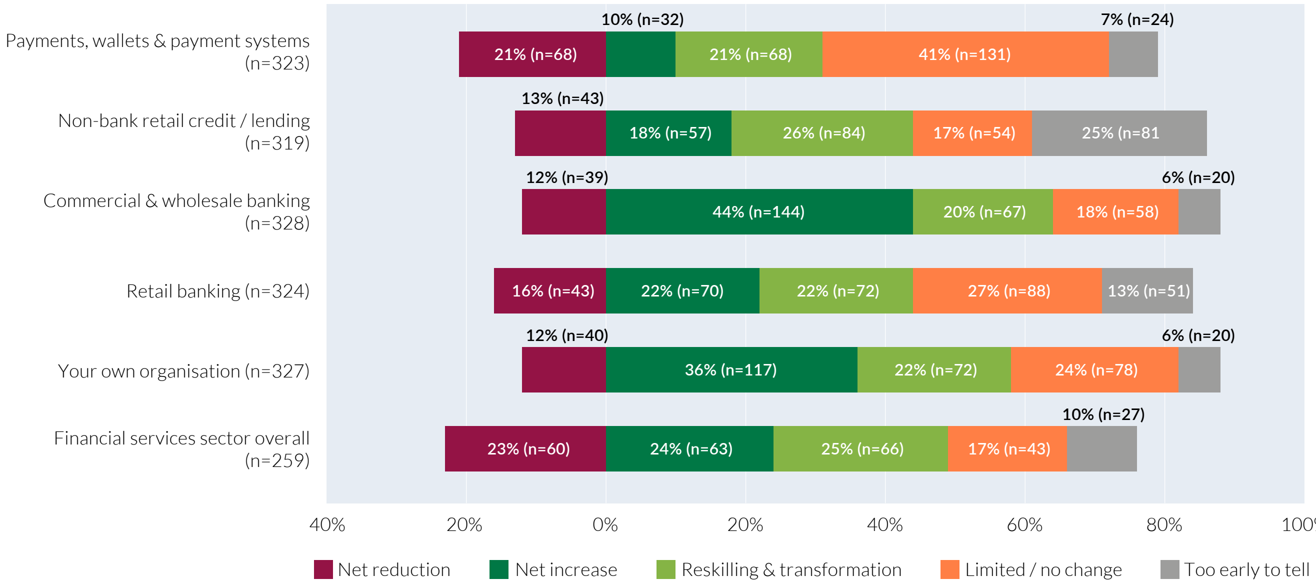

Collectively, these findings indicate a fragmented transition towards 2030. While 'reskilling and transformation' remain a consistent thematic anchor across all sub-sectors (ranging from 20% to 26%), there is a distinct lack of clarity and consensus about the future impact of AI on jobs in the financial services sector.

Complementing this macro-level outlook, OECD analysis of labour-market[21] data indicates that workforce transformation outcomes observed by 2030 will likely depend on the interaction between AI deployment strategies, task re-organisation, and reskilling investments, rather than on AI adoption alone. This reinforces the interpretation of the survey results as hypotheses that warrant continued monitoring and scenario-based analysis as AI diffusion accelerates.

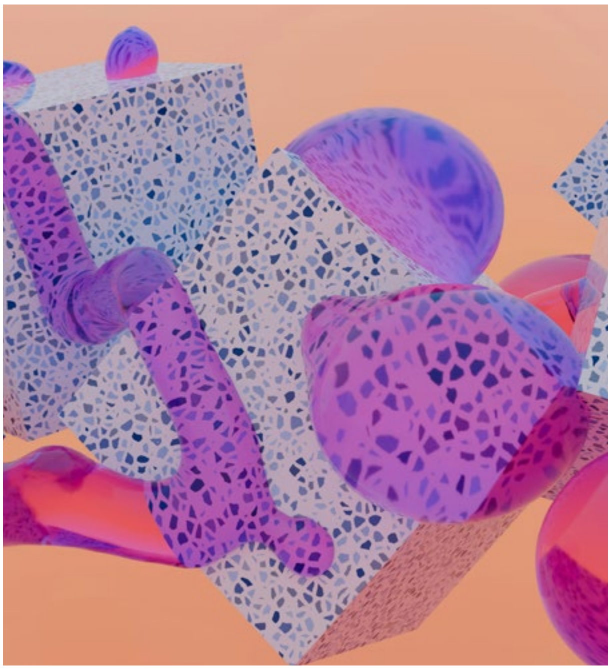

Partner perspectives: WEF – Demonstrated value and barriers to scaling

By Drew Propson, World Economic Forum

Analysis of the survey data confirms that AI's value has been well-established, and that the main question now is whether organisations can bring together all key components needed to scale (data quality, talent, architectural readiness, accountability and supervisory capability). While this finding is unsurprising, it is reaffirming and allows organisations to more confidently build out long-term and large-scale AI integration plans.

Profitability: Related to the above, as ROI and profitability dominate many of today's AI conversations, it is encouraging that 40% of organisations have already noted profitability gains. This insight, alongside the finding that higher spend is linked to greater impact (62% of organisations spending >USD 100,000 annually on AI reported increased profitability, compared with 37% among lower-spending peers) offers leaders a compelling, evidence-based case to support budget discussions and stakeholder buy-in, especially where resources are constrained. Fintechs continue to outperform traditional FIs on AI-driven profitability, a dynamic that is likely to sharpen competitive pressure across the sector.

Pain points for adoption: While AI in financial services still far from reaching its full potential, having insights on obstacles to greater scaling is essential. The survey data reveals that the core constraints are similar for all ecosystem actors, with data quality, talent and infrastructure at the top of the list. This convergence points to a clear opportunity for cross-sector collaboration and collective solutions. However, there will be nuances to approaches taken, as while the challenges are common, their severity varies notably across groups. The most significant divergence emerges between industry participants and regulators (the latter experiencing these more strongly), while fintechs and traditional FIs report a broadly comparable experience of these barriers.

WEF strategic alignment

The World Economic Forum is committed to exploring how technology is reshaping financial services and to ensuring that industry, policymakers and regulators have the insights and tools needed to drive meaningful innovation, while maintaining financial system stability.

In recent years, a significant area of this work has focused on the evolution of fintech globally. Through a longstanding collaboration, the Forum and CCAF have conducted several Future of Global Fintech studies tracking fintech's growth, performance, and role within the broader financial system. A standout finding from the latest edition, published in June 2025,[22] was that 84% of fintechs were partnering with traditional FIs, driven by shared ambitions in technology and infrastructure. Notably, 80% of fintech respondents were also implementing AI across multiple business domains. Together, these findings pointed to the importance of better understanding how AI is being deployed across the full financial services ecosystem, and this current study is a direct response to that goal.

While the Forum has a substantial body of work examining AI in financial services spanning the past decade, the pace of technological change makes it essential to continuously build upon this foundation. This collaborative study significantly advances that work, offering rich data and insights from across industry and government actors. Complementing the quantitative findings, the Forum is also compiling case studies that bring real-world AI applications to life.

Collectively, these efforts aim to equip ecosystem actors with a comprehensive and up-to-date view of AI deployment across financial services, supporting more informed -decision making in this fast-evolving space.

Challenges

The transition to an AI-driven financial services sector faces several constraints around data quality, fragmented systems, technology and infrastructure challenges, and limited institutional capabilities. These barriers shape not only how AI systems are developed and adopted, but also how effectively they can be monitored and governed.

AI adoption pain points

Industry versus regulator adoption pain points: There is a fundamental difference of perception between regulators and the industry in terms of the challenges they are facing in adopting AI.

Regulators emphasise issues around resourcing, especially talent. A lack of AI training to enable current staff to be competent in its use is considered an issue for 48% of regulators surveyed, followed by challenges in recruiting AI talent (47%) and data availability and quality (46%).

The biggest regulator-industry gap is in AI training and capacity building, with 48% of regulators citing this as the top issue, while only 17% of the industry see this as a major challenge. This gap is also significant in terms of the challenge of current technology and infrastructure, with 45% of regulators also seeing technology and infrastructure as a key blocker compared with 28% in industry. For AI vendors, the lack of data availability and quality (66% of respondents) stands out distinctly as the leading AI adoption barrier facing their clients.

5.10: Top six pain points for AI adoption by stakeholder group – industry (n=344), AI vendors (n=145), regulators (n=130) (multi-select)

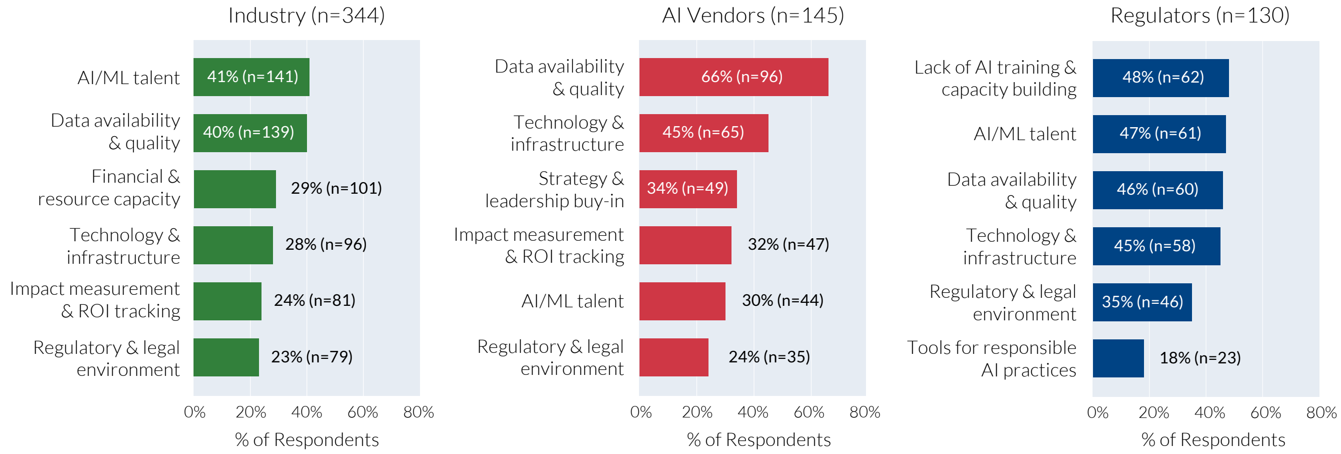

Data-related challenges: perceptions of AI vendors about their clients

AI vendors most frequently identified data quality and completeness as the leading data-related challenge facing their financial institution clients, cited by 72% (n=105) of respondents. Following this, 46% of vendors highlighted legacy systems and siloed data environments, and 41% of respondents see data sharing restrictions, including privacy and data residency laws as key data-related challenges, suggesting that barriers may arise from both formal restrictions on data use as well as internal weaknesses in data governance.

Figure 5.11: Data-related challenges to AI solution development – AI vendors (n=145) (multi-select)

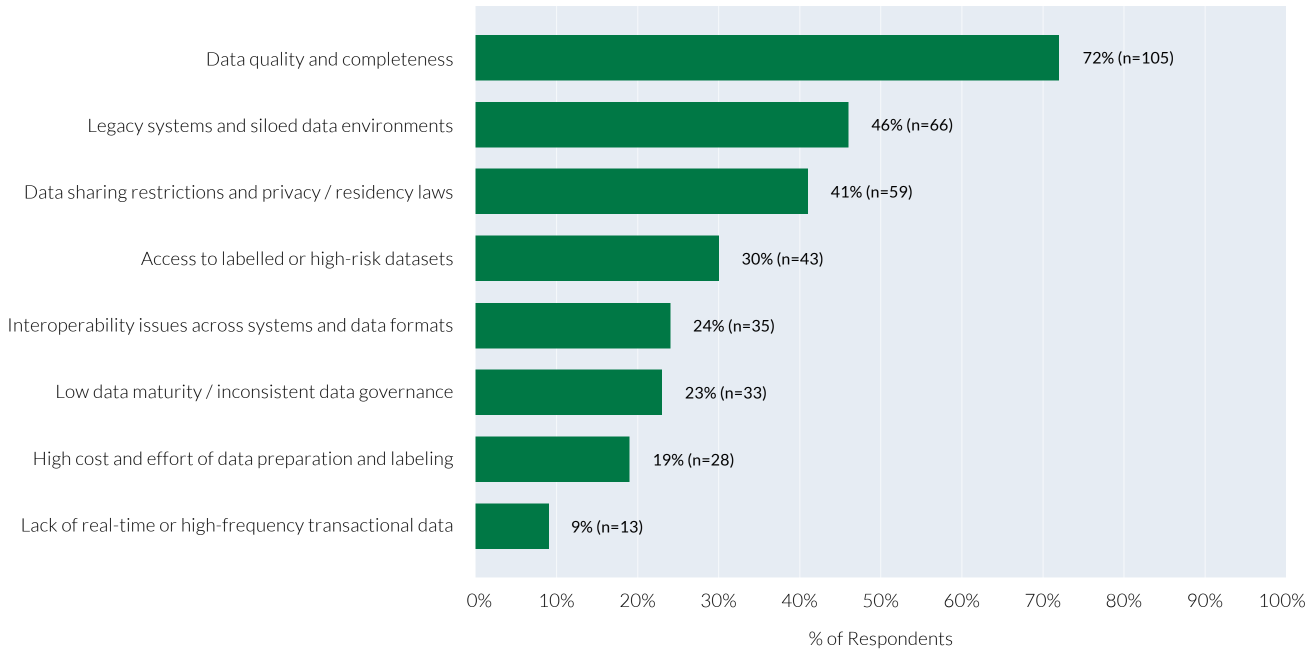

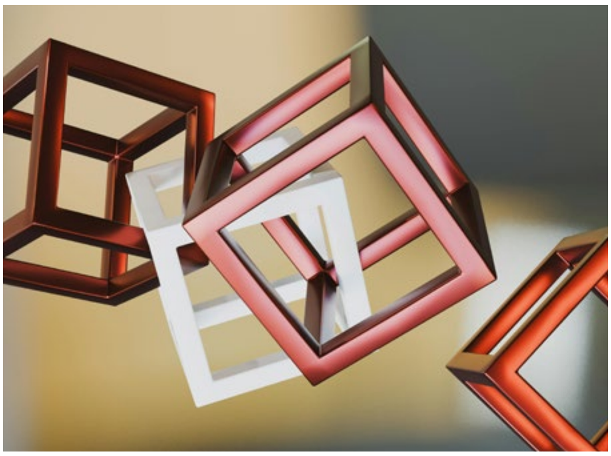

As highlighted in the IMF's report on opportunities and risks of AI in financial services, many AI/ML use cases depend on the availability of large volumes of good-quality, timely data, and shortcomings in data cleanliness, accuracy, relevancy, and potential biases pose material risks for institutions deploying AI systems.[23] Where AI systems are developed and adopted in environments characterised by low-quality data, fragmented architecture, with limited access to data, supervisory authorities may face a more complex monitoring task both to access the data they need.

Partner perspectives: CGAP – Data quality scaling constraints

By Arisha Salman and Camila Quevedo-Vega, CGAP

This study reinforces a persistent reality: data quality remains the single most significant barrier to scaling AI in financial services. These findings are striking not because they are new, but because they are persistent. Despite rapid advances in AI capabilities, the underlying data foundations have not kept pace. CGAP's forthcoming working paper, "Powering AI with Inclusive Data: A Roadmap for Financial Inclusion," argues that this is not incidental: "AI adoption is fundamentally constrained by the strength, inclusiveness, and usability of data systems not necessarily by the availability of algorithms

The constraint is data availability as much as quality. For large segments of the population, particularly women, informal workers, and micro and small enterprises, data trails remain thin, fragmented or entirely absent. Even where digital activity exists, it is often not captured or structured in ways that financial institutions can use. This creates a data availability constraint, limiting the ability of AI systems to generate reliable and generalisable insights.

At the same time, where data is available, it is often incomplete, siloed or not fit for purpose. Weak data foundations, characterised by poor quality, limited interoperability, and gaps in governance, directly constrain the effectiveness of AI systems by reducing model accuracy and reinforcing bias.

CGAP argues that the result is a dual constraint. AI systems are being built on datasets that are both insufficiently available and insufficiently reliable. Advancing towards data-driven financial inclusion therefore requires strengthening both dimensions simultaneously, expanding the availability of data trails while improving their quality, structure, and governance.

Consequently, AI performance and its inclusiveness depend on solving for both at the same time.

The 'connected but invisible' gap is driving data constraints: A central reason these challenges persist is that data gaps are concentrated among underserved populations.

Across many markets, individuals are digitally connected but remain effectively invisible within financial datasets. Their economic lives, often informal, irregular or outside traditional financial systems, are not adequately captured or recognised. This creates a connected but invisible dynamic, where participation in the digital economy does not translate into visibility within data systems.

As a result, financial institutions continue to rely on narrow, traditional datasets that fail to reflect the realities of large customer segments. When AI systems are trained on these datasets, they do not correct these gaps. Instead, they inherit, perpetuate and scale them.

This dynamic is already reflected in broader risks highlighted this survey and in CGAP's work, including bias, exclusion and lack of explainability in AI-driven financial services. These risks are not purely algorithmic, they are rooted in who is represented in the data, and who is not.

Structural barriers continue to limit data usability: Persisting data constraints reflect three structural issues across the financial ecosystem. Together, these barriers create a vicious cycle in which data remains underutilised, even where it exists, and AI adoption remains uneven.

1. **First, data ecosystems remain fragmented**: Data is siloed across institutions, stored in incompatible formats, and difficult to integrate. Even when data exists, it cannot easily be combined to generate meaningful, system-wide insights.
2. **Second, legacy infrastructure constrains data use**: As the survey highlights, outdated systems limit real-time processing and integration of alternative data sources, making AI deployment costly and operationally complex.
3. **Third, governance and trust gaps slow data flows**: Privacy concerns, data-sharing restrictions and unclear accountability frameworks limit how data can be accessed and used. At the same time, regulators face capacity constraints, particularly in technical expertise and in developing standardised approaches to evaluate AI systems.

Reframing AI adoption: from models to data systems: This survey's findings point to the need for a fundamental shift in how the industry approaches AI adoption. The question is no longer how to deploy more advanced AI models, but how to build data systems that make AI viable, reliable, and inclusive.

The persistence of data-related constraints makes one point clear – AI's trajectory in financial services will be determined less by advances in algorithms and more by the availability, quality, and governance of the data systems that underpin them.

Until these foundations are strengthened, data will remain the binding constraint to scaling AI. However, it is also the greatest opportunity. Institutions that invest in building richer, more representative, and better-governed data ecosystems will not only unlock AI's potential. They will define what responsible and inclusive AI looks like in practice.

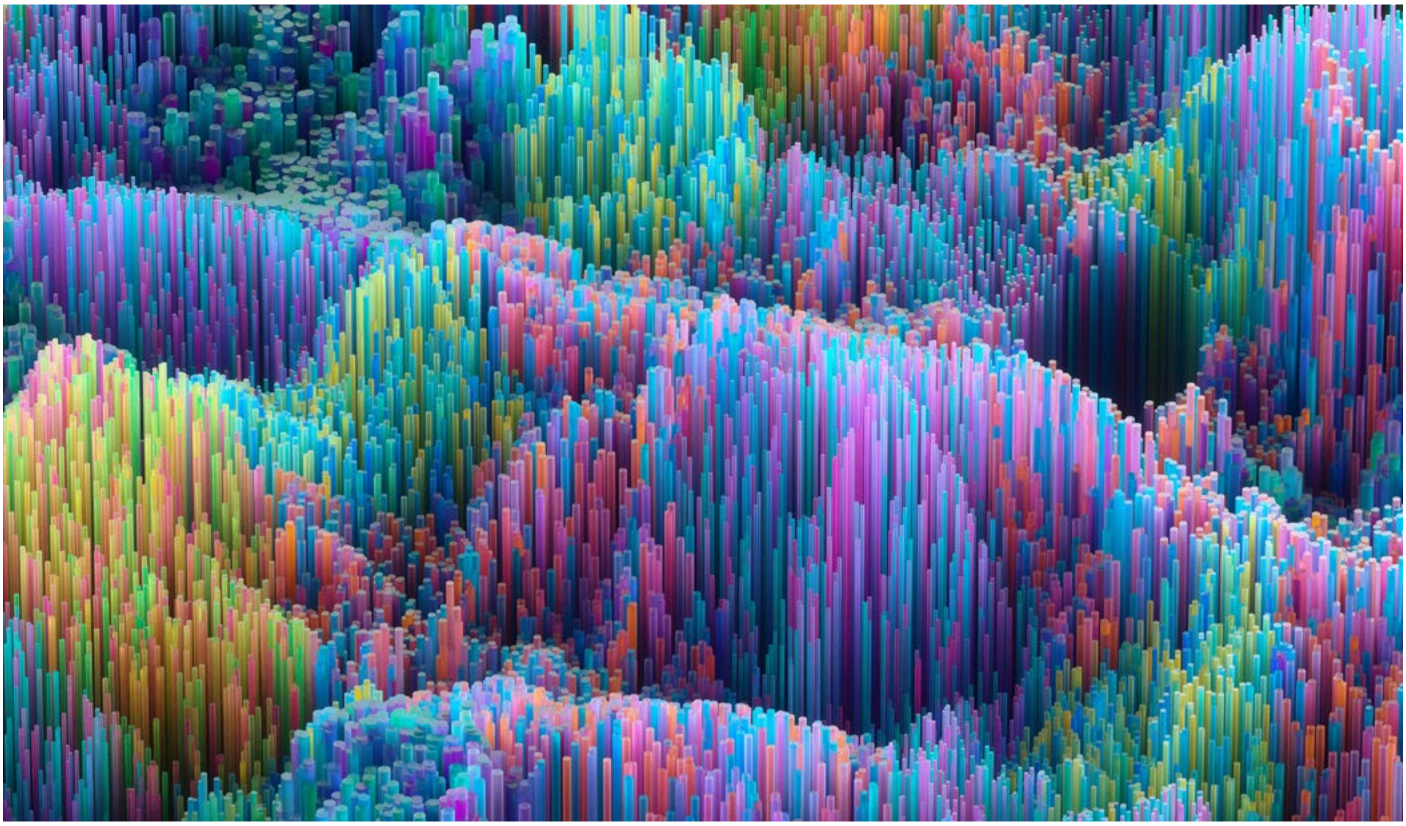

Capacity constraints challenges in regulators

Most regulators face constraints in monitoring AI in financial services, most frequently citing insufficient AI technical expertise, tools, and a lack of standardised metrics. These constraints are considered particularly important by regulators in EMDE jurisdictions when compared with those of AEs. This is particularly evident in the gap around limited expertise or resources (75% versus 57% in AEs), insufficient AI technical expertise (62% versus 41%), and insufficient tools and infrastructure to support effective data collection (48% versus 35%). This could indicate that constraints on effective AI oversight are driven predominantly by capacity and capability gaps. If these constraints are more pronounced in EMDE regulators, they may contribute to a widening gap in AI supervisory readiness between jurisdictions.

Figure 5.12: Regulator capacity constraints in monitoring AI in financial services, by economic development – AEs (n=51) versus EMDEs (n=79)

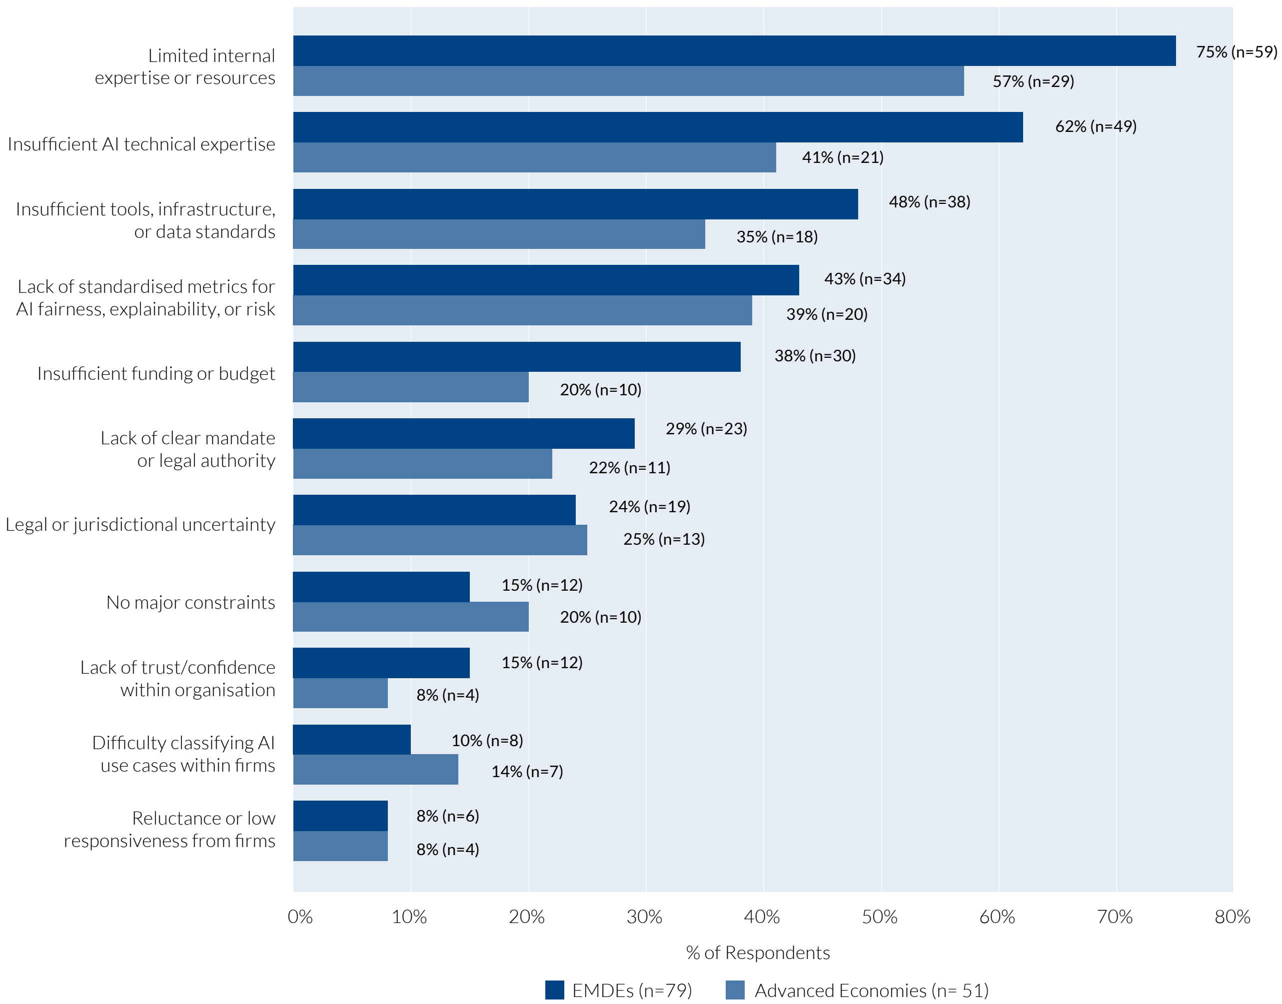

The IMF has stressed that the effectiveness of AI-enabled supervision depends heavily on foundational data readiness, noting that supervisors require strong data foundations to use AI tools safely. It also highlights that resource, and skills gaps represented structural barriers to supervisory adoption, reinforcing that limited expertise and insufficient tools were not unique to specific jurisdictions but reflected systemic challenges across the supervisory landscape.[24]

In summary, the data informing this chapter suggests that AI value in financial services is real but unevenly captured. Respondents broadly report productivity gains while profitability gains are more selective and harder to attribute. The measurement infrastructure to distinguish one source of return from another is, at all maturity levels, underdeveloped. The following sections explore the notion that this gap is not limited to performance management. Firms that cannot reliably observe what their AI systems are doing may also be less equipped to manage what those systems might do wrong. Regulators, many of whom are operating without clear AI mandates in many jurisdictions, may not be well positioned to compensate.

Chapter 6

Risks

The deployment patterns documented in the preceding chapters imply specific risk exposures. Concentration in a small number of AI infrastructure providers generates dependencies that cut across firm types and geographies. Frontier applications, such as autonomous agents, are being deployed into environments. While these AI systems offer unprecedented efficiency, they introduce a systemic chain of technical and market risks, especially Cyber-related risks, that traditional manual oversight can no longer contain. Which also leads to uncertainty around liability allocation. The measurement deficits that make AI value hard to capture, may also make AI risk equally hard to observe.

Although all stakeholders are aware of the opportunities created through the use of AI, there is also broad awareness that it's rapid development is creating a wide range of new risks and exacerbating existing ones. The survey data show that the strongest risk-related concerns are tied to reliability, system control, and data privacy and protection risks derived primarily from cyber-related risk.

As AI systems scale, the risks extend beyond individual firms to macro, interconnected system-level effects. Issues such as cyber and operational resilience, governance gaps, third-party dependencies, and systemic spillovers become more relevant, particularly from a regulatory perspective.

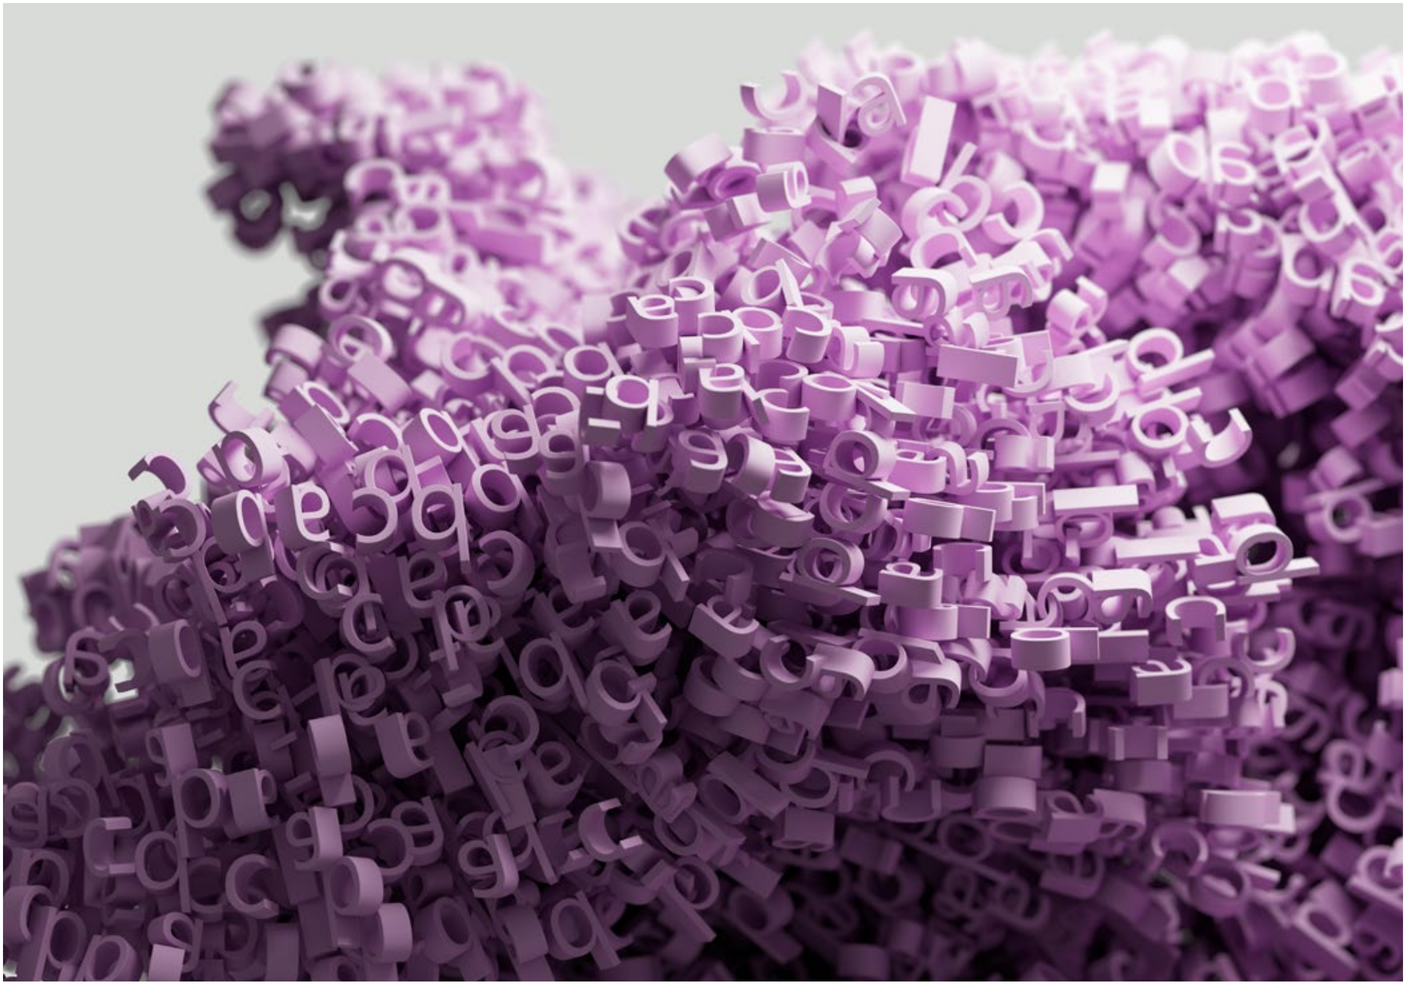

Top combined risks across all stakeholders

Stakeholders broadly align on key AI risks, with data privacy and unreliable outputs leading concerns, alongside a growing risk of reduced human oversight due to automation.

When evaluating technical and financial market risks, a strong consensus emerges across all stakeholders about the most pressing threats. The top two risks highlighted relate to data privacy/protection cited by 73% (n=447) of the total survey respondents to this question and 65% of AI vendors. Of equivalent concern for 69% of respondents (n=424), are risks relating to unreliable results – hallucinations – which are of major concern to the three stakeholder groups: financial industry, AI vendors, and regulators.

A third key risk worth highlighting is around the concept of collective forgetting or loss of human oversight whereby the automation of processes means organisations can lose the institutional memory and capabilities to execute processes manually if required, particularly when crises hit. This is of greater concern to the financial industry (55%) and AI vendors (51%) than regulators (42%).

Figure 6.0: Top Combined AI risks across industry (n=342), AI vendors (n=142), and regulators (n=127) (Technical (T), Financial (F))

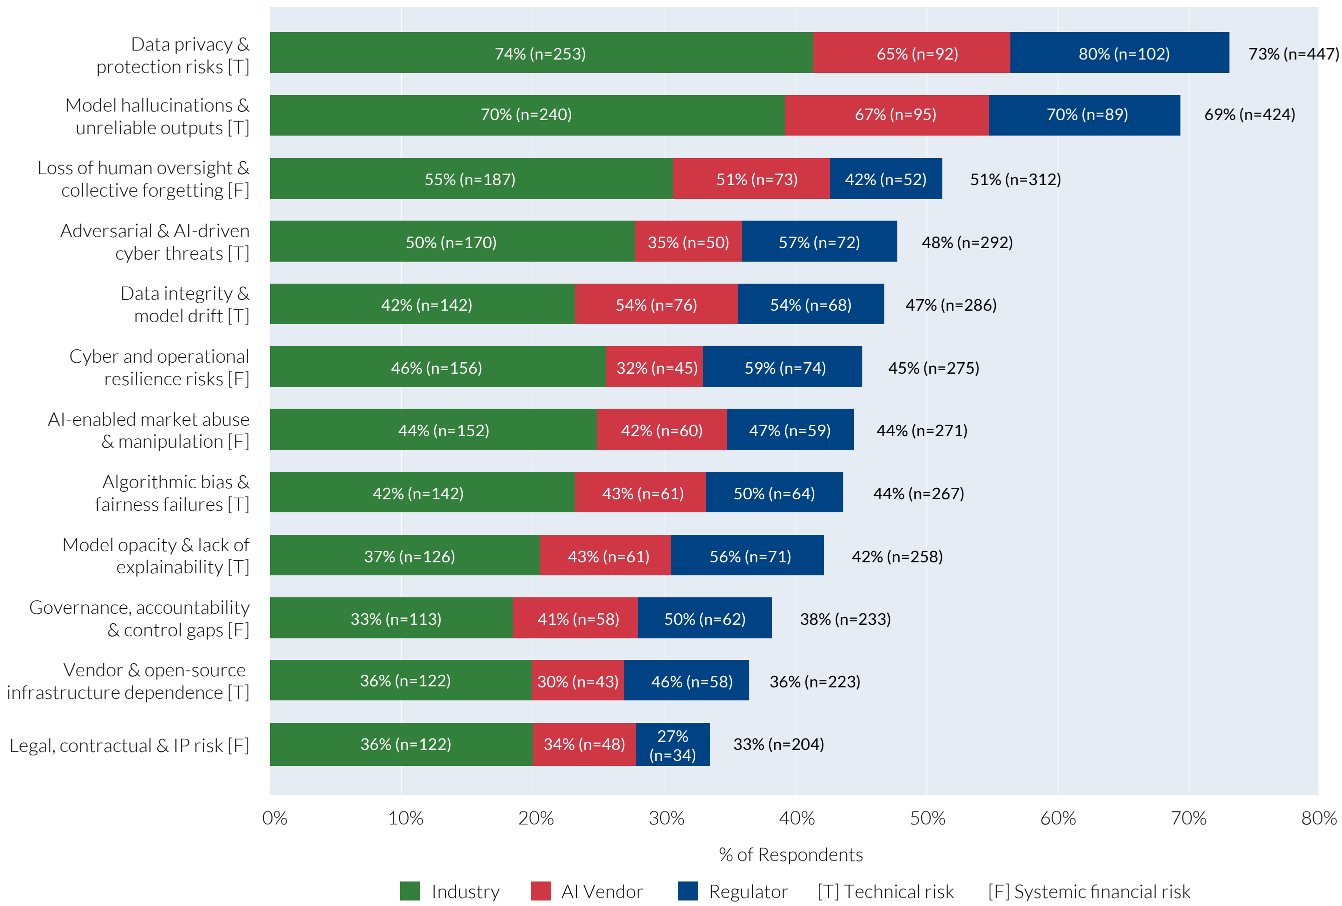

Evolving cyber risks vulnerability and challenges

The rapid deployment of agentic AI (52% of survey respondents) compounds cyber vulnerabilities and is rendering manual oversight increasingly ineffective. Software engineering is the financial industry's most mature AI application (42% fully deployed, 33% in development) and is also a primary cyber risk transmission vector. Financial institutions are moving beyond static tools toward systems that plan and execute actions independently. This transition fundamentally changes the risk landscape, moving from isolated output errors to complex, non-deterministic, multi-step operational failures that unfold in real time.

As reliance on AI grows, institutions gain speed but lose granular visibility and control. Vulnerabilities can be introduced in ways that are difficult to detect through traditional human review. Loss of human oversight and collective forgetting is cited by 51% of respondents, pointing to a reduced ability to intervene when systems behave unexpectedly – the third highest AI risk overall as rated by all respondents. In software engineering specifically, the unprecedented volume and velocity of AI-generated code make traditional manual reviews near impossible. This can also lead to the loss of institutional memory and the ability of teams to step in to execute processes manually should a crisis emerge.

Further complicating this problem space is a notable perception gap: AI vendors place less priority than industry and regulators on both adversarial AI threats (35% versus 50% industry, 57% regulators) and cyber/operational resilience (32% versus 46% industry, 59% regulators).

These differences may reflect roles and incentives but underline the need for better alignment in how risks are understood and managed.

Traditional controls focus on perimeter security, access control, patching, and monitoring known vulnerabilities. However, attackers can now use AI to identify weaknesses faster, generate exploit code, and automate large scale system probing, shifting institutional focus toward defending against AI enabled attacks rather than hardening internal systems. Anthropic's recent 'Mythos'[25] disclosures point to an imminent future where next generations of AI models are set to be incredibly capable of exploiting software vulnerabilities. This presents both firm-level cyber resilience as well as systemic financial risks.

The limits of human-in-the-loop governance mechanisms strain traditional accountability and liability frameworks in financial services. Regulators generally maintain the principle that financial firms should remain accountable for harms, including cyber-attacks, whether AI is built in-house or supplied by third parties, but that position becomes harder to apply in the context of more autonomous systems that are provided and managed by third party vendors.

The survey data shows that regulators are more than twice as likely as industry participants to say that responsibility should remain with the regulated financial institution (38% vs 18%). By contrast, industry respondents and AI vendors are more supportive of shared or joint liability arrangements across parties (22% and 24%, versus 9% among regulators).

Call for action: AI-enabled supervision

In this fast-evolving context, traditional approaches to oversight by regulators may no longer be sufficient. The scale, speed, and autonomy of AI systems make purely human-led supervision less effective. Regulators must therefore adopt agentic AI capabilities themselves to match the systems they oversee. Without embedding agentic supervisory systems that can monitor, test, simulate, and respond to risks at machine speed, regulatory capacity may lag technological reality. Without this transition, the gap between technological advancement and regulatory capacity will continue to widen, increasing the likelihood of rapid, system wide disruptions, risks and potential harm to market participants.

Diverging risks priorities by stakeholder group

Beyond data privacy and hallucinations, risk priorities diverge for each stakeholder group. Regulators are more focused on cyber and explainability, vendors on model performance and bias, and industry on operational resilience and human oversight.

Risk perceptions seem to be driven by each stakeholders' interest and role in the financial sector. Regulators are most concerned about cyber threats (59%) and the lack of explainability (56%). In contrast, most AI vendors consider model drift and data integrity as key issues (54%) as well as a large majority consider algorithmic bias as a concern. Industry respondents tend to consider technical issues such as loss of human oversight and collective forgetting (55%) and cyber threats (50%) as key risks.

Figure 6.1: Top five AI risks by stakeholder group – industry (n=342), AI vendors (n=142), regulators (n=127)

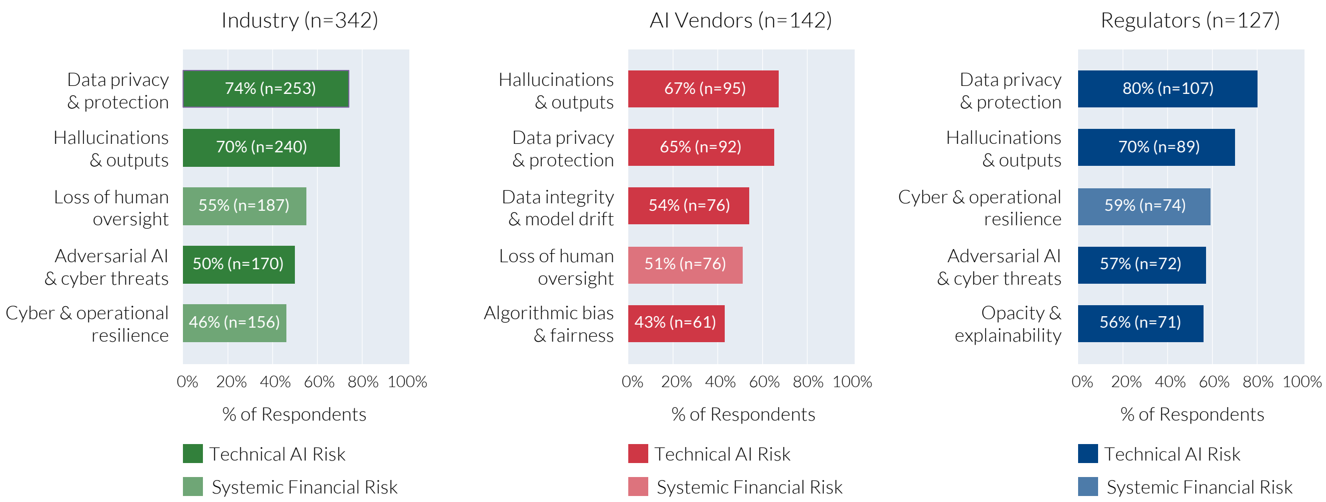

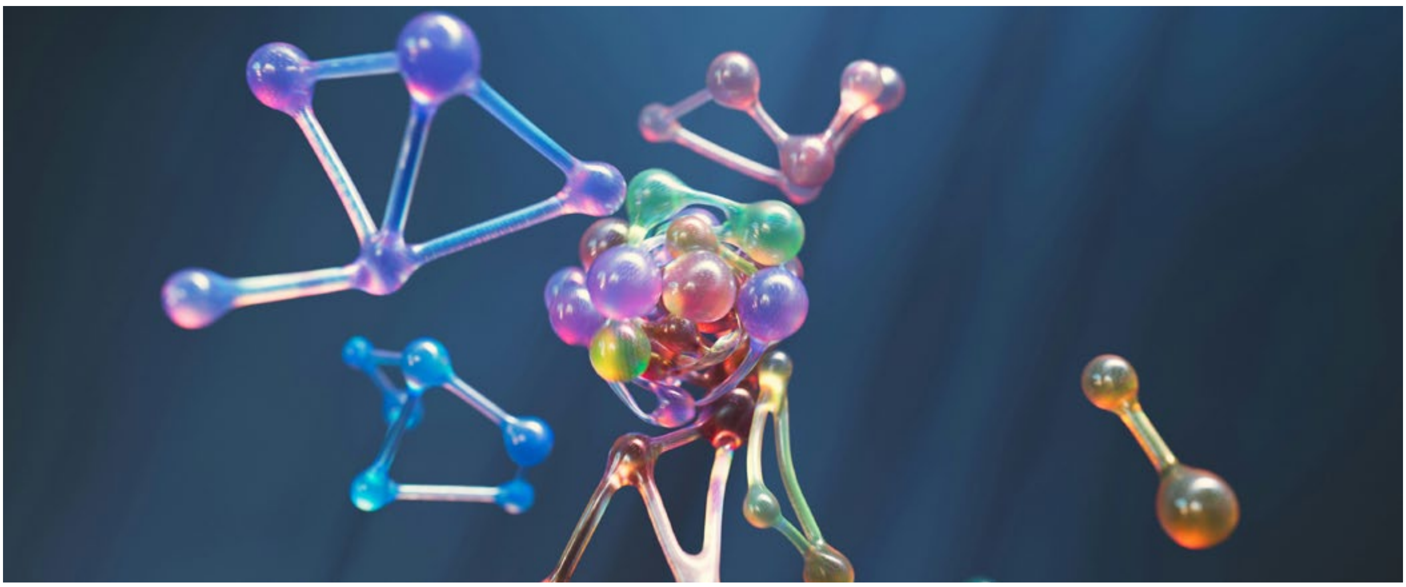

Technical AI risks and the AGI paradox

All stakeholders prioritise immediate technical risks like explainability and model reliability over AGI, despite expecting its near-term emergence, highlighting a gap between current risk focus and future disruption.

The technical AI risks focused on data quality, model reliability and system performance. Within this category, the most pronounced divergence concerns model opacity and the lack of explainability (regulators 56%; vendors 43%; industry 37%). A relatively high proportion of regulators view these as critical risks. This may be because explainability and interpretability of AI model outputs are necessary conditions for regulators to assess compliance and uphold accountability. In contrast, firms and vendors may place less emphasis on fully tracing the model's decision logic provided they are confident and benefit from the accuracy and reliability of outputs.

Interestingly, all stakeholders do not prioritise the risk and emergence of AGI as an immediate threat to financial services (where AI reaches parity with humans across all tasks) instead, placing emphasis on more immediate technical priorities. However, this reveals a paradox: while AGI is not ranked as a top current risk today, 50% of survey respondents expect the emergence of AGI by 2030, implying that this risk emphasis may change very quickly over the next 2-3 years.

Figure 6.2: Technical AI risks by stakeholder group – industry (n=342), AI vendors (n=142), regulators (n=127)

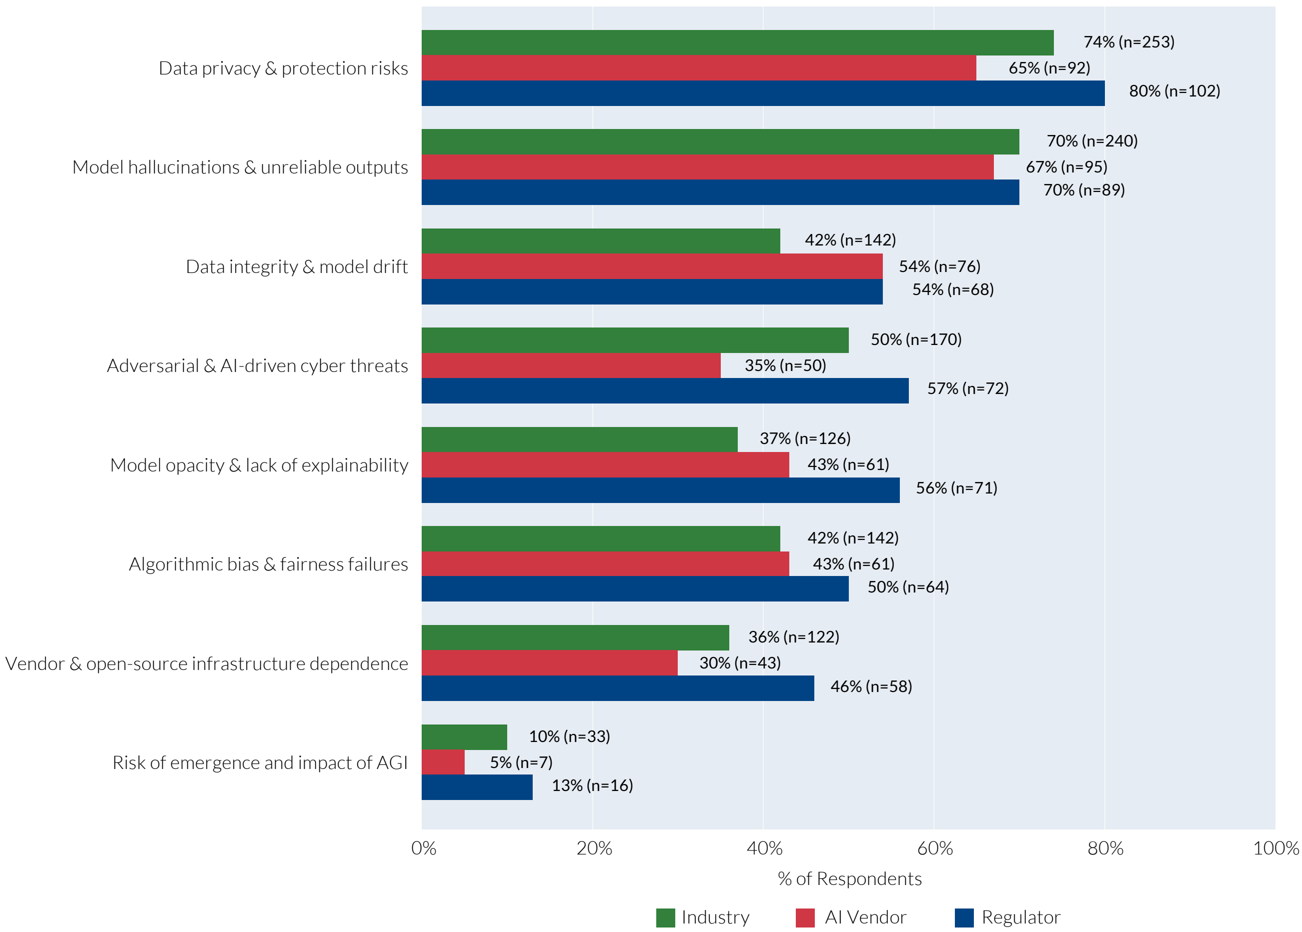

Systemic financial AI risks

Respondents see AI as shifting financial risk from isolated issues to system wide impacts, with concerns around reduced human oversight, market abuse and governance gaps, and reliance on shared third-party infrastructure that can spread disruptions, while longer term risks remain a lower priority.

Figure 6.3: Systemic financial AI risks by stakeholder group – industry (n=342), AI vendors (n=142), regulators (n=125)

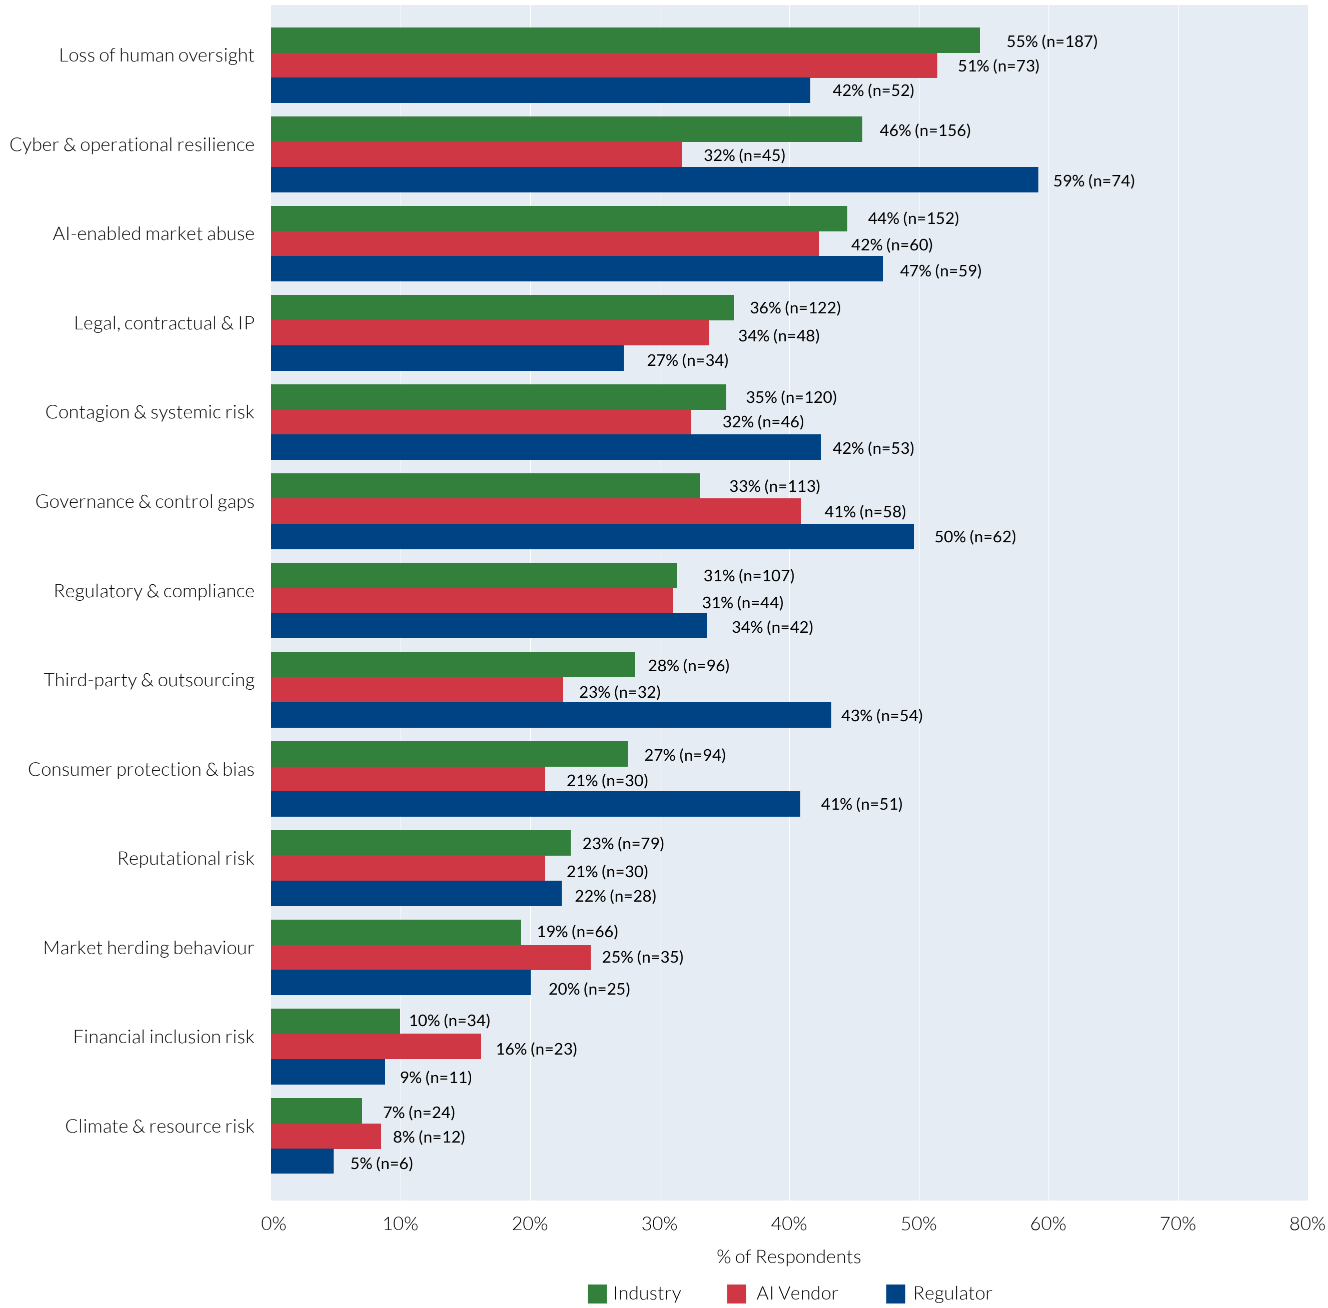

The emphasis on loss of human oversight and operational resilience (industry 55%, vendors 51%, regulators 42%) suggests that respondents see AI as changing how control is exercised in financial products, services and systems. When decision making becomes more automated, errors become harder to detect and may also spread faster across processes, products, and institutions.

The prominence of AI-enabled market abuse, manipulation, (cited by 44% of industry respondents, vendors 42%, regulators 47%) and control and governance gaps (industry 33%, vendors 41%, regulators 50%) in the results, indicates that AI in financial services is perceived as an amplifier of existing financial risks. Automated models can augment deployed trading strategies, pricing decisions or customer interactions in ways that are difficult to monitor in real time, which raises concerns around market conduct and accountability.

Cyber and third-party risks appear consistently across all three stakeholder groups (industry 46%, vendors 32%, regulators 59%) because AI systems often rely on shared infrastructure, external data and vendor ecosystems. This creates concentration points as seen in the data relating to cloud and foundation models providers, where failures or attacks can have wider effects beyond a single institution. Regulators emphasise these risks because they can transmit shocks across the system rather than remain contained (it should be mentioned, however, that only 18% of regulators collect data on third-party AI dependencies as the next chapter documents). The data examined earlier, showing that most firms rely on a small group of foundation model providers, further reinforces the potential concern of concentration-related systemic shocks.

Conversely, broader risks such as financial exclusion, geopolitical risks, and climate-related AI risks (industry 7%, vendors 8%, regulators 5%) indicate that attention is still centred on risks with direct and near-term implications for firm operations and financial stability.

Perception gaps between industry and regulators

Results highlight clear perception gaps, with regulators placing significantly higher emphasis on explainability and governance risks than industry, reflecting their focus on ensuring accountability, transparency and control over AI-driven decisions in critical financial activities.

Regulators and firms seem to diverge substantially in risk perception on some systemic risks.

The chart below highlights some of the key divergences between industry and regulators. The most significant gaps relate to model opacity and lack of explainability, with just 37% of industry compared with 56% of regulators citing this as a top five key systemic financial risk. Most regulators expect greater visibility in the decisions with AI to ensure fairness, especially in high-risk domains such as credit provision and compliance.

A similar gap exists around governance and control where regulators expect firms to be in control of the decision-making capabilities of AI systems they use. This aligns with the regulatory focus on whether clear accountability, oversight structures and control frameworks exist around AI systems. This is particularly so given growing reliance on shared technology providers, such as cloud and AI vendors, where gaps in governance can quickly translate into concentration and system-wide risk.

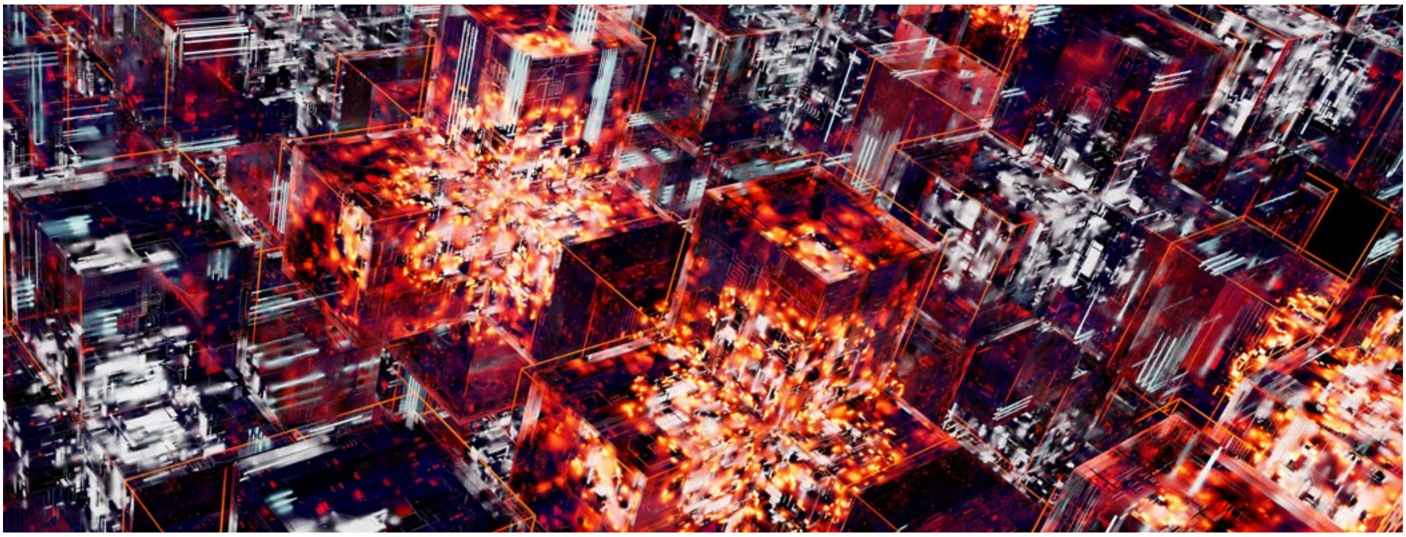

Figure 6.4: Risk perception gaps between industry (n=342 technical / 342 systemic) versus regulators (n=127 technical / 125 systemic)

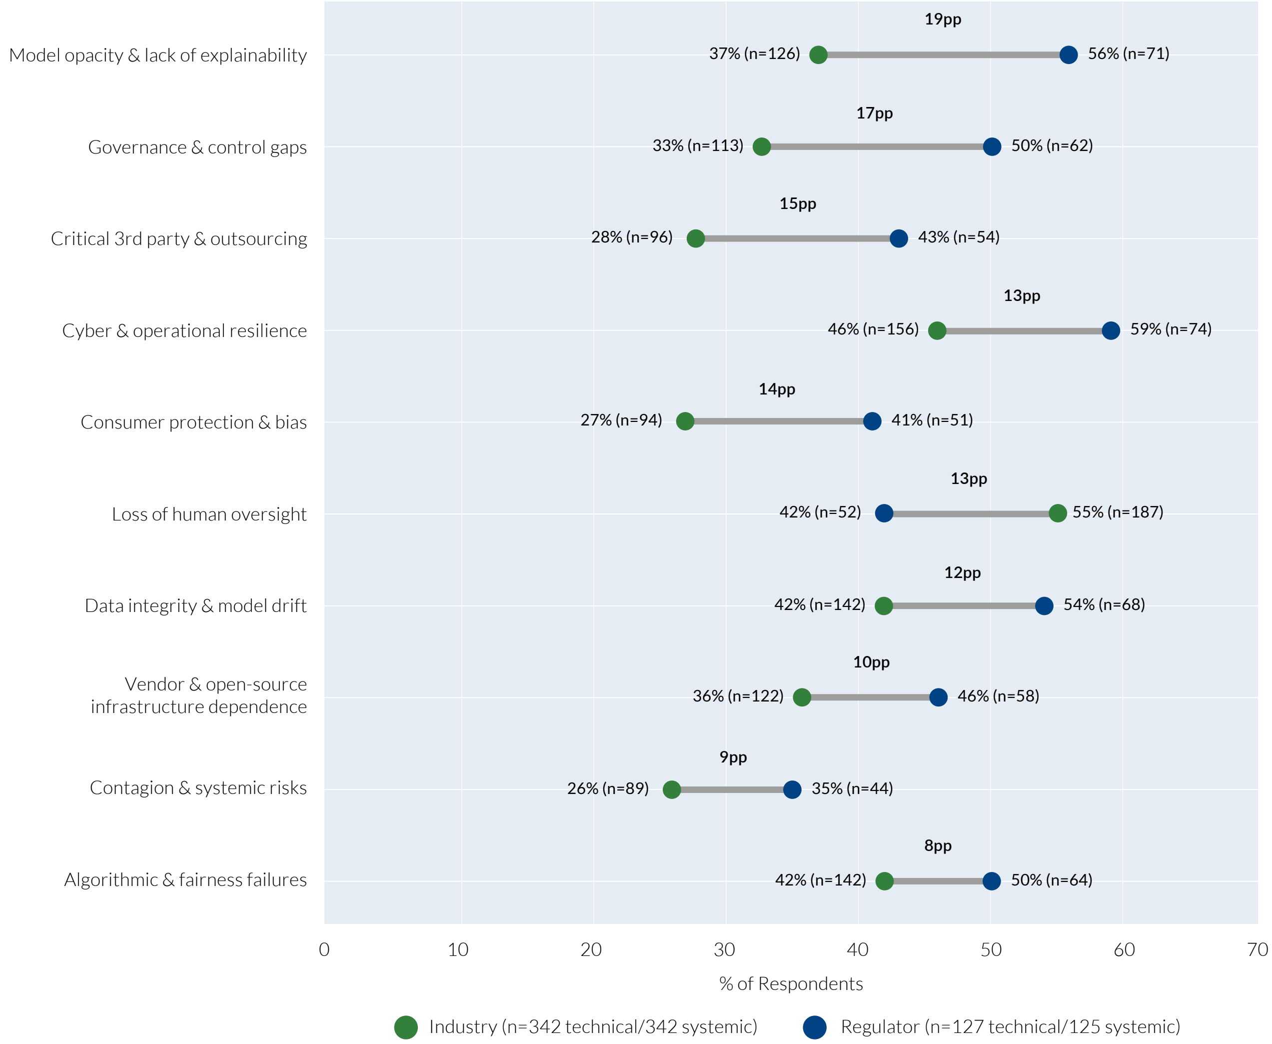

Partner perspectives: CGAP – AI's impact on consumer protection

By Eric Duflos, CGAP

AI brings significant risks as well as potential benefits for consumers: The survey confirms that AI is a double-edged sword for inclusive and responsible financial systems. While 78% of regulators believe AI will be transformative for achieving their objectives by 2030, the findings raise concern across several consumer protection areas.

Respondents ranked AI-enabled market abuse and manipulation as the top systemic risk, with cyber-resilience also near the top. Both industry (74%) and regulators (80%) identify data privacy and protection as a leading concern. CGAP's research on evolving consumer risks in digital finance shows similar issues:

- AI exacerbates those risks by increasing the scale, speed and sophistication of fraud, particularly through deepfakes combined with the accelerating use of social media by consumers and fraudsters alike.
- AI also creates new exclusion risks, as algorithmic biases and hallucinations can systematically shut out vulnerable consumers. The survey reflects this: 70% of regulators and industry actors see hallucinations as a major issue.

Most regulators also flag the lack of explainability in AI-driven financial services as a significant risk, yet only 23% of industry actors are actively addressing it. CGAP's research highlights the same tension with AI models becoming black boxes, which makes it very difficult for regulators to hold providers accountable, or for consumers to challenge AI-generated outcomes that affect their financial lives.

Considerations for financial authorities and other actors in the ecosystem: The risks that AI exacerbates will require ecosystem-wide solutions. As argued in CGAP's Responsible Digital Finance Ecosystem (RDFE) framework, reducing fraud and data misuse, the two fastest-growing risks in digital finance, requires collaboration across the ecosystem, including financial and non-financial authorities, diverse providers, fintechs and consumer representatives. AI will also play a significant role in detecting, preventing and mitigating fraud: half of the successful consumer protection solutions CGAP has identified are AI-powered.

Building the capabilities of ecosystem actors is equally critical. This means supporting authorities in adopting AI-powered tools, strengthening industry players' defences against AI-powered fraud and improving consumers' ability to protect themselves. Attention is also needed to close the knowledge and resource gap between EMDEs and high-income countries on AI adoption, as leapfrogging cannot always be assumed. This will come with costs and culture change.

Finally, AI raises important questions about the relationship between data protection and consumer protection. We found that AI can greatly help protect consumers from fraud but often at the cost of increasing data privacy and protection risks. This tension deserves further exploration and calls for closer collaboration between consumer protection and data protection authorities as well as industry. When AI misuses consumers' data, 'liability' gets diluted across multiple actors in the ecosystem (AI developers, providers, deployers and users). Authorities need to develop rules that establish who is responsible when things go wrong (from data breaches, data misuse and data misappropriation) to ensure good protection for consumers.

The explainability and interpretability challenge

The study results highlight that explainability and interpretability are essential for regulators to assess fairness and enforce accountability in AI-driven financial decisions. However, there is a clear gap between regulatory expectations and industry capability, with limited adoption and expertise in explainability methods by industry.

Explainability and interpretability are important to financial regulators. There is no consensus in what defines the two overlapping concepts, but generally, 'interpretability' refers to how well we understand the inner workings of AI models, while 'explainability' deals with how the stakeholders, including lay people and non-technical experts, understand the results and reasoning outputs of AI systems.[26]

In highly regulated environments like financial services, organisations must often need to be able to explicitly explain why a specific decision was made (for example, a loan denial or a trade execution) to regulators. This drives their strong preference for predictable, deterministic models over the non-deterministic, probabilistic nature of most large AI models, which inherently carry a level of variance and unpredictability.

In cases where AI models affect financial decisions such as credit, compliance and fraud detection, regulators cannot assess fairness and enforce accountability if those decisions cannot be explained. Consequently, 79% (n=98) of regulators stated AI model explainability and interpretability are important or critically important for achieving their regulatory objectives.

Notably, 46% cited that they are important but they are not the only factor, indicating that regulators generally view them alongside other factors when dealing with AI systems. This aligns with studies that see explainability as only one of the operational trade-offs faced by financial institutions.[27]

Figure 6.5: Regulator views on the importance of explainability and interpretability – regulators (n=125)

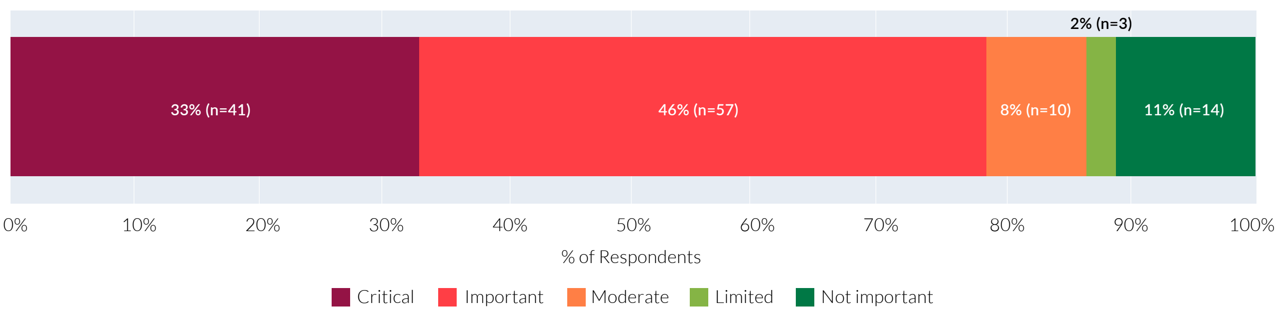

Whether interpretability is really feasible is also a notorious debate. Complex AI models, such as deep neural networks, are known to perform better than traditional regressions, but at the cost of opacity or lack of transparency, with critics arguing for using only inherently interpretable models.[28] This tension possibly drives the critical mismatch existing between surveyed regulators and vendors.

As a response to the opaque processes of black box models, explainable AI or XAI models have been increasingly studied and applied in finance. Most of these explainability methods are called surrogate models that estimate the model behaviour of black boxes (for example, SHAP and LIME.[29]

A substantial disconnect exists between regulator expectations and industry capabilities. Only 50% of surveyed financial institutions use explainability methods while the other 50% do not use explainability methods or are unsure if they do.

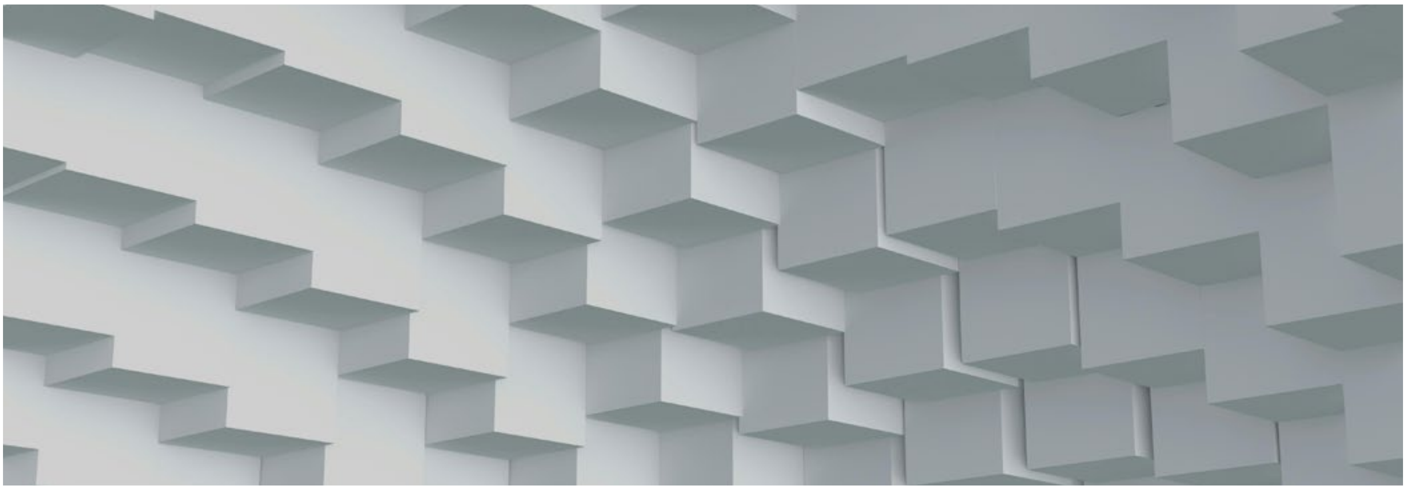

Figure 6.6: Industry use of explainability methods for AI models – industry (n=340)

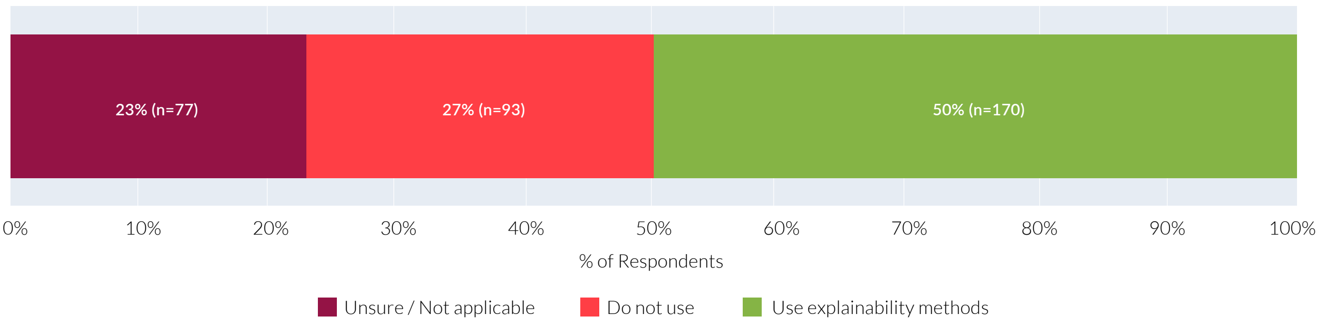

Furthermore, AI vendors indicate that their financial industry clients have limited capacity to use explainability methods, with over 50% rated as having low or no expertise in applying these techniques. This highlights the need for more human-centric approaches to AI explainability as existing XAI techniques require sophisticated technical expertise to be used which may not be relevant or feasible for financial domain experts or the end users affected by AI decisions.[30]

Figure 6.7: Client expertise in using AI explainability tools – AI vendors (n=131)

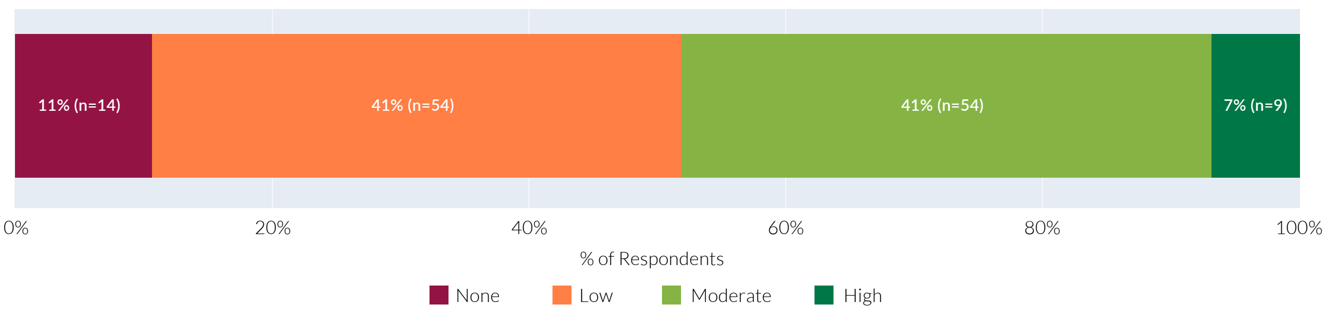

Partner perspectives: BIS – AI explainability with Project Noor

By Bénédicte Nolens and Jack Lee, BIS Innovation Hub

Project Noor is an initiative of the BIS Innovation Hub,[31] developed in collaboration with the Hong Kong Monetary Authority (HKMA) and the Financial Conduct Authority of the United Kingdom (FCA). The project seeks to equip financial supervisors with independent, practical tools to evaluate and interpret the inner workings of AI models used by banks and other financial institutions. By combining explainable AI (XAI) methods with risk analytics, the project aims to deliver a prototype through which supervisors can verify model transparency, assess fairness, and test robustness.

Project Noor could enable greater transparency in AI adoption and encourage responsible innovation, allowing financial institutions to adopt new technologies with practical, privacy-preserving checks. It is important to note that financial institutions retain responsibility for model explainability and that Noor does not aim to prescribe definitive standards or replace existing practices. Instead, Noor strives to equip supervisors with methods and benchmarks to form their own informed opinions.

Insights from Noor engagement with banks/ regulators: Under Project Noor, a series of structured workshops were held, comprising central banks, banking supervisors, insurance supervisors and financial institutions. During the workshops, participants provided their views on Project Noor. Many of the findings from these workshops echo the survey results.

- **Caution in high-stakes deployment**: Workshop participants noted there is a notably low risk-appetite for deploying generative/agentic AI in high stakes use cases. Workshop participants also noted that banks generally remain cautious as some supervisors have not yet provided detailed guidance or standardised frameworks for these advanced models. This echoed findings from the survey that 'clarify and update AI guidance' is a top priority for the industry and regulators.
- **Explainability as an operational challenge**: Some regulators and financial institutions noted a degree of maturity in the use of XAI, but noted that these practices are not standardised. There is a lack of a common 'supervisory language'. Participants expressed that Noor may serve as a starting point for a structured, repeatable assessment framework.
- **Human-in-the-loop as an anchor**: Consistent with the survey's finding that human-in-the-loop approaches being critical, workshop participants emphasised that XAI tools (and by extension, Noor) should not replace but instead empower humans. Noor aims to provide the 'technical diagnostics' that allow human supervisors and risk officers to provide meaningful, rather than perfunctory, oversight.

Project Noor seeks to better understand these dynamics and explore how explainability tools could support supervisors in managing these emerging risks.

The CCAF survey results provide a valuable global perspective that complements the qualitative insights from Project Noor. Together, they highlight several areas for further work. These include the governance of third-party AI providers, clearer allocation of responsibilities across the AI value-chain, and the implications of increasingly advanced agentic AI systems for supervisory frameworks.[32]

The BIS Innovation Hub has developed a suite of projects leveraging artificial intelligence across use cases. In financial crime, Project Aurora applies AI, machine learning and privacy-enhancing technologies to enable cross-border collaborative analytics for improved money laundering detection, while Project Nadim explores AI-driven sharing of fraud patterns to strengthen international detection efforts. In climate and sustainability, Project Gaia uses AI and large language models to extract and standardise climate risk data from disclosures, and Project Symbiosis applies AI and big data to enhance supply chain sustainability insights. In monetary policy, Project Neo applies advanced data science and machine learning to novel high-frequency data for real-time economic insights, and Project Spectrum combines generative AI and machine learning to improve large-scale price classification for inflation nowcasting.

AI Accountability and Liability

A clear divide on accountability for AI-related harm exists, with regulators placing primary responsibility on financial institutions, while industry and vendors favour shared or joint liability models, reflecting differences in how responsibility for AI outcomes should be assigned and enforced.

Figure 6.8: Primary accountability for AI-related harm by stakeholder group – industry (n=334), AI vendors (n=134), regulators (n=130)

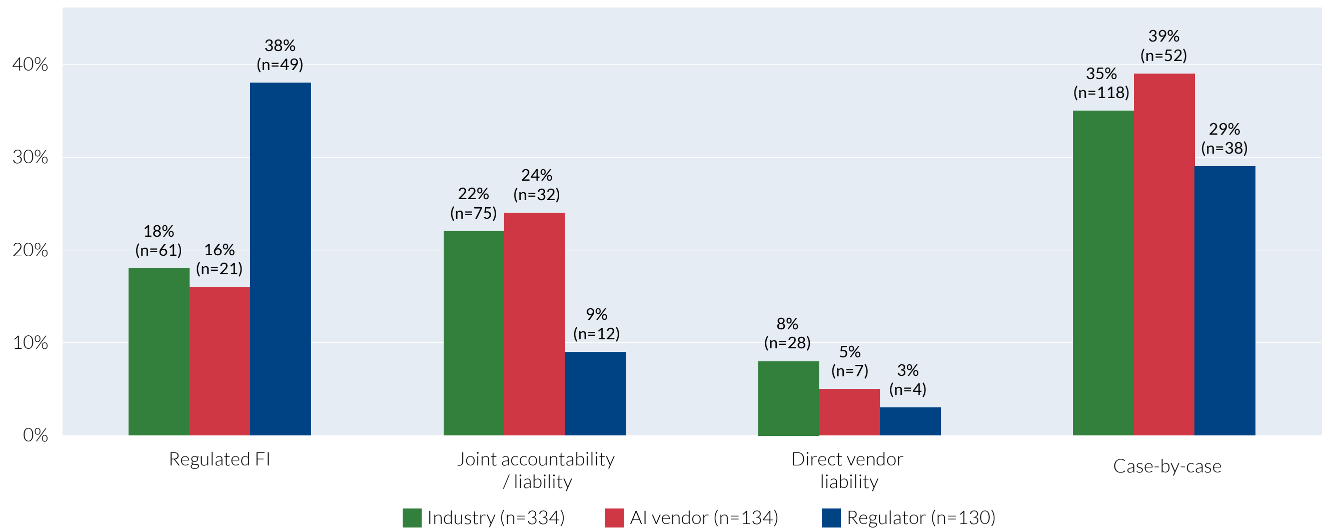

Figure 6.9: Primary accountability for AI-related harm: industry (n=334) versus regulators (n=130)

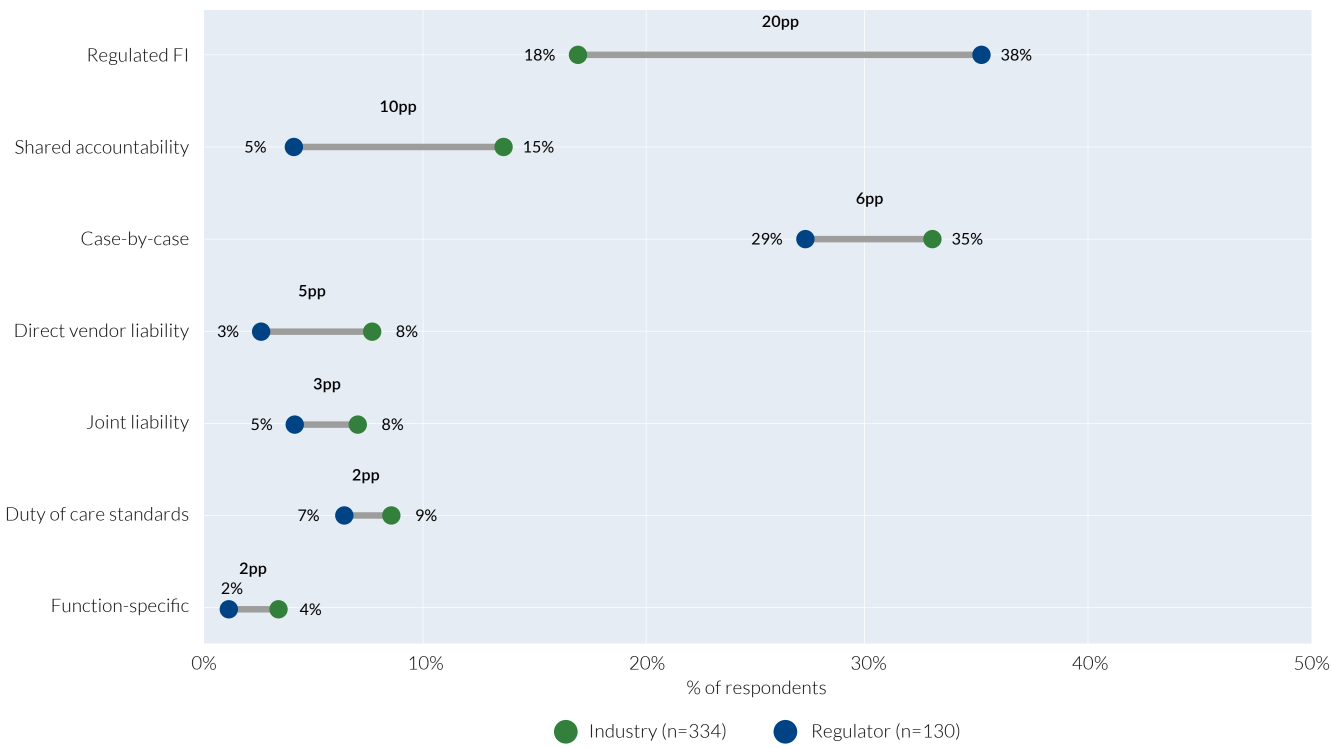

The regulatory stance is consistent with the longstanding principle in financial regulation that firms remain responsible for the services they deliver and any related harms, regardless of whether AI is developed in-house or by an external provider.

When AI models are supplied by third parties, the financial industry is expected to understand, monitor and challenge their providers, under outsourcing rules. Recent regulatory initiatives, including regimes that place critical third-party providers under financial authority oversight (for example, the EU's Digital Operational Resilience Act), do not change this principle, in fact, they are intended to strengthen it by increasing the obligations of vendors towards financial institutions.

By contrast, industry respondents and AI vendors show greater support for hybrid forms of accountability, and, importantly, for distributing resulting liability. The survey distinguishes between shared, joint and case-by-case liability. Shared liability means each party is responsible only for its share of the harm, while joint liability means any party can be held fully responsible for the entire harm. Case-by-case liability refers to situations where allocation of responsibility is based on the specific circumstances and is typically defined contractually.

Industry respondents and vendors favour shared or joint liability arrangements (22% and 24%, compared with just 9% among regulators). However, a significant share across all groups (29 to 39%) selected 'case-by-case', suggesting widespread recognition that AI use cases differ in complexity and risk and may require flexible liability arrangements.

It should also be noted that even where accountability and direct regulatory liability for AI-enabled decisions remain with financial institutions, financial firms may still seek indirect recourse from vendors. This may be enabled through product liability regimes or private contractual arrangements. Further research would help clarify how these accountability and liability relationships evolve as AI adoption deepens.

With the increasing complexity of AI technologies and increasingly complex AI supply chains, accountability will be a key issue to address, requiring careful attention to the challenges of assigning responsibility, monitoring system behaviour, and evidencing and defending outcomes.

Lower levels of regulator support for shared accountability or co-responsibility frameworks should also be interpreted cautiously. This hesitation may not reflect a rejection of the concept itself, but uncertainty around how such frameworks would operate in practice and whether they align with existing supervisory mandates. Financial sector regulators are typically structured to place primary responsibility and liability on regulated institutions, including accountability to clients and counterparties. Depending on their design, shared accountability models could blur these established lines of responsibility, potentially weakening clarity around institutional liability rather than strengthening oversight. As such, regulator scepticism may stem from concerns about legal coherence, enforceability, and mandate compatibility.

Critical third-party dependency

Results highlight increasing concentration and vertical integration among cloud and foundation model providers, raising concerns around dependency on a few dominant players and potential risks to competition, resilience, and technological sovereignty.

The concentration of providers across the AI supply chain is another potential source of risk, particularly where cloud infrastructure and foundation models are tightly integrated within the same technology stack. There is significant variation in concentration levels across major cloud and foundation model providers, the survey shows. Google stands out, with 88% of respondents using Google Cloud also relying on Google's foundation models. This is more than 2.5 times Microsoft (35%) and nearly four times Amazon (23%).

Figure 6.10: Cloud-to-foundation model provider dependence by stakeholder group

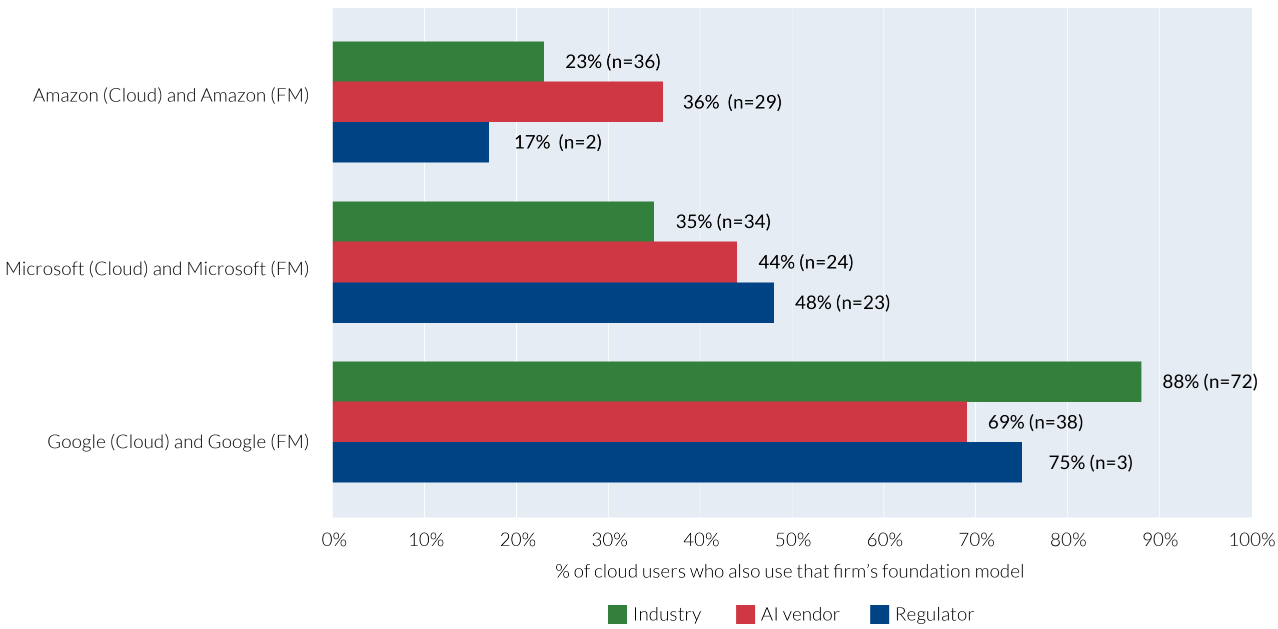

A breakdown by respondent type shows that in Google's case, cloud-foundation model integration is higher in the industry (88%) than among regulators (75%). For Microsoft, the pattern reverses, with the financial industry reporting 35% cloud-FM integration, against 48% of regulators.

Vertical integration raises concerns about growing dependencies on a small number of dominant providers, particularly from the United States. As AI adoption accelerates, regulators and competition authorities, particularly in jurisdictions where dominant providers are not headquartered, may seek to curb concentration risks and preserve contestability across the AI supply chain. This is increasingly intertwined with concerns around data and AI sovereignty, as governments look to ensure strategic autonomy over critical digital infrastructure, data and model capabilities.

The findings in this chapter show that AI risks in financial services are widely recognised, but with clear differences across stakeholders in how they are prioritised, managed and assigned responsibility.

Building on these insights, the next chapter examines how regulators are responding to these challenges, including the role of AI in achieving regulatory objectives, the evolving and fragmented global regulatory landscape, approaches to monitoring risks such as bias and discrimination, data collection practices, and emerging priorities for AI oversight across jurisdictions

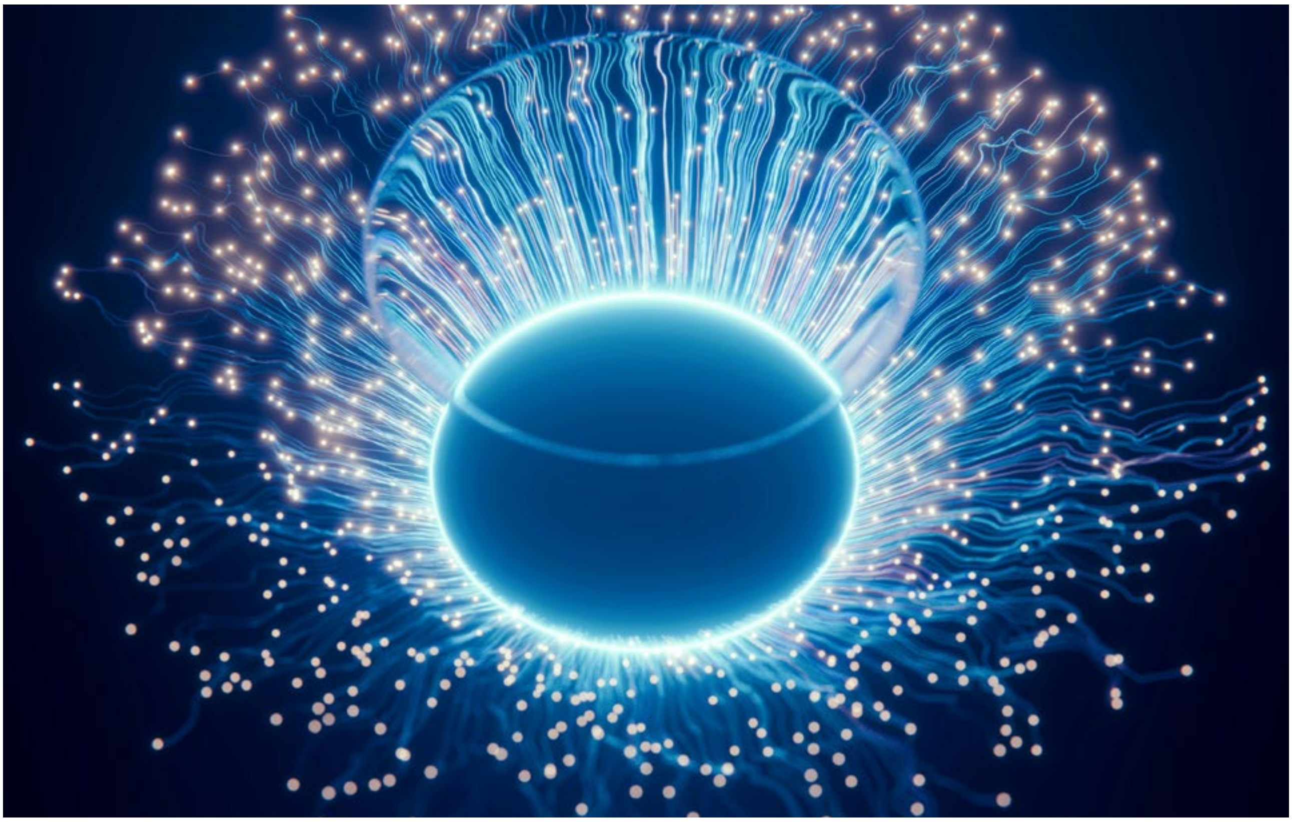

Chapter 7

Regulation and Use of AI by Regulators

AI regulation and policy developments

Whether the risks documented in the preceding chapter can be narrowed depends substantially on regulatory capacity. This chapter examines what the regulatory environment looks like in practice: the frameworks adopted, the supervisory data being collected, and how regulators are deploying AI in their own operations.

Regulatory responses require coordination between financial regulators and government departments, both domestically and internationally, in ways that existing supervisory frameworks were not designed to accommodate.

This complexity is also coupled with uncertainty. It is not clear to regulators and supervisors what AI use cases are being pursued by industry, nor the potential emergent risks for which they will need to craft responses to. Most of the surveyed regulators are not yet collecting data on AI adoption, which may further limit their visibility into the potential risks.

For industry and AI vendors, there is a general lack of guidelines on the use of AI from their respective financial regulators. Nearly 50% of the surveyed authorities do not have a national AI strategy in place. Even those that do, the guidelines are general or broad rather than specific to AI in financial services. However, there is unanimous recognition from all surveyed stakeholders on the need to clarify and update AI guidance with 79% identifying it as the top priority for regulators and supervisors to focus on.

Europe leads while Latin America and the Caribbean (LAC) lag: Europe clearly leads on developing its more principle-based regulations with 59% of its respondents reporting active frameworks, reflecting the momentum of the comprehensive European Union AI Act.

Mixed results from Asia-Pacific (APAC): The APAC region is highly mixed – with 42% reporting no rules and 37% having them – mirroring a diverse environment that ranges from strict, targeted algorithm laws in China to agile, voluntary guidelines issued by regulators like the Monetary Authority of Singapore (MAS).

Strong emerging adoption in some regions: Meanwhile, the Middle East and North Africa (MENA) and Sub-Saharan Africa (SSA) show strong emerging adoption, with over half of their respondents (62% and 53%, respectively) reporting established regulations driven by foundational frameworks like Rwanda's National AI Policy.[33]

Regulatory optimism despite fragmentation: Despite these complex challenges and fragmented rules, regulators are broadly optimistic about the potential impact of AI adoption, both within their own organisation as well as the use of AI by industry, on helping them to achieve their regulatory objectives.

AI Importance for achieving regulatory objectives

78% regulator respondents consider that the use of AI can be transformative or significantly impact their regulatory objectives. This percentage is largely consistent among central banks (82%) and other regulators (75%) surveyed.

Figure 7.0: Importance of AI for regulatory objectives by 2030 – regulators (n=124)

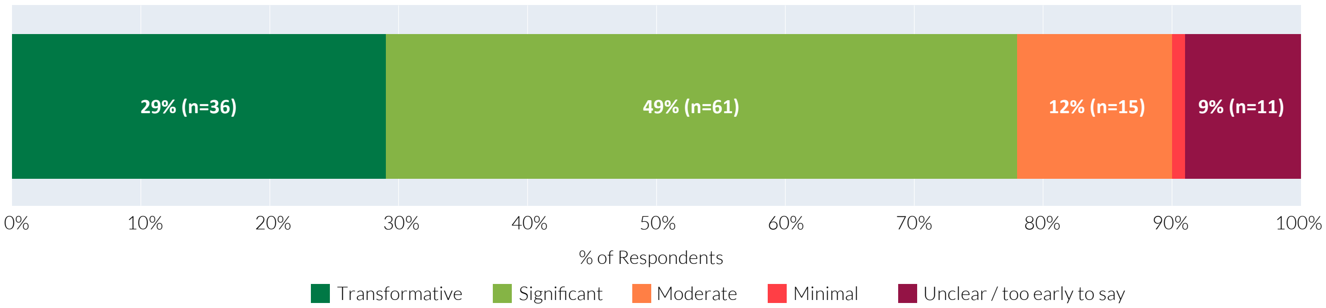

Advanced economies (AEs) versus emerging markets and developing economies (EMDEs): Interestingly, regulators in EMDEs are far more likely to consider the use of AI in financial markets to be transformative for their impact on their regulatory objectives (36%) than regulators in AEs (17%). This is despite the finding that surveyed AE regulators are ahead in AI adoption (33% in Scaling or Transforming phase and 43% in Piloting) compared with surveyed regulators at EMDEs (10% in Scaling or Transforming and 27% in Piloting). This optimism could be due to the ambition for AI applications to be used to help regulators in EMDE jurisdictions to leapfrog existing legacy, expensive regulatory and supervisory technologies currently available on the market. It is also possible EMDE regulators view AI as a means to supplement limited human supervisory resources, which may be more constrained than those available to AE regulators.

Figure 7.1: Importance of AI for regulatory objectives by economic development – Advanced economies (n=47) and EMDEs (n=77)

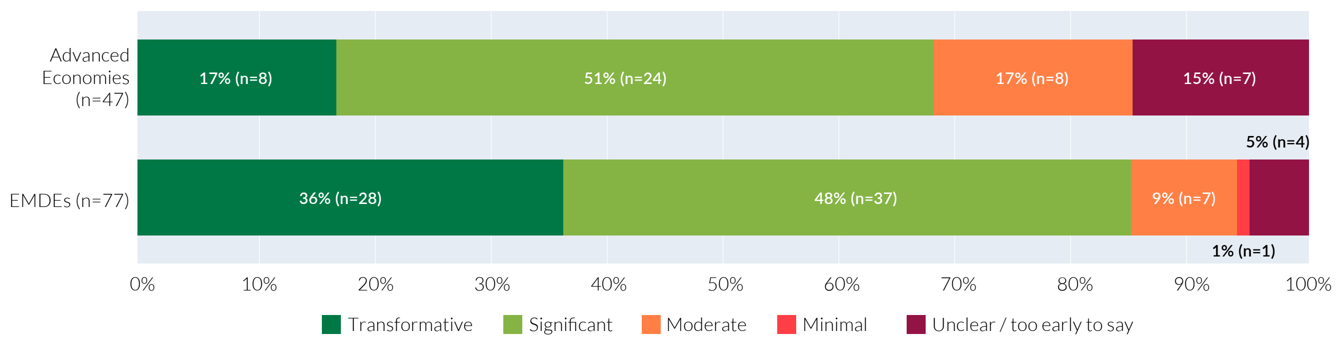

Regulators’ views on impact of AI usage to their regulatory objectives

Views on impact of industries’ adoption of AI: Regulators see the use of AI by the financial services industry as broadly supportive of their regulatory objectives. However, there are a few areas which seem to be a potential challenge, with 27% of regulators seeing cyber and technology resilience being a challenge. 27% also see increasing challenges to consumer protection because of the use of AI by industry. Challenges around market conduct, AML/financial crime, competition and international coordination also raised concerns.

Figure 7.2: Regulators’ views on the impact of industry AI use on regulatory objectives

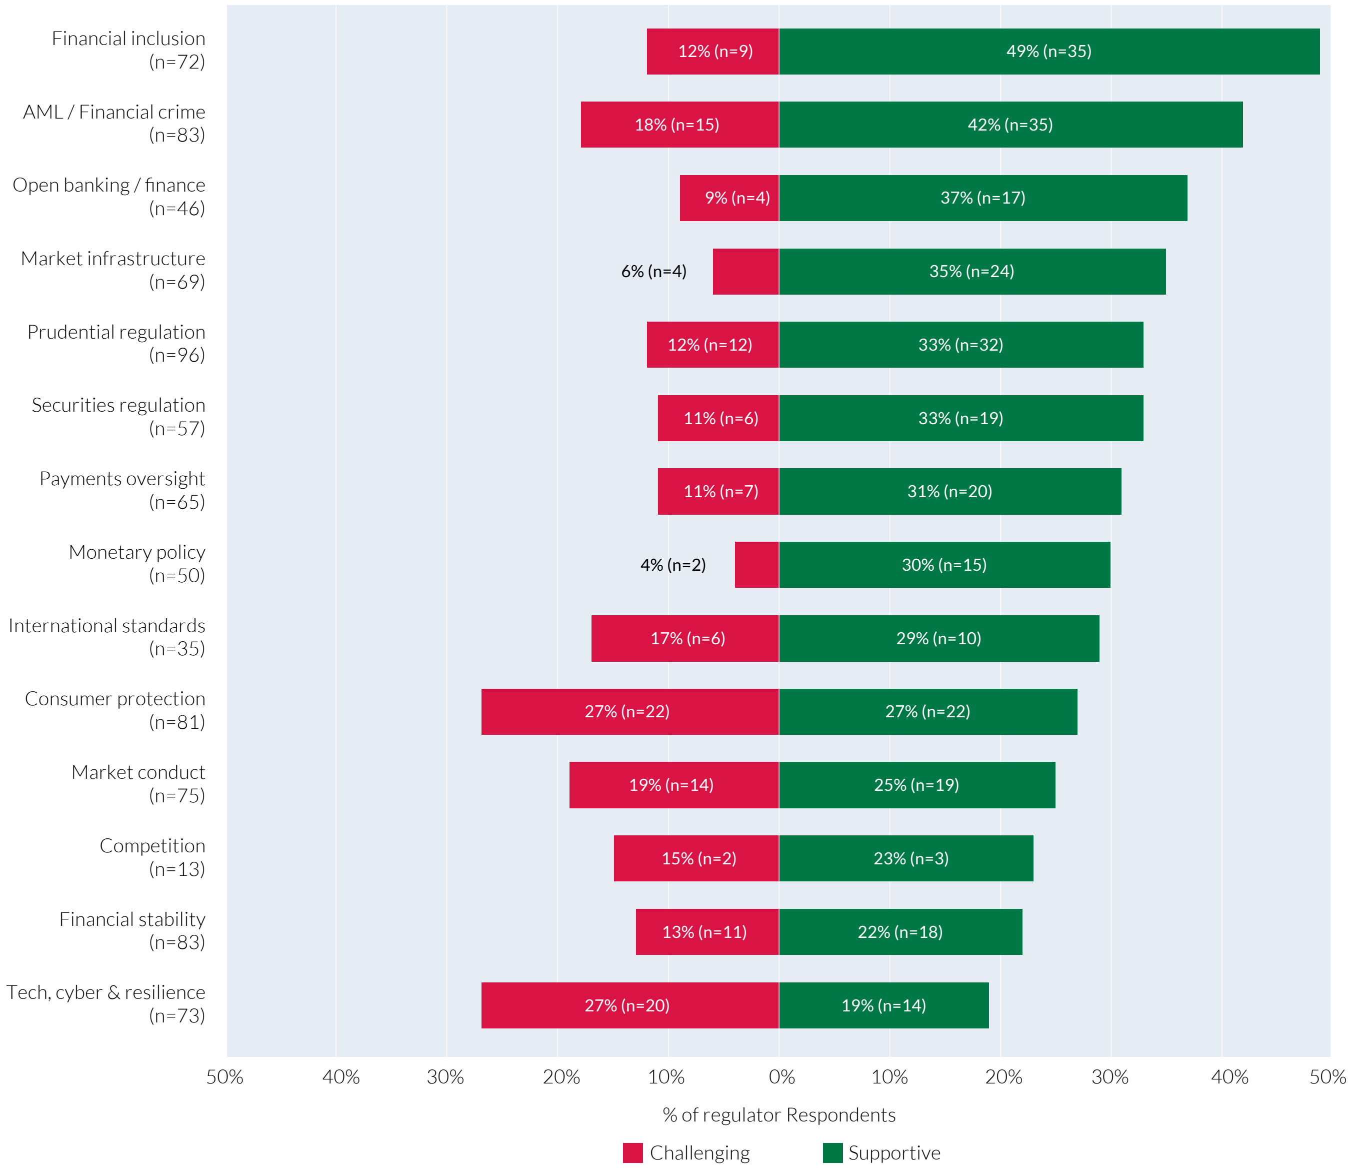

Views on impact of regulators' own adoption of AI: In terms of regulators' own use of AI in challenging or supporting them in achieving their regulatory objectives, the trend across different regulatory areas was far more consistently positive across the board. The only exception was for tech and cyber resilience (with 14% of regulators) and AML/ financial crime (with 10% of regulators) stating this could be a challenge. These are both understandable given the increased risks that AI enables for these regulatory areas.

Figure 7.3: Regulators' views on the impact of internal AI use on regulatory objectives

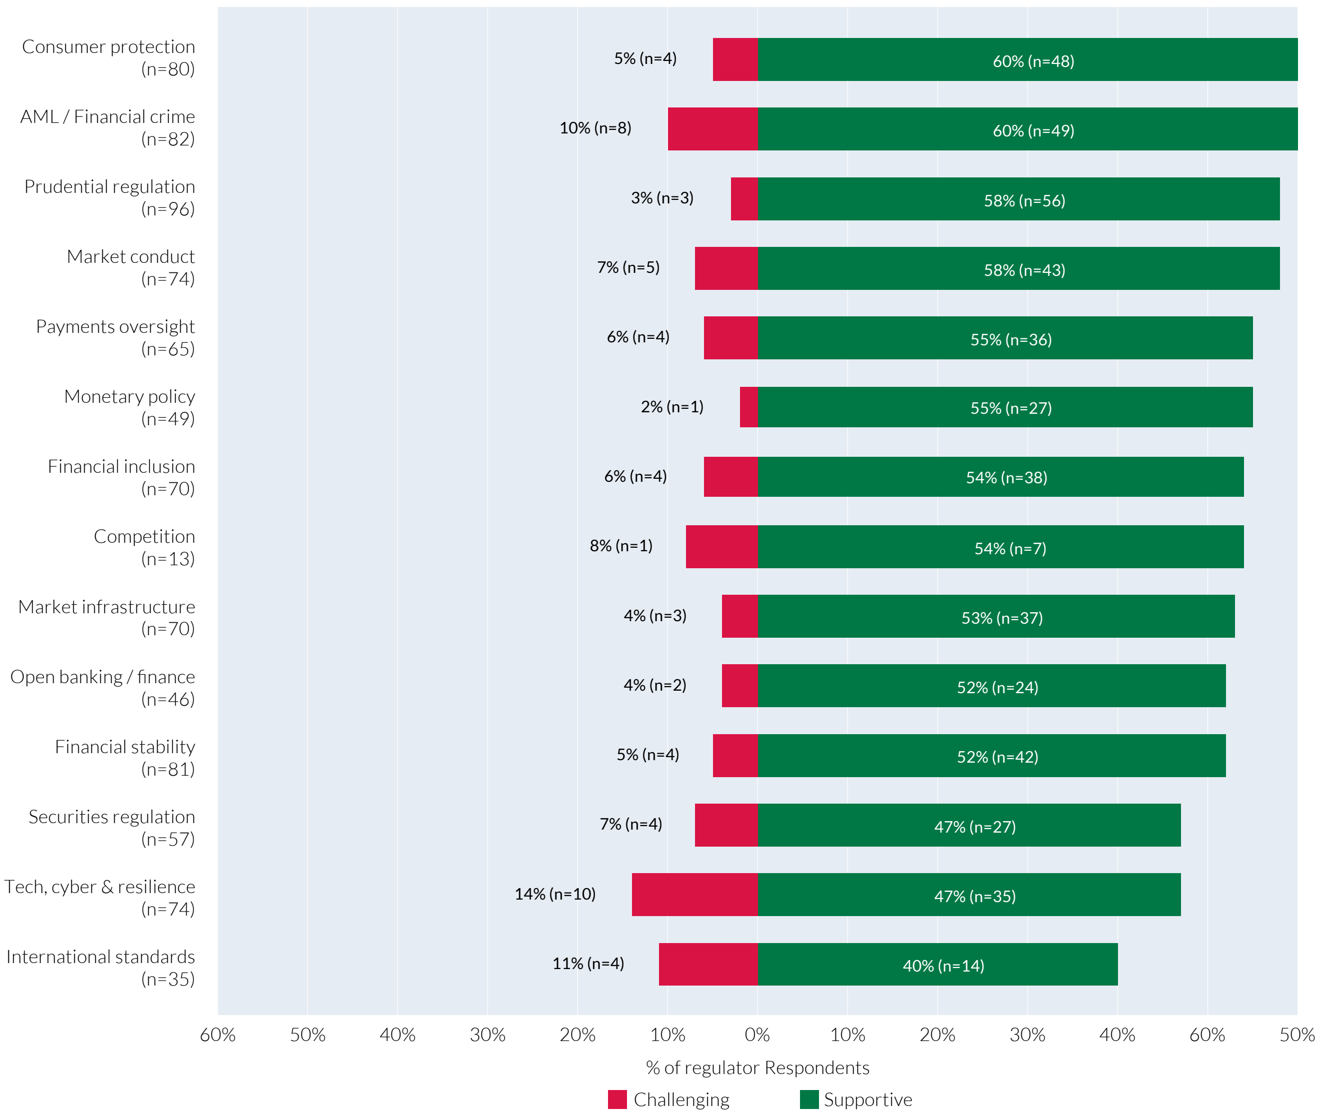

These findings corroborate a recent 2025 IFC report,[34] which suggested that central banks expect significant benefits from AI in performing their various tasks.

Collecting data on AI adoption: While there is a broad sense of optimism about the potential for AI to achieve regulatory objectives, there is a distinct lack of data being collected by regulators on AI adoption to enable informed data-driven decision making and insight.

Limited collection of data: The survey (n=120) reveals that only 24% of authorities collect data on industry AI adoption, and 43% have no plans to start within the next two years. Active monitoring is even rarer for complex risks, with just 5% tracking AI-driven market correlations or systemic bias.

Figure 7.4: Regulator collection of data on AI adoption and related risks – regulators (n=120)

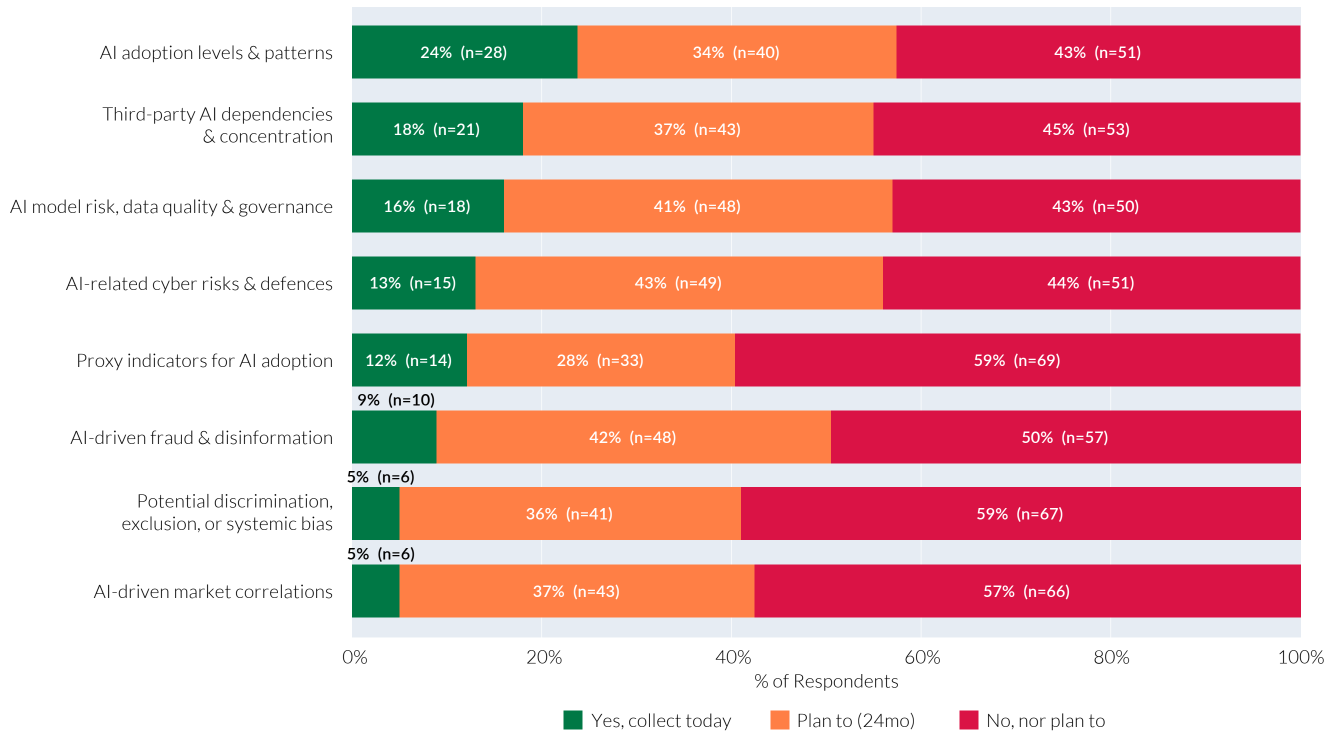

This monitoring gap reflects recent Financial Stability Board (FSB, 2026)[35] findings, which recently noted authorities remain at early stages of monitoring. To safeguard financial stability, the FSB explicitly urges regulators to formalize metrics and address data gaps, particularly for challenging vulnerabilities like market correlations, model risks, and data governance.

Ultimately, this empirical blind spot may undermine the prevailing optimism. Authorities cannot successfully harness or oversee AI if they are navigating its adoption and risks without hard data.

Monitoring for bias and discrimination

One area that highlights the lack of data collection and closely relates to consumer protection is AI-related bias and discrimination. Across the financial ecosystem, algorithmic bias is widely acknowledged as a critical threat. Half of surveyed regulators (50%), alongside 43% of AI vendors and 42% of industry respondents, rank algorithmic bias among the top five risks associated with the progressive deployment of AI in financial services. Despite this high level of shared awareness, there is a stark dissonance when it comes to actual operational practices.

Figure 7.5: Algorithmic bias as a top five AI risk by stakeholder group

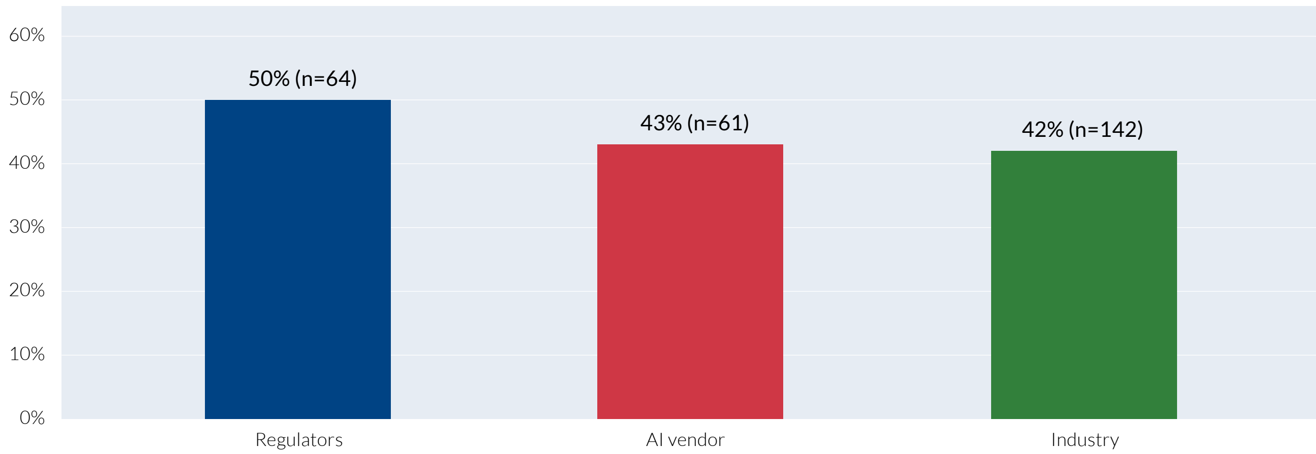

Lack of industry monitoring for bias: A significant majority of industry organisations are failing to address the issue proactively, with 65% of surveyed firms reporting that they do not currently monitor their AI models for bias or discrimination.

Figure 7.6: Industry monitoring for bias or arbitrary discrimination in AI models – industry (n=263)

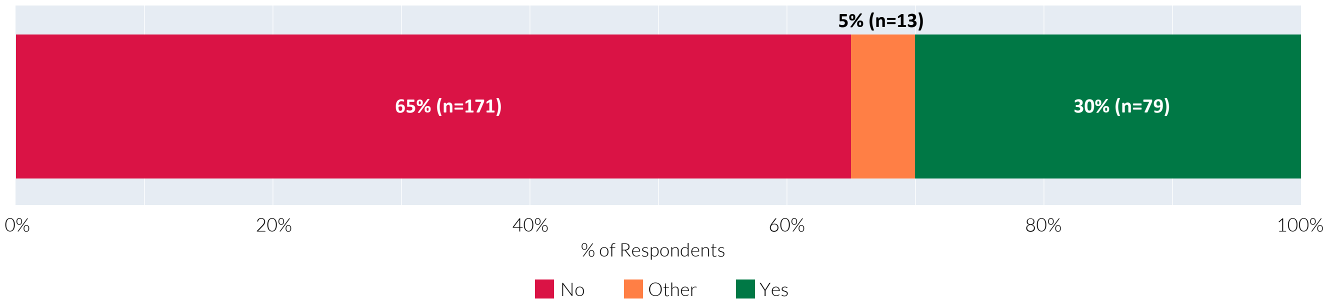

Regulators also lagging in data collection on bias: This operational inaction extends directly into the regulatory sphere as well. Even though regulators frequently flag algorithmic bias as a primary concern, only a marginal 5% currently collect data regarding arbitrary discrimination, exclusion or systemic bias within the use of AI models.

Figure 7.7: Regulator collection of data on bias and discrimination – regulators (n=114)

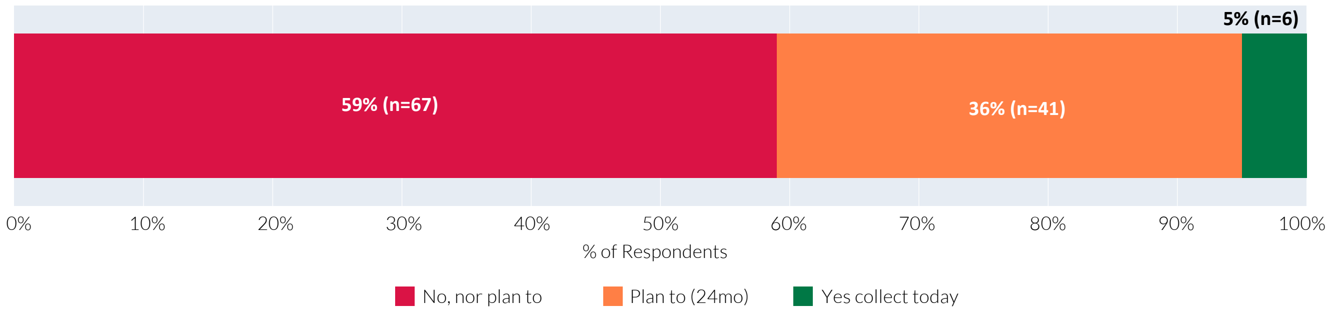

Furthermore, 59% of regulatory authorities report that they have no plans to begin collecting this data within the next two years. The BIS highlighted some of the challenges around collecting data relating to bias and discrimination which is certainly difficult.[36]

While stakeholders widely recognise the severe risks AI systems pose regarding fairness and discrimination, there is a profound lack of concrete, data-driven action being taken to measure, monitor or mitigate these risks.

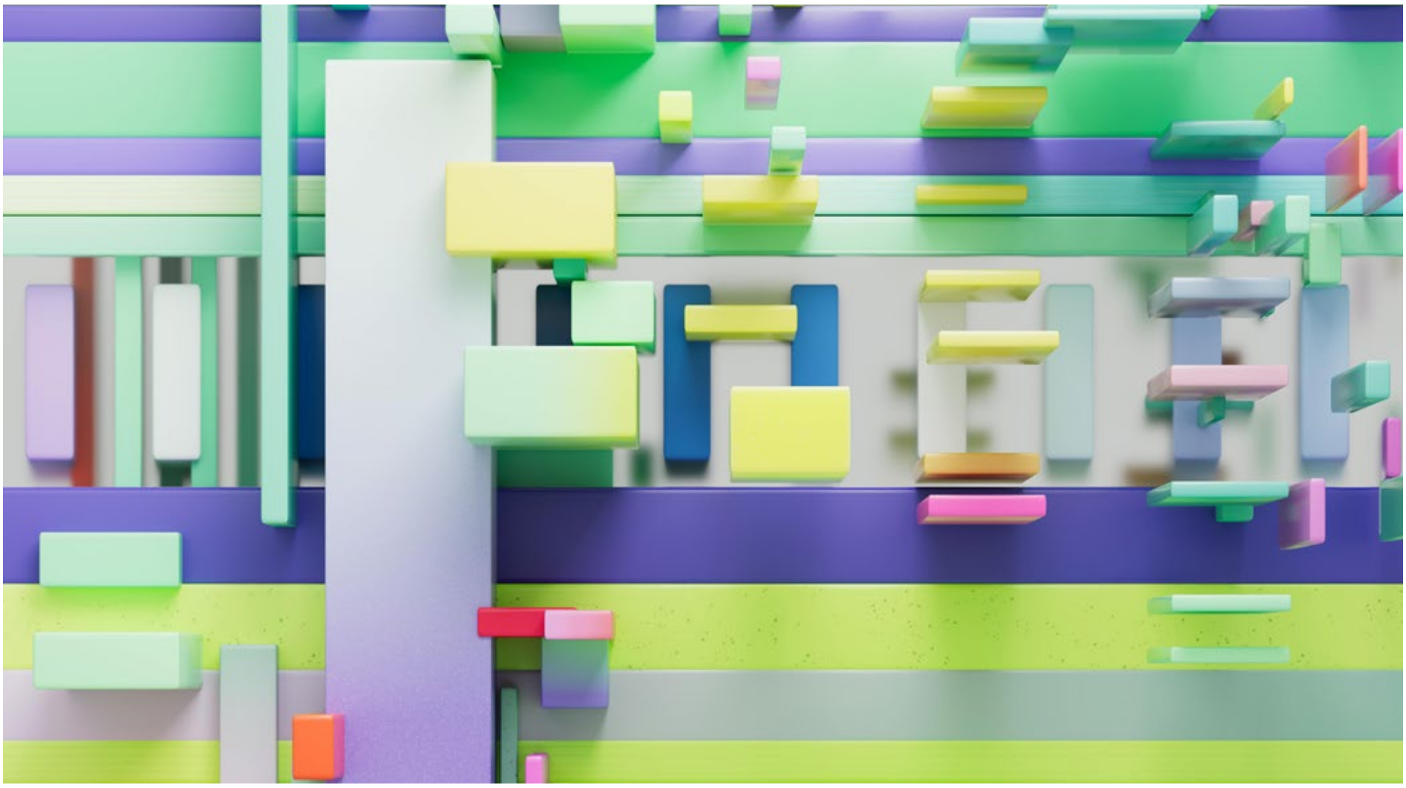

AI regulation and policy developments

AEs versus EMDEs: Currently, 48% of the regulators surveyed already have a national AI strategy and guidelines in place with AEs slightly ahead (55%) compared with EMDEs (44%).

Figure 7.8: Applicable national AI regulatory frameworks for the financial sector – regulators (n=126)

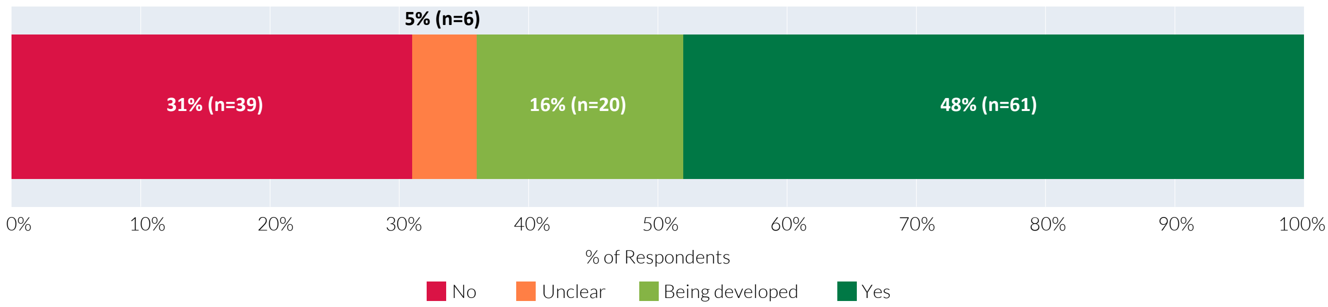

Variation across regions: Most regulators report basing their policy on broad cross-sector guidelines rather than specific frameworks for financial services. However, there remains comprehensive variation across regions with MENA, Europe and SSA reporting the highest proportion of specific policies for AI in the financial sector.

Figure 7.9: National AI policies for the financial sector by region

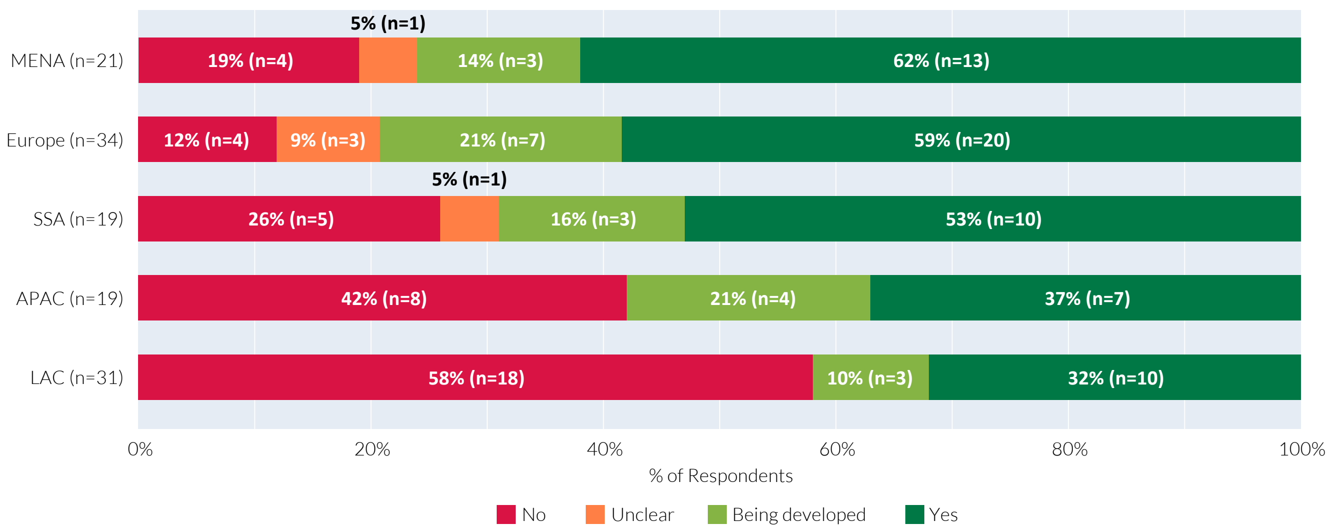

The AI references that are informing regulators

It is of interest to see which regions have regulatory frameworks in place, with MENA and Europe having the highest percentage of regulators who have AI frameworks. Some of these frameworks are used as a reference for other authorities when developing regulatory frameworks. For example, as the table below indicates, the EU AI Act is acting as a reference for other jurisdictions in Europe and central Asia, even though the implementation of the EU AI act is still being addressed.

Figure 7.10: AI regulatory frameworks referenced by regulators – regulators (n=125) (multi-select)

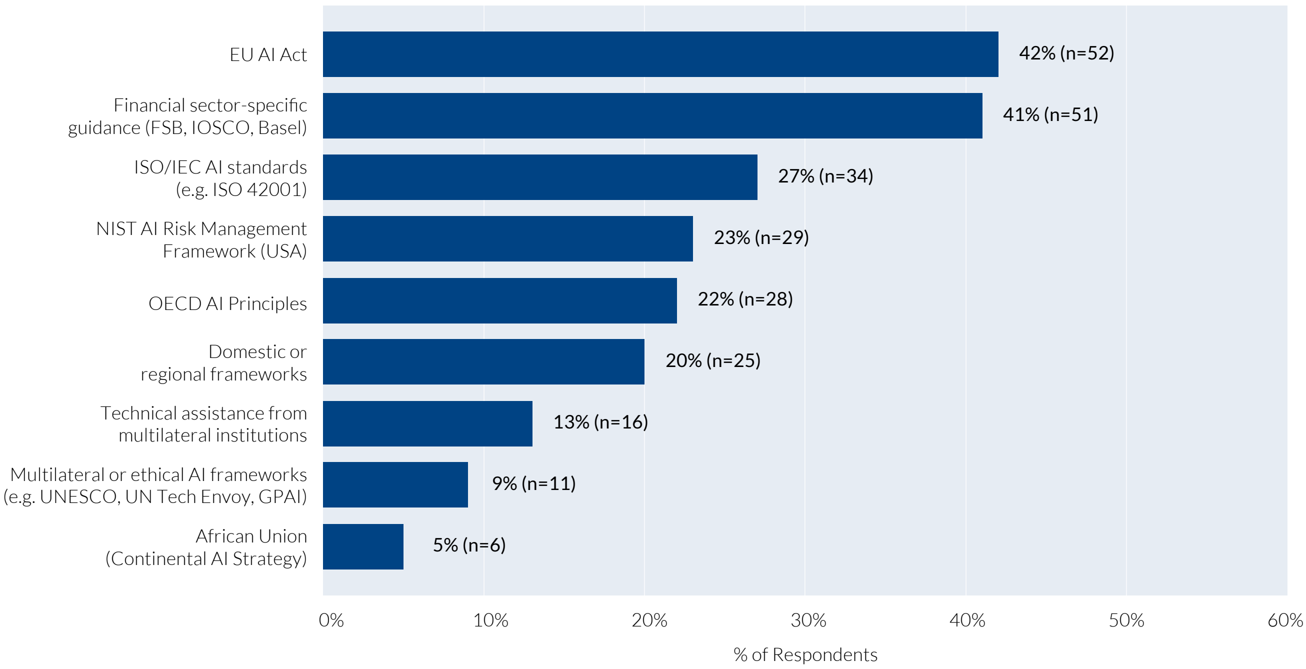

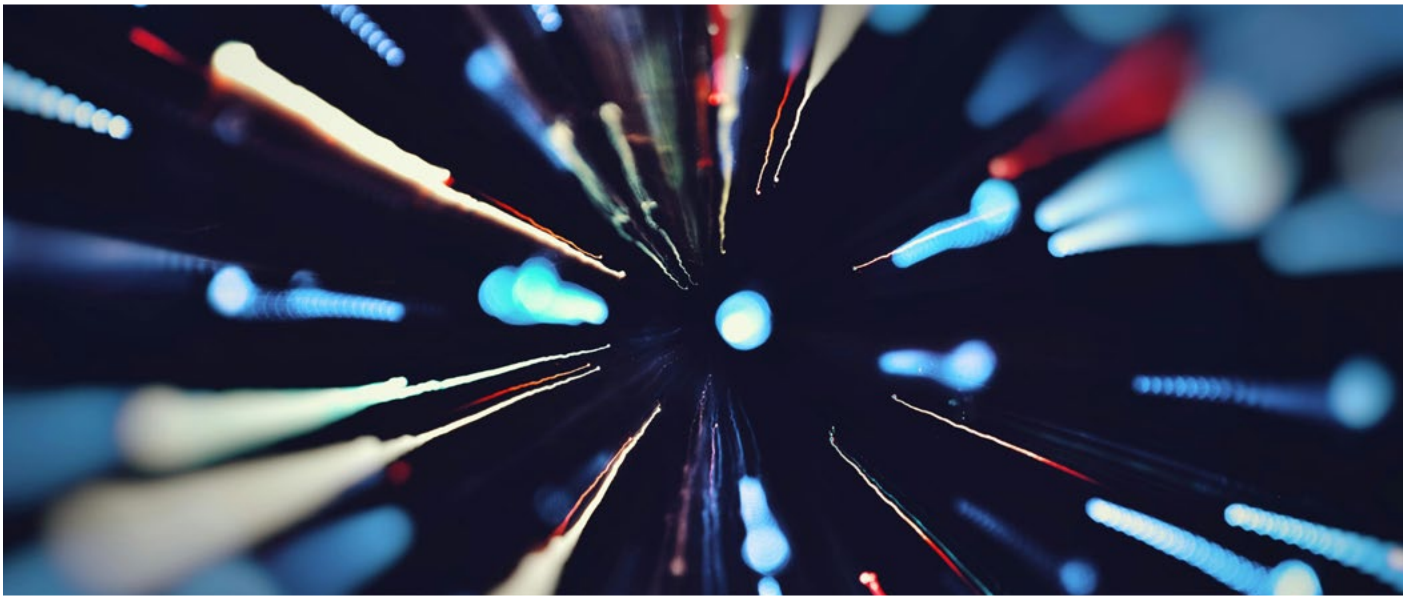

Perception of AI regulatory environment by industry and vendors

Applicable regulatory frameworks for AI, as perceived by industry respondents, appear to be in a period of transition. A sizeable share (65%) of industry respondents report already being subject to both AI-specific and financial sector rules, while 58% say they are not yet subject to AI regulation and 40% expect to be soon regulated.

Figure 7.11: Regulatory frameworks currently applying to industry use of AI – industry (n=333) (multi-select)

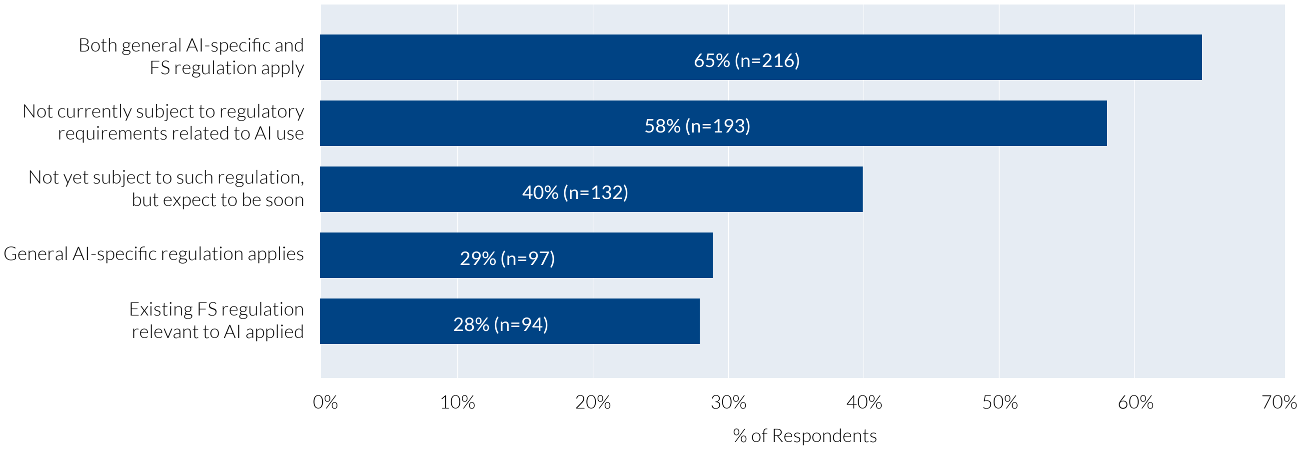

This suggests that, from the industry's perspective, AI oversight is present but uneven and is often being experienced as an extension of existing financial regulation rather than as a wholly separate regime. This perception is consistent with recent OECD and BIS work showing that many authorities are first applying existing financial sector rules and supervisory tools to AI use cases, while AI-specific rules are emerging more gradually.[37, 38]

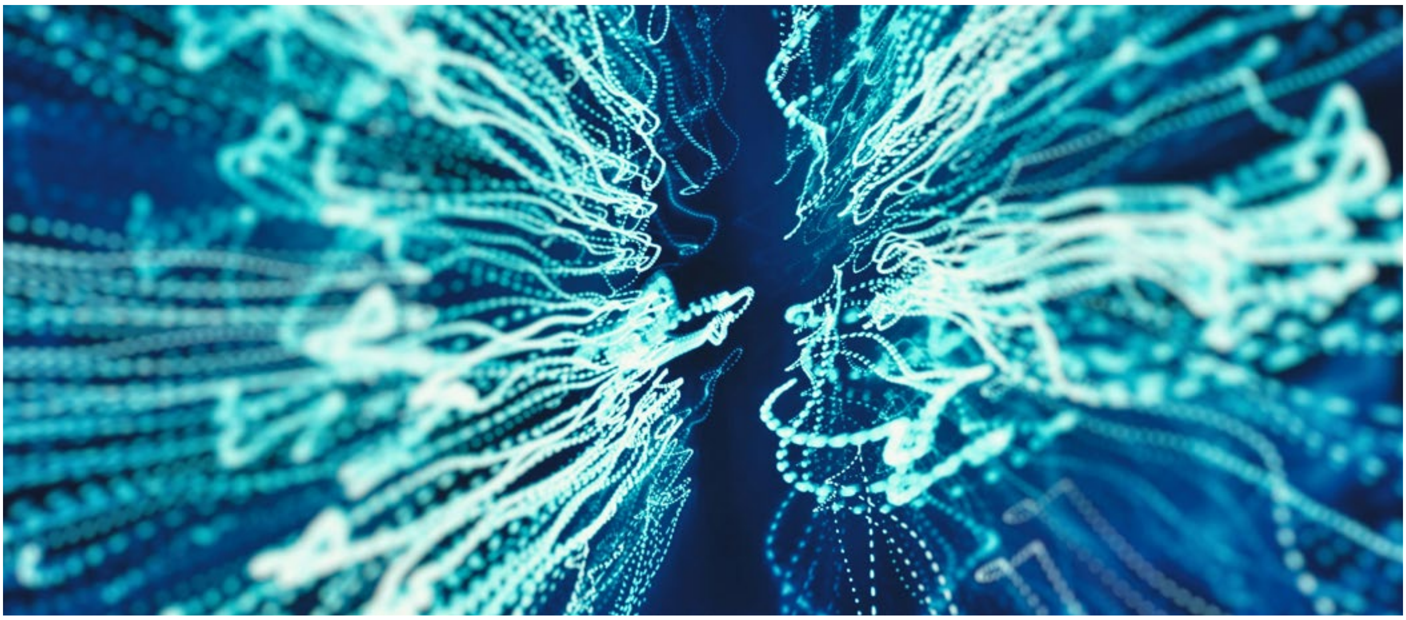

Priorities and internal use of AI by regulators

There is near unanimous consensus from industry, regulators and vendors on the immediate priority for oversight: regulators need to "clarify and update regulatory guidance for AI use". This was a top priority for 69% (n=236) of industry and 67% (n=95) of vendors and 79% (n=98) of regulators.

Figure 7.12: Regulatory priorities for AI oversight by stakeholder group – industry (n=341), AI vendors (n=141), regulators (n=124)

Regulators also highlighted the need to "invest in supervisory tools, skills and test environments" with 79% (n=98) listing this as their joint top priority alongside clarifying and updating AI guidance. However, the same was echoed on a lesser scale by industry respondents (44%, n=151) and AI vendors (47%, n=66).

Use of AI within regulators

Given the rapid developments of AI in technology and industry, regulators themselves have been deploying their own internal uses of AI to execute their functions.

However, 81% of regulators are still Exploring or Piloting or are very limited in their overall use of AI. Hence, the mass implementation of AI use cases for regulation is still at an early stage, despite the promise of increasing the effectiveness of AI use cases for regulators.

Figure 7.13: Regulator AI adoption maturity – regulator (n=130)

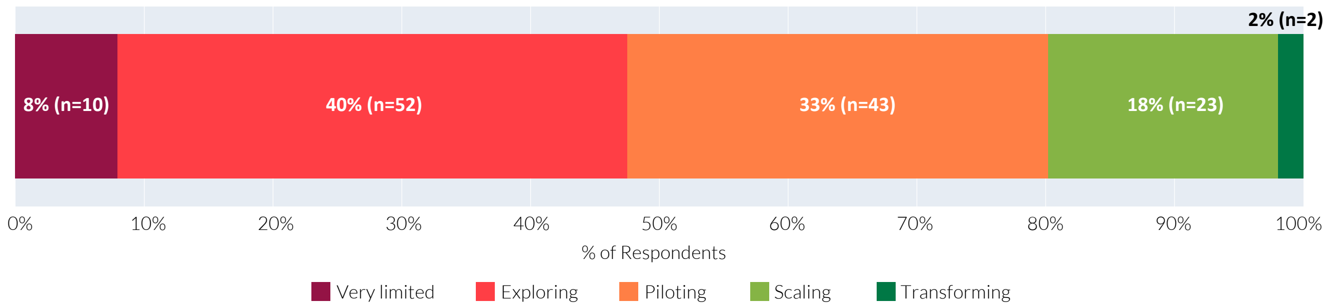

Regulator AI use case maturity

Regulators in AEs are further ahead in terms of scaling or transforming their organisational approach to AI. The gap between the AE jurisdictions and EMDE jurisdictions is stark in terms of internal AI use case adoption. Just 9% of the surveyed EMDE jurisdictions are scaling AI use cases, in contrast to 31% of AE regulators. It may be noted EMDEs' regulators report a greater level of capacity constraints (for example, insufficient technical expertise, limited internal expertise and insufficient funding/ budget) in adoption of AI. Accordingly, the gap in AI adoption between EMDE regulators and in AEs may be reflective of the greater capacity constraints faced by EMDEs in adopting AI use cases.

Figure 7.14: Regulator AI adoption maturity by economic development – advanced economies (n=51) versus EMDEs (n=79)

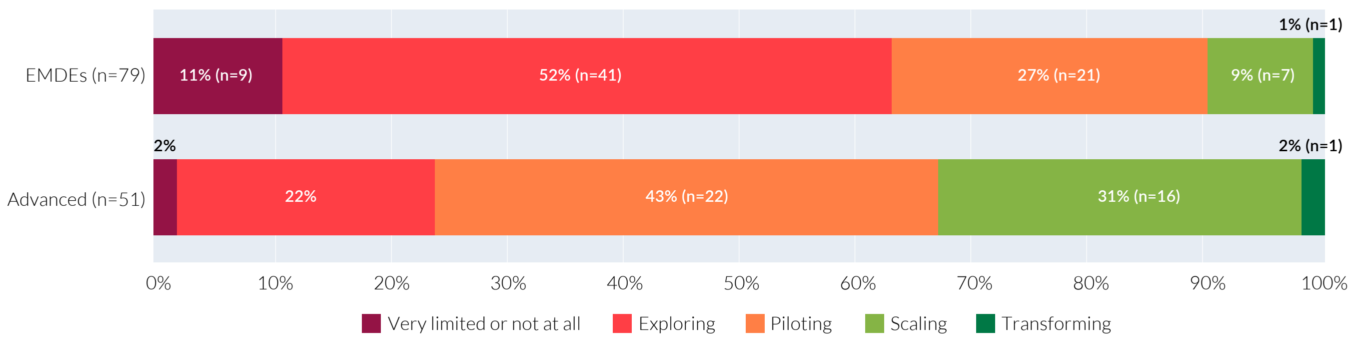

Partner perspectives: IDB – AI for supervision in Latin America

By Diego Herrera, Henrique Chitman and Ana María Zárate Moreno

In LAC, it is increasingly common for financial regulators and supervisors to explore the use of AI to support their oversight mandates. The benefits of employing AI as a SupTech tool are significant.

1. For data collection, AI enables the processing of large volumes of structured and unstructured information, rapidly identifying patterns and anomalies.
2. In the domain of data processing, AI can automate routines for data validation, document classification and both simple and complex data quality checks, thereby allowing supervisory personnel to focus on higher order judgment and policymaking. For instance, in jurisdictions where liquidity risk data informs supervisory actions, AI can support the rapid generation of early warning indicators, facilitating supervisory off-site decisions.
3. Enhanced data processing capabilities also improve the detection of misconduct, including schemes that are difficult to identify through manual review.
4. Moreover, AI is well suited to detecting patterns and outliers relevant to AML/CFT oversight, making enforcement a matter of judgement and action.

Overall, the benefits are substantial: AI can help financial authorities allocate scarce supervisory resources more efficiently and concentrate efforts on analytical review, on and offsite examinations and enforcement activities.

Data constitutes the foundational element for the application of AI in supervisory processes. In this sense, data is the indispensable precondition, from which any effective AI-enabled supervisory framework emerges. This touches institutional capacity in at least four aspects.

1. First, AI integration requires the pre-existence of a proper supervisory model, preferably risk-based, to arrive at the on and off-site examinations, decisions and enforcement when necessary.
2. Second, this condition is based on a regulatory mandate to request the data beyond the legacy ones (balance sheet, income statement and the notes to financial statements): credit, market, liquidity and operational risks are necessary.
3. Third, it requires governance: who, where, how the data is received, processed and informed (including privacy and security mandates), also data modelling rules, cybersecurity mandates and vendor-related conditions when there are no in-house developments.
4. Fourth, AI integration also requires technology to integrate AI into supervisory processes and systems, model training and verification.

Incorporating AI into financial supervision demands sustained institutional effort, and ultimately, the effectiveness of AI tools depends on the presence of a well-defined supervisory model and informed human judgment.

LAC: AI adoption in financial regulators

While there is a growing interest in using AI for regulatory and supervisory functions in LAC, financial authorities are moving slowly along the maturity curve due to foundational gaps in data quality, infrastructure and talent.

LAC regulators lag behind global peers in overall AI adoption maturity, and AI activity is heavily concentrated in early-stage experimentation. The share of regulators with very limited or no AI adoption is twice as high in LAC (12%) compared with the global average (6%). Also, 85% of LAC regulators remain in exploration or pilot phases, compared with 70% globally. Only 3% of regional authorities have reached the Scaling stage, far below from 22% of global regulators. Notably, none of the LAC regulators have yet achieved the most advanced "transforming" level of maturity, where AI becomes a key building block in supervisory and regulatory functions. These trends show that the region is progressing through the maturity curve at a slower pace.

Therefore, there is a significant opportunity to leverage AI in strengthening oversight, detecting and assessing risks, and improving regulatory processes in LAC. Overall, these trends point to institutional capacity gaps, including quality data and structural constraints that hinder digital transformation in supervisory functions. These constraints include limited internal resources and expertise, including specific AI expertise, insufficient tools/infrastructure and insufficient standardised metrics. In terms of infrastructure, LAC adoption of cloud and external foundation models is lower than global averages, reinforcing slower AI maturity relative to other regions. This mirrors global findings that regulators (LAC included) remain more reliant on on-premises systems.

LAC's regulatory and supervisory use of AI mirrors global trends but lags in operational readiness, largely because the foundational data, governance and infrastructure needed for AI, remains underdeveloped. Addressing these challenges will require a sustained focus on capacity building, as LAC regulators acknowledge AI's central role in meeting their mandates by 2030.

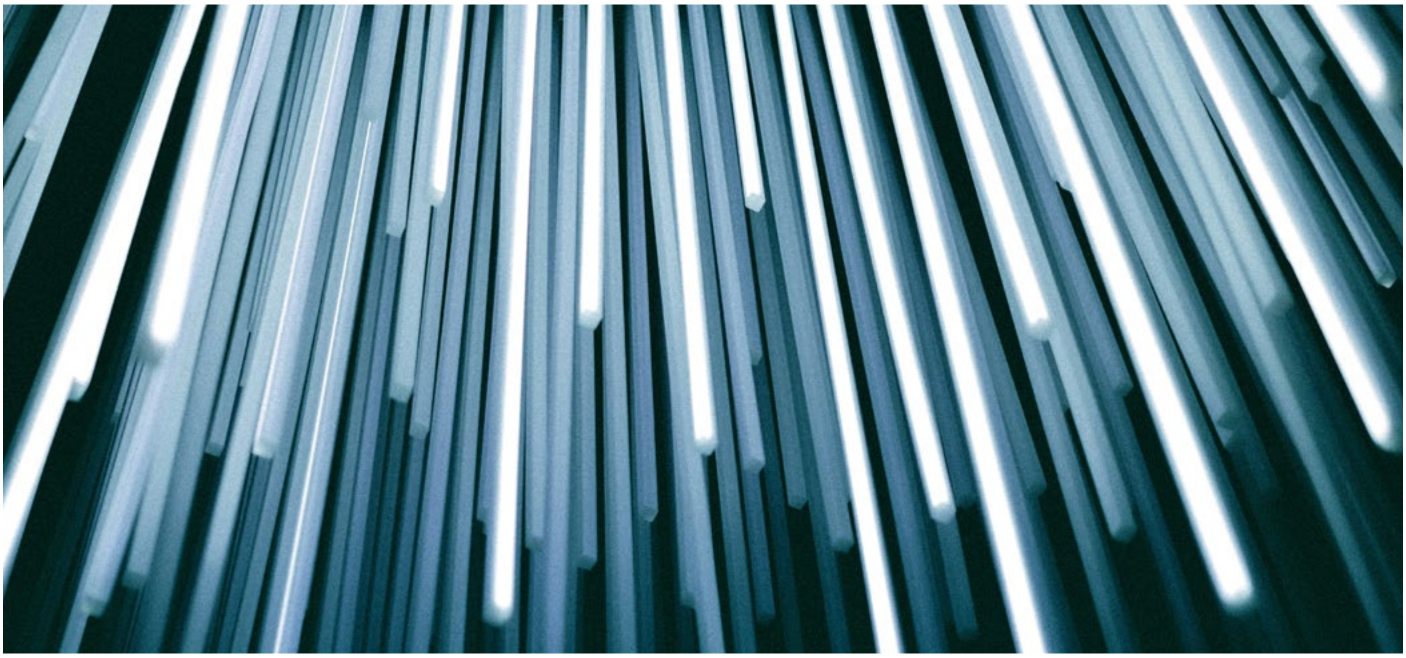

Regulator AI use-case maturity

When considering specific use cases, the supervision and enforcement function has seen the highest adoption of AI by the surveyed regulators.

However, the proportion of regulators using AI for achieving their regulatory objectives, including for supervision and enforcement, remains low. This is despite the high optimism displayed by regulators as described in the previous chapter.

Figure 7.15: Regulator AI use case maturity by regulatory domain

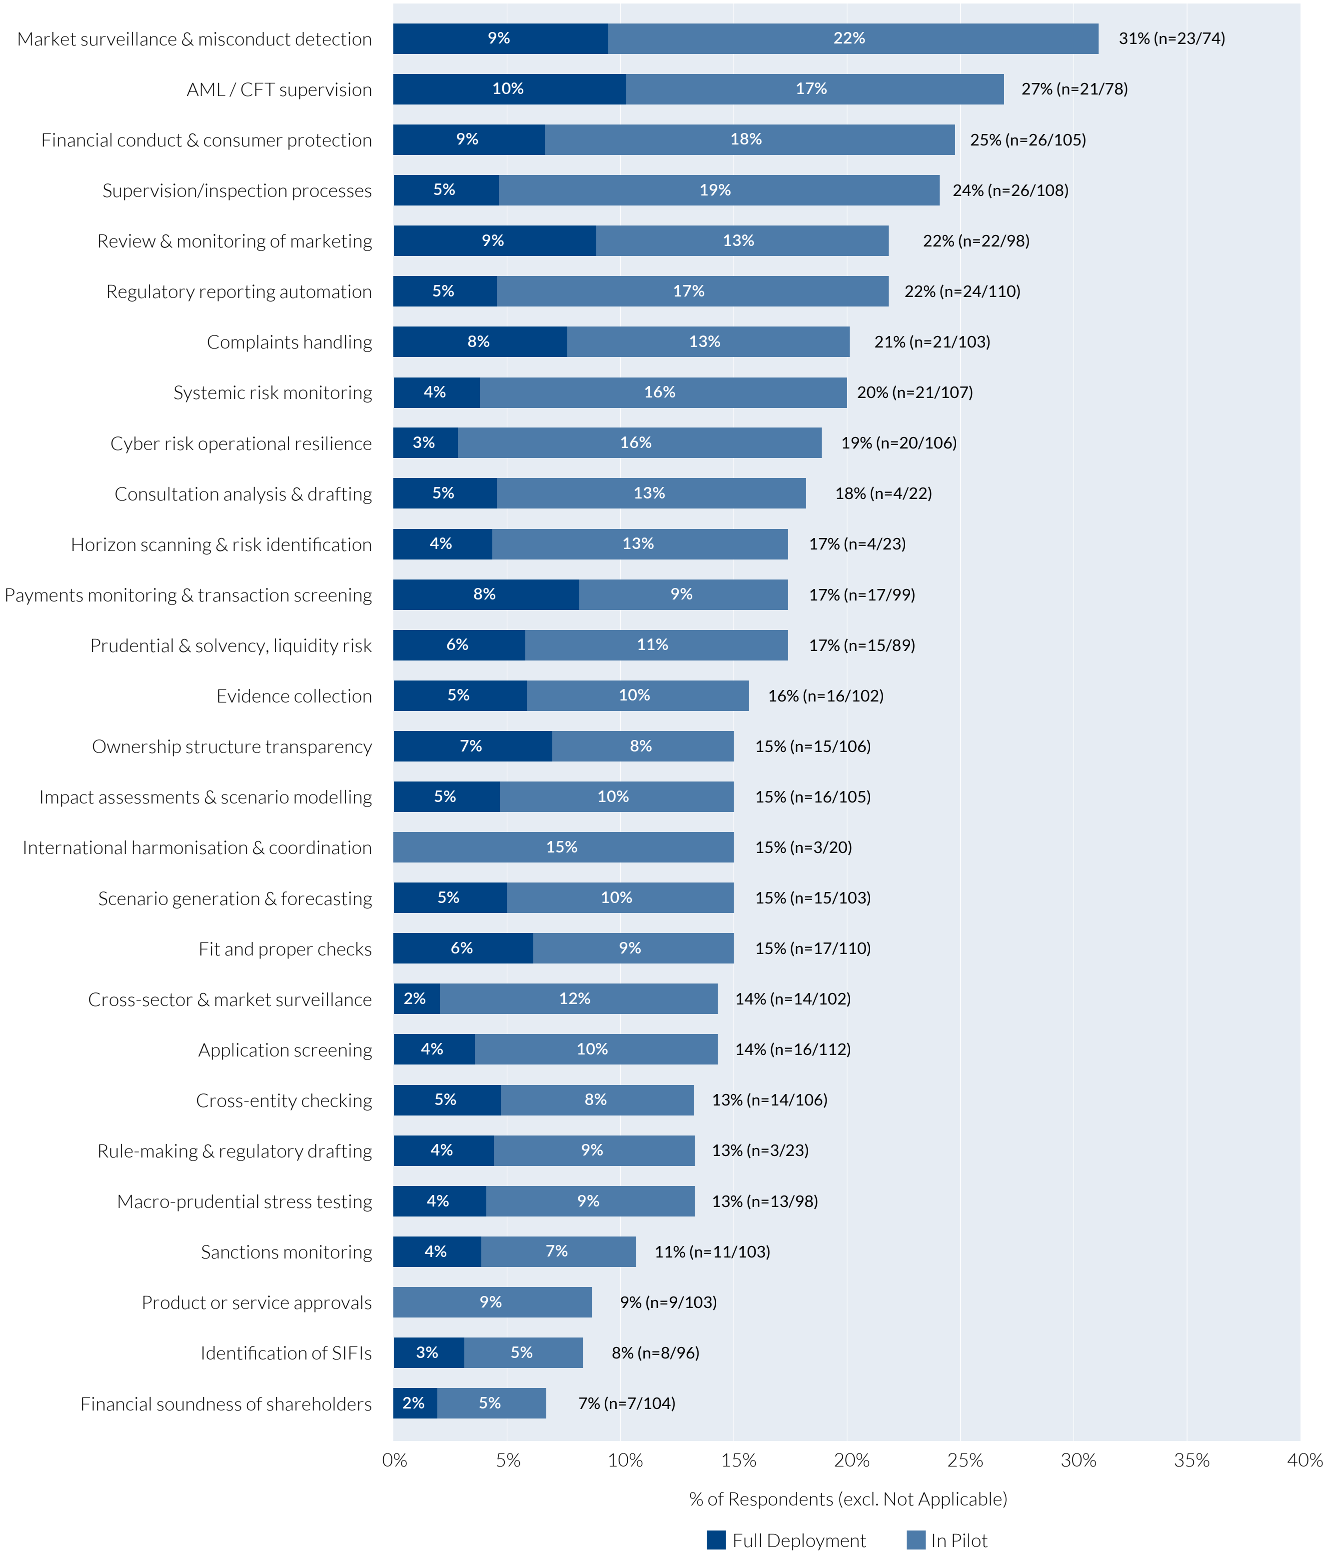

The low usage is also consistent with a recent BIS survey (2025),[39] which suggested that despite high expectations, current AI-based applications across regulators remain limited.

Licensing and authorisation

The usage of AI within licensing and authorisation remains among the lowest across regulatory application areas. Fit and proper checks (15%), application screening (15%) and ownership structure transparency (15%) remain the top three use cases within this function. Furthermore, the proportion of regulators who have a visibility of adopting AI within each specific use case remains quite low, in single digits.

There might be several reasons explaining this lower use including Principle 5 of the Basel Committee on Banking Supervision (BCBS) Core Principles for Effective Banking Supervision standards, which sets minimum criteria for licensing banks. At a minimum, the licensing process should consist of an assessment of the ownership structure and governance of the bank and its wider group, its strategic and operating plan, internal controls, risk management and projected financial conditions. Thus, the licensing criteria across regulators would usually be qualitative once the basic minimum quantitative requirements (such as minimum capital) have been met. Accordingly, it may be difficult to establish objective measures to train AI as to in what instances a license may be considered favourably or unfavourably.

Figure 7.16: Regulator AI use case adoption in licensing and authorisation

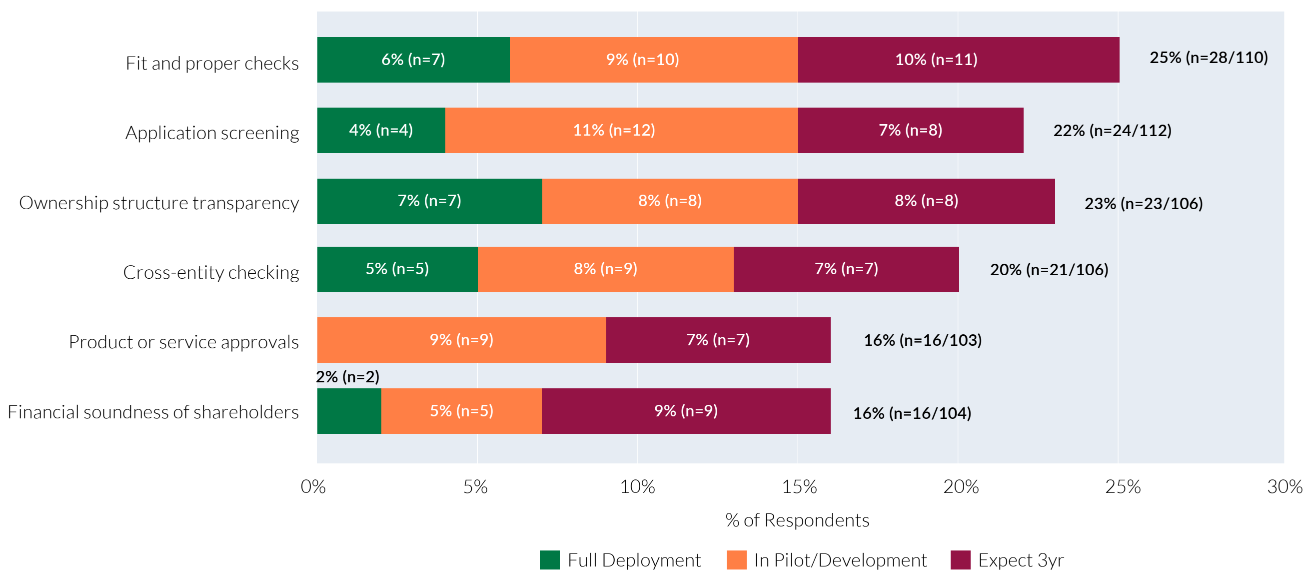

Another general challenge may be a tension between predictive accuracy and 'interpretability/ explainability'. While sophisticated ML models may deliver accurate predictions, the complex, non-linear interactions of many variables make it hard to assess which inputs are key to driving the output. Reliable forecasts therefore come at the price of accepting that the underlying model is a black box (BIS, 2025).[40] This may specifically be an issue within the licensing and authorisation function of regulators, where any rejection of a license may be subject to intense scrutiny and require clear explanation around reasoning for acceptance or rejection.

AI for policy and rulemaking

'Policy and rulemaking' is an area where the surveyed regulators have the lowest proportion of use cases but have the highest expected future AI deployment.

The proportion is particularly high in the 'horizon scanning and risk identification' use case where nearly 30% of regulators expect to deploy AI within the next three years. This may be reflective of the capabilities of AI/ML models to work with unstructured data and information, which regulators grapple with while framing policies.

Figure 7.17: Regulator AI use case adoption in policy and rulemaking

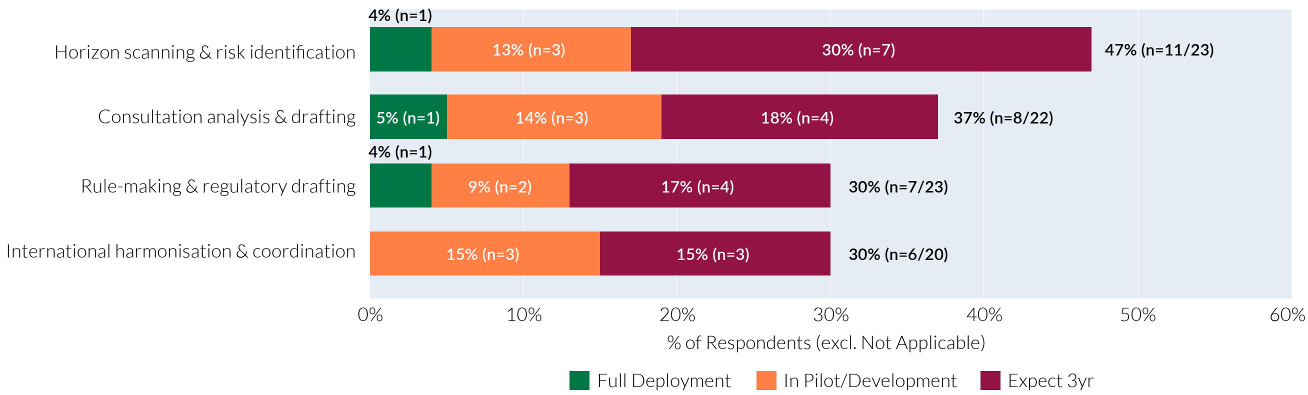

While AI should not be used to replace human judgement of regulators, AI may serve as an additional input tool in identifying the discrepancies and inconsistencies during regulatory drafting, especially across different regulations and jurisdictions.

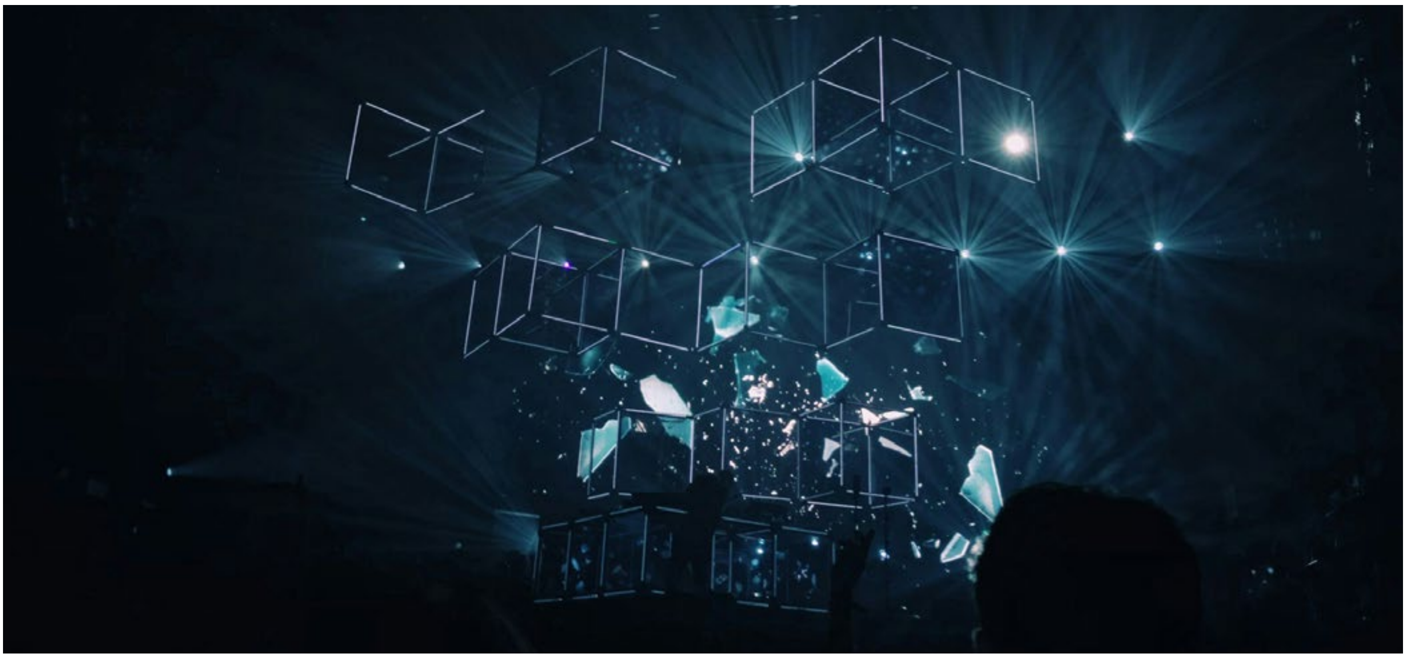

AI for supervision and enforcement

As stated above, AI for supervision and enforcement is the most advanced area for regulators. The most common use cases within supervision and enforcement with full deployment or in pilot are market surveillance and misconduct detection (31%), AML/CFT supervision (27%) and financial conduct and consumer protection (25%).

Supervisory inputs may span from data collected through returns from regulated institutions, news stories, internal bank files and previous supervisory reviews. Manually mining this trove of data for actionable insights is labour-intensive and as volumes grow it verges on impossible. AI can help in tasks such as document processing, knowledge management and document review (BIS, 2025).[41]

Figure 7.18: Regulator AI use case adoption in supervision and enforcement

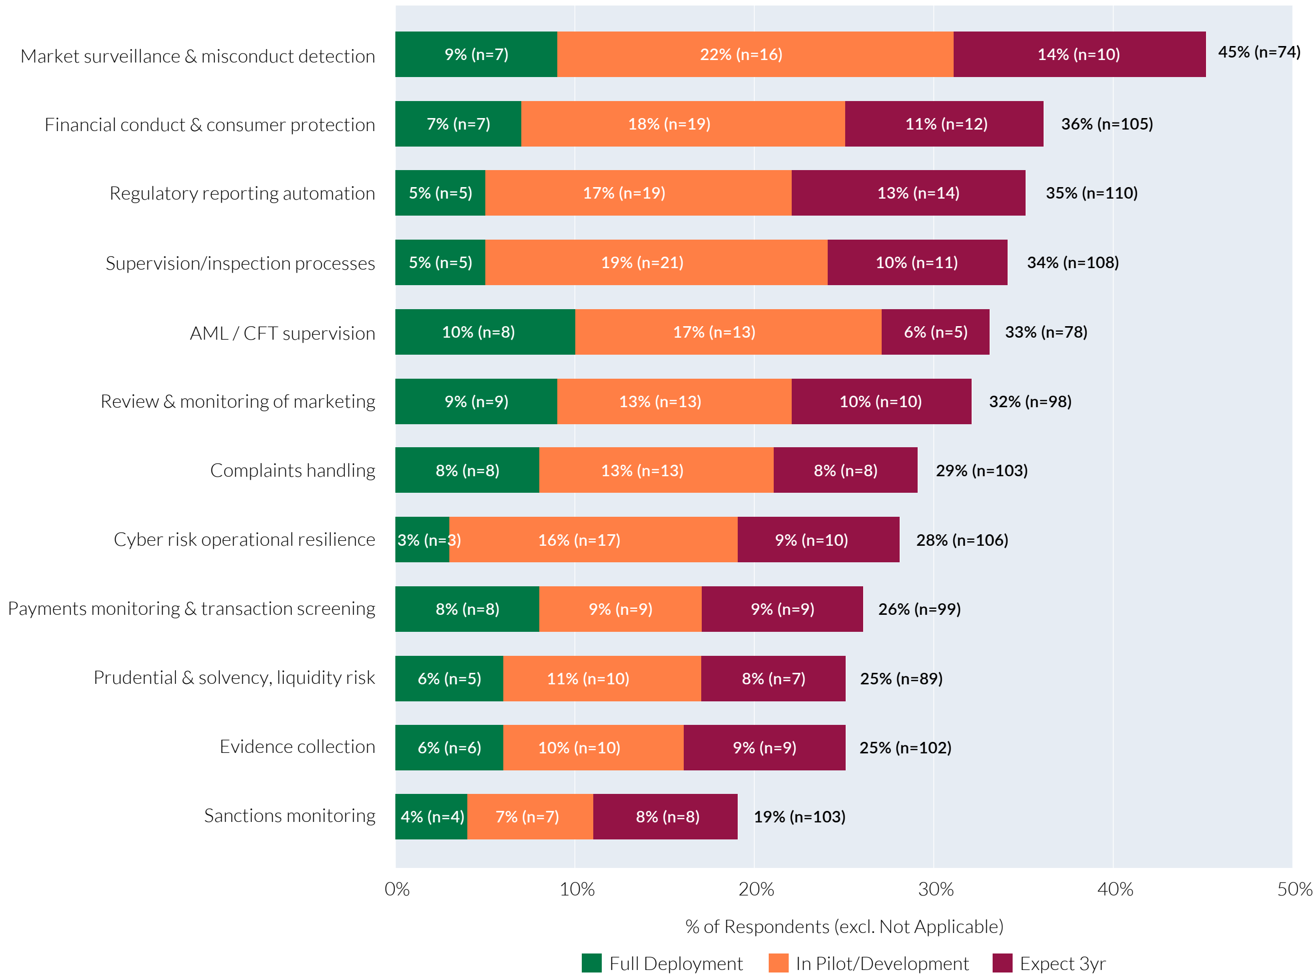

AI for financial stability and oversight

The usage of AI by regulators for financial stability and oversight remains one of the least developed, with most of the use cases remaining under Pilot/Development stage. The top three use cases are systemic risk monitoring (20%), impact assessments and scenario modelling (15%) and scenario generation and forecasting (15%).

However, use of AI for financial stability and oversight may hold potential. This has been echoed by the BIS in their report on the use of AI for policy purposes (BIS, 2025).[42] The report states that financial stability risks usually stem from a slow build-up of vulnerabilities. Fed with rich data sets that offer ample scope for pattern recognition, AI can help craft early warning indicators that alert supervisors to brewing pressures linked to system-wide risk. AI can also help design more realistic scenarios by learning from historical crises or even generating plausible crisis narratives and can also accelerate the processing of results.

Figure 7.19: Regulator AI use case adoption in financial stability and oversight

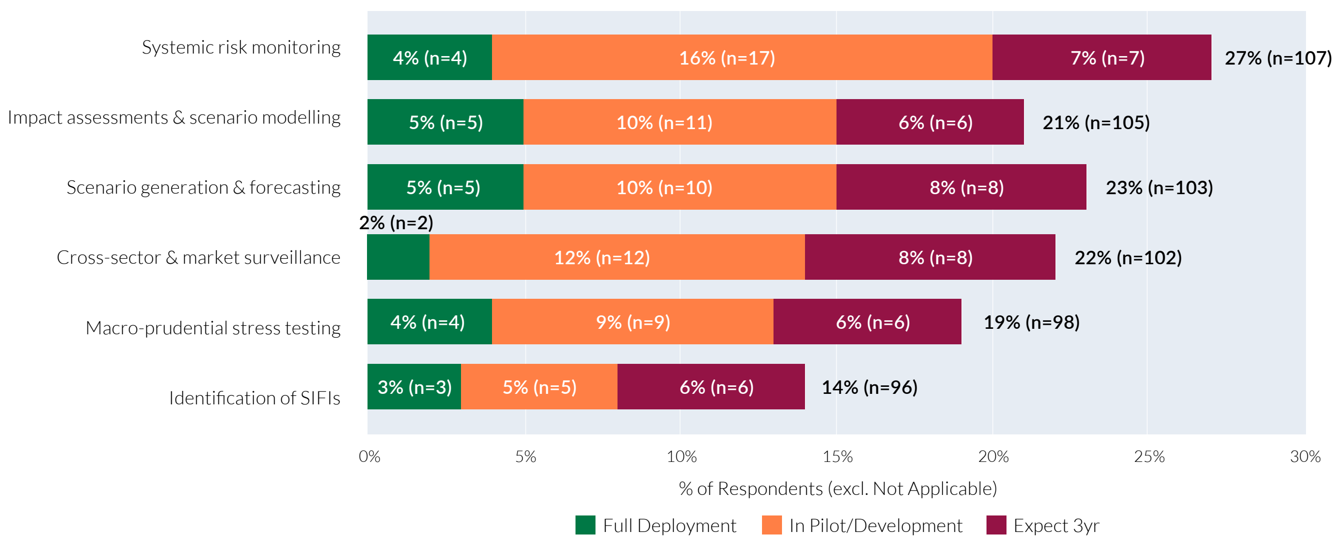

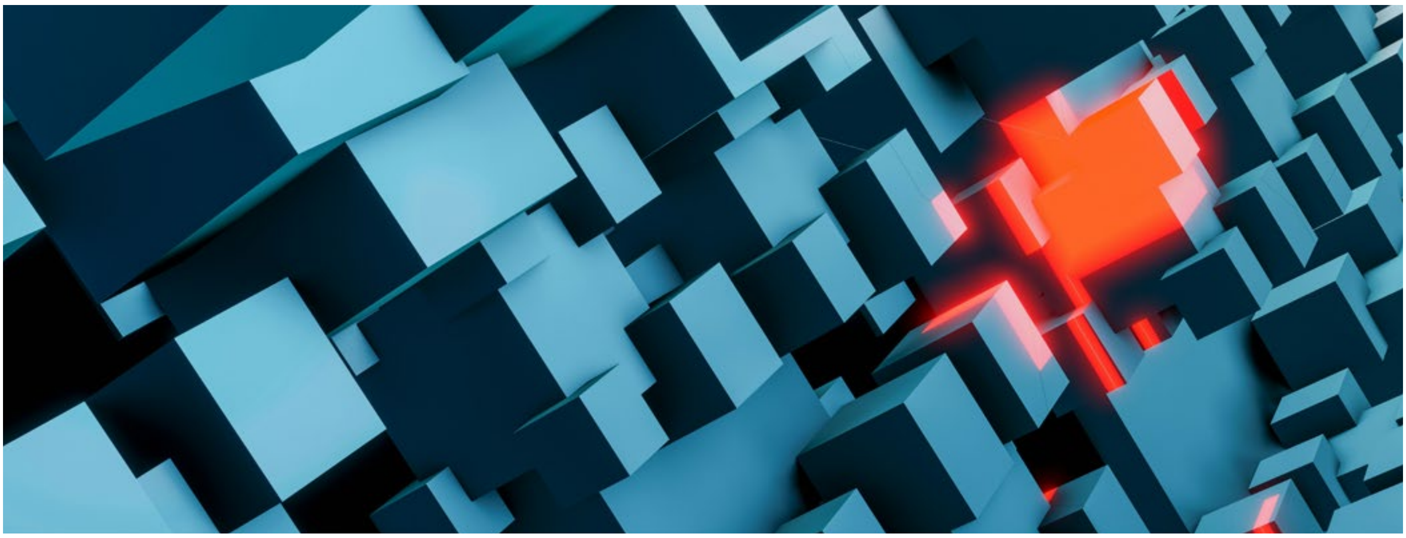

Summary of report to G20 South Africa by World Bank

By Gian Boeddu , Ezio Caruso, Erik Feyen, Serafin Martinez Jaramillo, Karol Karpinski, Sergio Mesquita, Yasemin Palta, Srishti Sinha and Alexandra Gutiérrez Traverso

The content in this sub-section is drawn from a recent World Bank Group report to the G20 South Africa Presidency.[43]

A central pillar of the World Bank Group's engagement with financial sector authorities in Emerging Market and Developing Economies (EMDEs) is supporting them in strengthening supervisory frameworks and processes, and in implementing the technology that facilitates such supervision. This includes the implementation of technology-enabled supervision tools (SupTech), as well as supporting supervisors in monitoring financial service providers' adoption of new technologies.

Artificial intelligence (AI) is increasingly central to this work, offering significant potential to make supervision more efficient, data-driven, and innovative. At the same time, the adoption of AI in supervisory contexts introduces its own set of risks and challenges that authorities must carefully navigate.

The benefit vs risk dimensions of AI in SupTech are an important consideration: on the one hand, AI can enable supervisors to process larger volumes of data, identify risks more quickly, and allocate resources more effectively; on the other, it raises concerns related to model governance, data quality and governance, institutional capacity, cybersecurity, technology providers concentration/lock-in, and the risk of over-reliance on automated outputs. Understanding both dimensions is essential for authorities seeking to harness AI responsibly and effectively.

The 2025 G20 Presidency commissioned the World Bank to prepare a report on AI in financial sector supervision in EMDEs, particularly in Africa, with these focuses in mind. The observations below reflect findings in that report.

AI adoption by EMDE financial sector authorities

As of mid-2025 only a limited number of EMDE authorities appear to be using AI for core supervisory tasks such as data collection, on- and off-site supervision, and anomaly/fraud detection: all of these are either at an early stage of adoption or conducting tests and pilots of AI solutions. However, EMDE authorities are optimistic about the potential benefits of AI and are increasing their invested resources and focus. Some authorities have been mapping supervisory processes to identify areas where AI could add the most value. Some are more proactive, encouraging departments to experiment broadly, while others are more cautious, limiting AI experimentation to certain types of projects or supervisory business lines.

Basic GenAI tools have seen widespread uptake by staff within EMDE authorities for general purposes such as drafting and summarization. But beyond this, some authorities have been cautiously working to deploy AI agents, chatbots, and other GenAI-based tools for more sophisticated tasks such as internal knowledge management, complaints analysis, and risk and compliance assessments of supervisory documents. It is important to note that these more advanced use cases are generally found in authorities that have been experimenting with a broad range of AI solutions for six to eight years, well before GenAI became mainstream. This accumulated experience has given those authorities a stronger foundation from which to advance in the current technological environment, underscoring the importance of building institutional AI capacity progressively and over time.

Collecting and managing large amounts of often sensitive data remains a significant challenge for EMDE authorities seeking to integrate AI into their supervisory processes. Data is often fragmented or not available in formats that are readily usable, and authorities must also navigate data privacy, security, and localization requirements. Cloud adoption – another enabler for AI – varies widely across jurisdictions, with vendor dependency and data sovereignty emerging as common concerns. In parallel, many authorities are attempting to integrate AI within existing, and often outdated, IT infrastructure. As a result, several have prioritized investments in strengthening foundational IT and data systems as a prerequisite for more advanced AI-enabled supervision. Skill gaps and difficulties attracting and retaining talent – cited as a particularly acute challenge by many authorities in Africa – are also a fundamental constraint, and a number of authorities are investing in their workforce's readiness, with some taking a more strategic approach to embedding both domain knowledge and technological expertise within their supervisory teams.

Adoption of AI within authorities also needs to be subject to appropriate governance arrangements. Effective governance goes beyond issuing internal written guidance or policies, although these are important, and requires implementation of clear decision-making structures, accountability mechanisms, and regular evaluation of AI tools and their impact on supervisory activities and the broader organization. EMDE authorities were increasingly adopting formal policies governing their internal use of AI, aiming for a sound risk mitigation strategy, although still allowing sufficient room for innovation. Most authorities that did not already have a formal strategy or policy expected to establish one by mid-2026.

Supervising AI-related risks in the financial sector

Many EMDE authorities currently have limited capacity to monitor AI developments in the financial sector and assess their implications for the different supervisory domains. Several risks are prominent, including potential financial stability implications from the rapid, widespread adoption of a small number of AI models and related offerings. The reliance on a small number of models and global vendors for critical IT and AI infrastructure can amplify herding behavior, concentration and third-party risks. Cybersecurity threats and data breaches are also a primary concern, prompting authorities to focus on financial institutions' safeguarding of systems and implementation of strong governance frameworks.

As financial institutions increasingly adopt AI, authorities must also sharpen their focus on potential consumer and investor protection implications. Some risks are familiar, even if they are manifesting through new technologies. For example, the use of AI in areas like credit decisions and fraud detection can introduce or amplify biases, leading to discriminatory or inequitable outcomes for consumers. Other areas may pose even more significant supervisory challenges. Limited transparency and explainability in AI models can not only make it more difficult for consumers to understand decisions that affect them, but also for authorities to conduct effective oversight. While many EMDE authorities recognize such risks, their capacity to address them is often still developing, reflecting the early stages of AI adoption in their markets. In this context, SupTech solutions leveraging AI may help support supervisory efforts, but supervisory approaches, methodologies, and skills must also continue to evolve.

Looking ahead

Two important observations emerged from interactions with a wide range of EMDE authorities. First, there is a consensus that AI should augment, not replace, independent supervisory judgement and discretion. Automated tools do not by themselves ensure the effectiveness of supervision. Overreliance on AI may severely impact supervisory effectiveness and lead to material reputational damage. Supervisors must retain final authority over AI-assisted decisions and be able to explain their rationale. Second, it is considered essential that supervisors ensure financial institutions understand their own AI applications, their inherent risks to external and internal stakeholders, and can be held appropriately accountable for decisions and customer interactions made using AI.

Effective oversight of AI in financial supervision requires strong collaboration and coordination—across agencies and across borders. Working in isolation can lead to blind spots, inconsistent rules, and duplicated work. By collaborating with each other and with the private sector, authorities can share expertise, align supervisory expectations, and respond better to both domestic and cross-border risks. Engaging with industry is especially important to understand how AI is being used, monitor emerging risks, and support responsible innovation by financial institutions, including Fintechs. In addition, many supervisors oversee financial institutions that operate across multiple jurisdictions, making bilateral cooperation—such as through supervisory colleges—and harmonized cross-border approaches essential for effective AI oversight.

The WBG's report to the G20 concludes with five key recommendations for financial authorities:

- Authorities should establish a board-level governance framework (and corresponding internal processes) to ensure that AI adoption aligns with organizational objectives and maintains public trust.
- Authorities should prioritize the development of integrated internal IT and data infrastructures to support effective AI adoption, paying close attention to the challenges as well as opportunities associated with cloud integration.
- Authorities will need to focus more intensively on developing systematic approaches to attract, retain, and nurture the necessary technical skills and expertise, and to integrate both domain knowledge and new digital skills within supervisory teams.
- Authorities should invest in more systematic monitoring of AI developments in the context of their supervisory methodologies, and in conducting risk assessments, including by addressing data gaps, to strengthen their understanding of the associated opportunities and risks.
- Finally, collaboration at both the domestic and international levels is essential to avoid regulatory fragmentation and arbitrage, mitigate the build-up of new risks, and ensure effective oversight as AI technologies and their applications evolve.

In summary, the survey data point to regulatory ambition and infrastructure still requiring better alignment. The case for stronger AI governance finds broad support across all three stakeholder groups. The investment required to build supervisory capacity in tools, training, and data collection has not yet been made at the scale the ambition requires. Regulators are acutely aware of those deficits, with the data showing regulators being consistently more focused than industry or vendors on the urgency of supervisory investment.

The preceding chapters document a sector broadly committed to AI but divided in its capacity to deploy it well. Real value is being generated in productivity, in operational efficiency and in risk management, but it accrues unevenly. It is strongest among higher-maturity, higher-investment institutions and among fintechs compared to traditional FIs.

The risk and regulatory picture show a parallel pattern. Awareness of AI's risk exposures and the governance needed to address them is broadly shared across all three stakeholder groups, but the frameworks and supervisory capacity to act on that awareness remain underdeveloped relative to the scale of deployment. The survey's forward-looking data, examined in the following chapter, document what the sector expects from this trajectory by 2030 and the distance between those expectations and the governance structures currently in place.

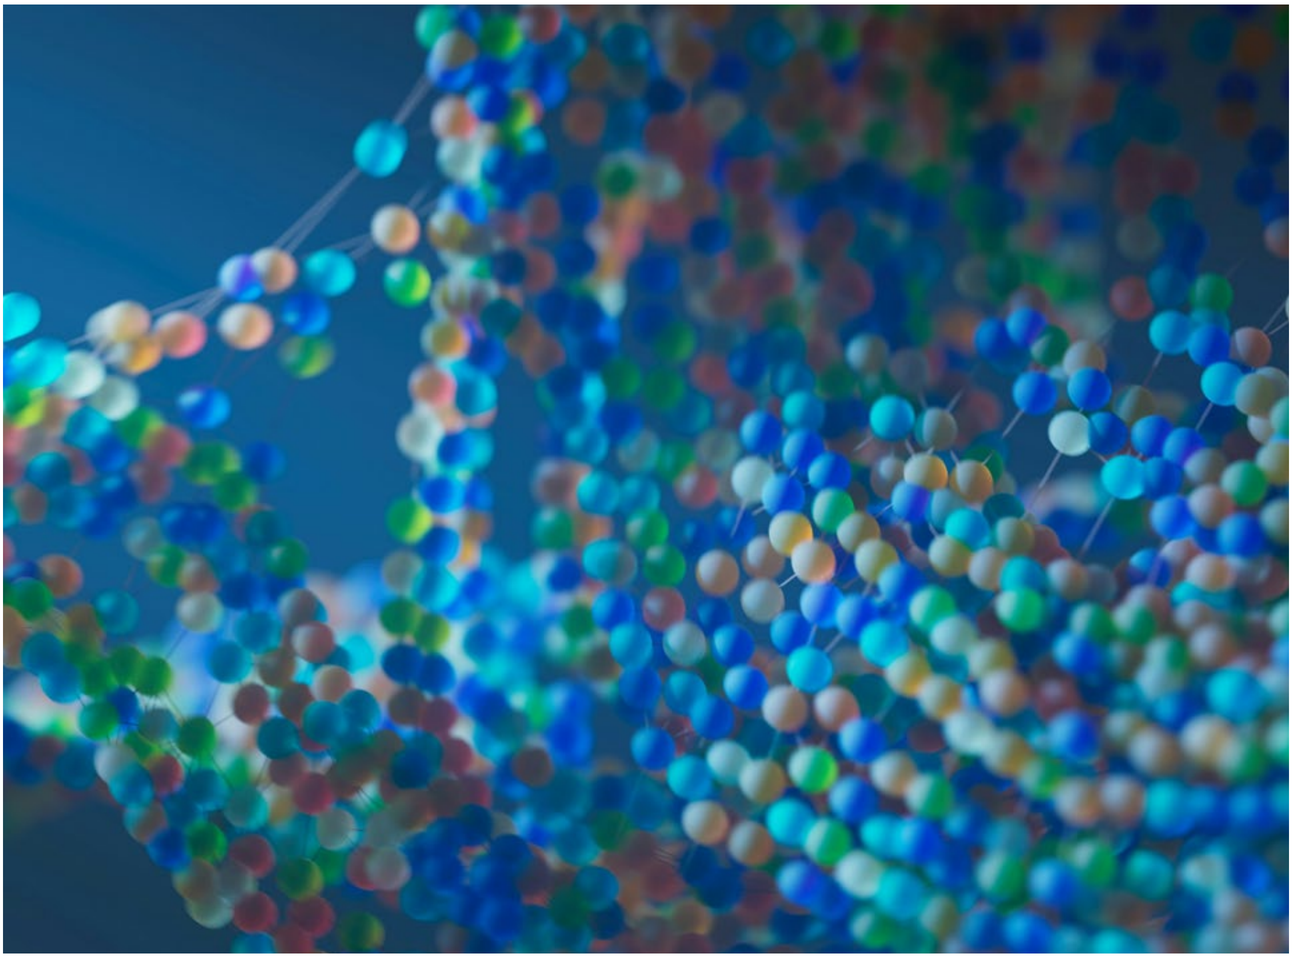

Chapter 8

AI in 2030

The survey's forward-looking questions record what financial institutions, AI vendors, and regulators expect the next five years to bring. These expectations are shaping investment decisions and regulatory priorities being made now. They also reveal where the distance is greatest between what the sector believes is coming and what its current governance infrastructure is designed to handle.

Expected impact on competition

There is a clear divergence in expectations across stakeholder groups in how AI will impact market competition. AI vendors expect most competitive disruption, with 58% expecting fragmented or balanced market dynamics, likely reflecting their own role in enabling new market entrants and challenging incumbents. 47% of industry respondents believe the market will become more competitive/balanced, while also showing the highest concern about winner-takes-all concentration (21%).

Figure 8.0: Expected impact of AI on competition in financial services by stakeholder group – industry (n=340), AI vendors (n=140), regulators (n=122)

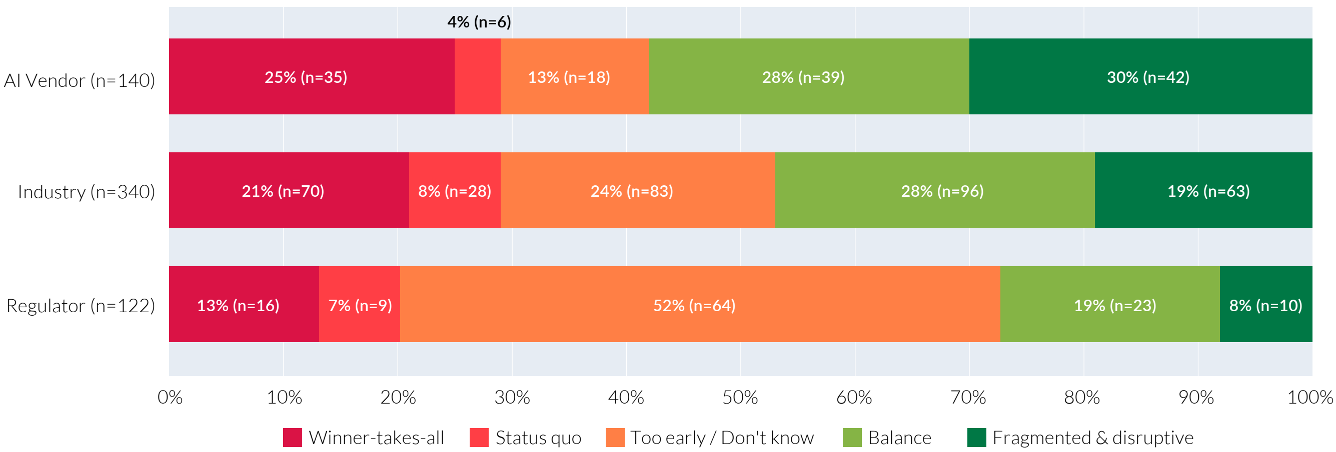

Regulators differ, with over half (52%) saying it is too early to tell or they do not know, and only 27% foresee more competitive outcomes. This caution may reflect the limited empirical evidence available to supervisors on how AI is reshaping market structures and suggests a need for better data collection on AI-driven competitive dynamics. As examined earlier in Chapter 6, merely 5% of regulators currently collect data on AI-driven market correlations, and 57% have no plans to begin within 24 months.

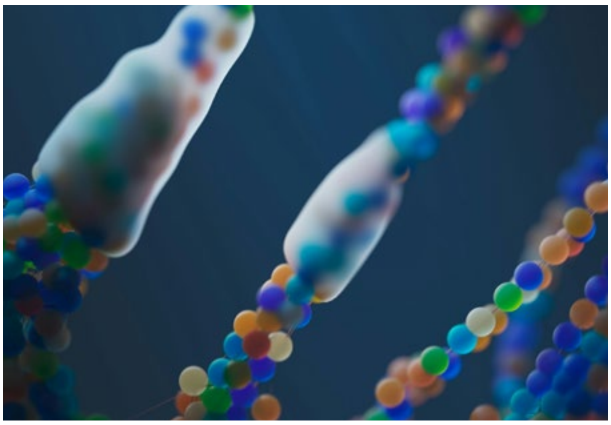

Expected AI benefits

All three stakeholder groups agree that operational efficiency is the top expected benefit of AI in financial services, though industry see this as a great benefit (73%) compared with regulators (56%).

Figure 8.1a: Expected benefits of AI in financial services by stakeholder group – industry (n=340), AI vendors (n=140), regulators (n=122)

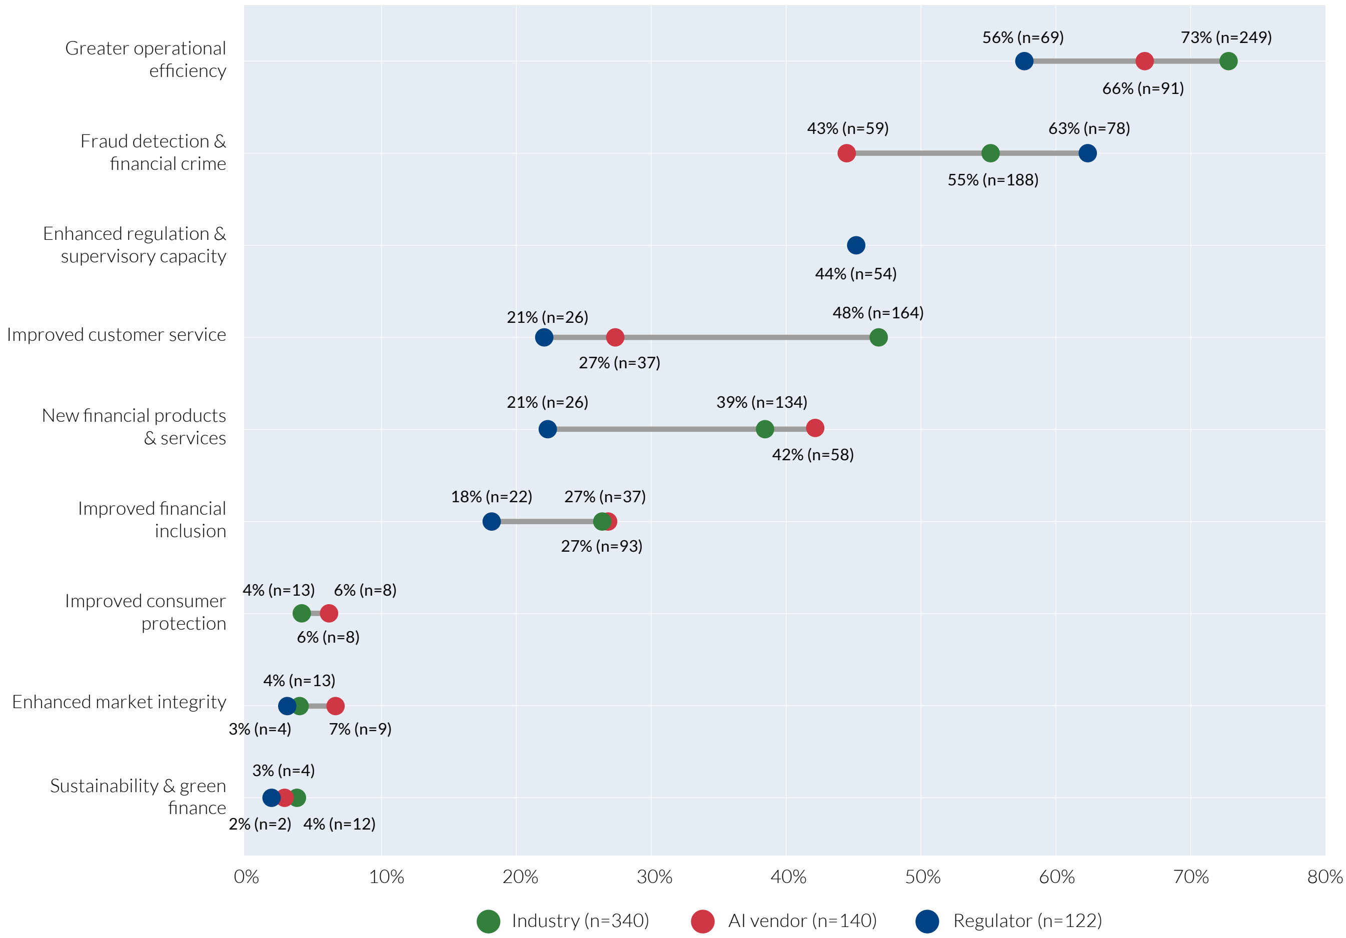

The most notable divergence in terms of perceived benefits is on fraud detection and financial crime prevention, where regulators rank it highest among their priorities (63%), even above operational efficiency. This likely reflects regulators' mandate regarding AML/CFT and their hope that AI can augment supervisory capacity in this area. Earlier findings from Chapter 2 show that fraud detection is already among the most mature AI applications in the sector, showing that this expectation has a basis in current deployments.

The World Economic Forum's 2025 report Artificial Intelligence in Financial Services highlighted that AI is increasingly being used across financial institutions to enhance fraud detection, including the ability to monitor transactions and identify suspicious activity in near-real time. This helps explain why regulators in the survey prioritised fraud detection and financial crime prevention more highly than industry respondents. Similarly, the International Monetary Fund's 2021 report Powering the Digital Economy: Opportunities and Risks of Artificial Intelligence in Finance documented the growing deployment of AI and machine learning systems in AML/CFT, market conduct supervision and regulatory compliance, where these technologies improved the efficiency, targeting and prioritisation of investigative and supervisory processes. Taken together, these findings reinforce the view that supervisors increasingly regard AI as a means to augment constrained investigative capacity amid rising and more complex financial crime risks.

Figure 8.1b: Potential benefits versus risks of AI and fraud

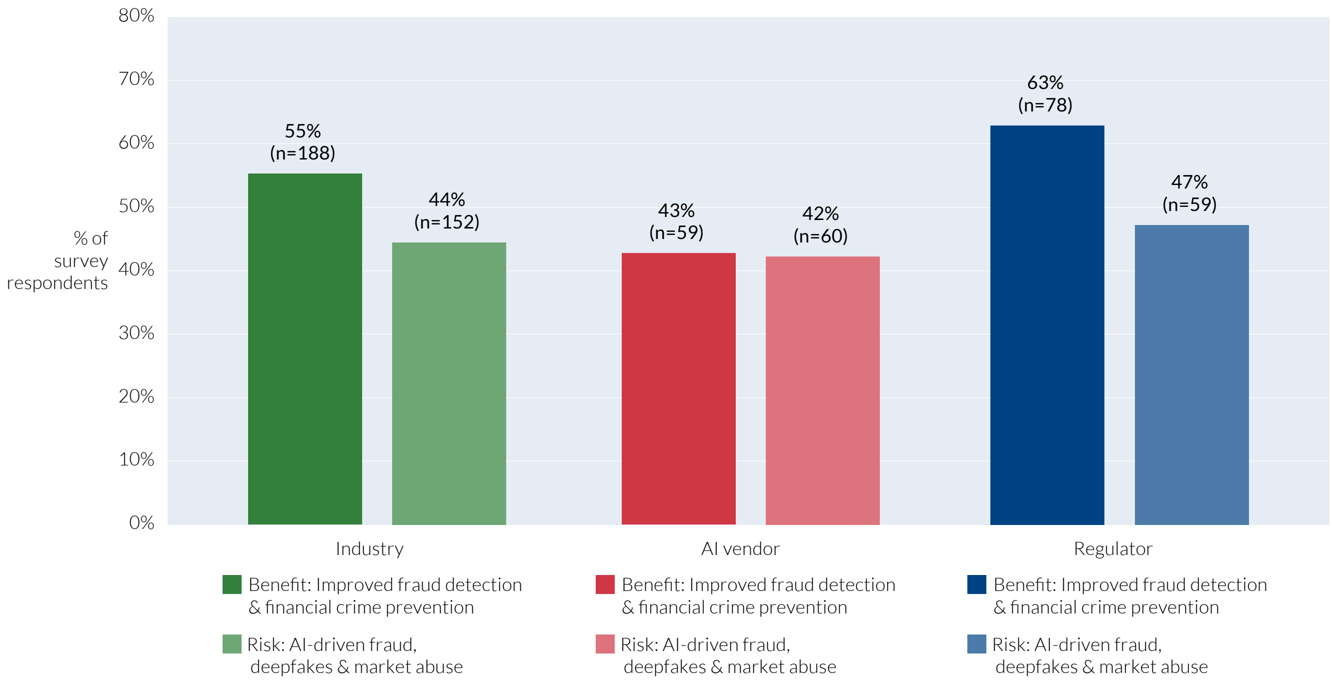

All stakeholders view AI as a net positive against financial crime, acknowledging that AI can be deployed both to commit and prevent illicit activities. Regulators show the highest net optimism (+16 points: 63% benefit versus 47% risk), followed by industry (+11 points). AI vendors show the narrowest margin (+1 point), perhaps highlighting their deep awareness of AI's offensive capabilities.

Consequently, fraud and AML remain one of the most universally supported use cases for AI adoption.

In terms of product innovation, vendors and industry align closely on the potential benefit of new financial products (42% and 39%), but regulators are far less convinced (21%), potentially reflecting their own supervisory and regulatory priorities and mandate rather AI-driven product innovation.

Despite the recognised potential of AI, all surveyed stakeholder groups identified sustainability and green finance, enhanced market integrity and improved consumer protection as much lower expected benefits by 2030. This suggests that stakeholders associate AI's value primarily in immediate operational and control-related use cases, with less of an emphasis on broader sustainability or consumer-facing outcomes.

Global cooperation for AI regulation by 2030

When looking towards 2030, regulators are cautiously optimistic about the potential for international AI cooperation. Nearly half of regulators surveyed (48%) acknowledge it is challenging today but expecting it to improve over time. A further 37% see progress as steady but slow, comparable to other regulatory areas.

Figure 8.2: Outlook on international regulatory cooperation for AI – regulators (n=118)

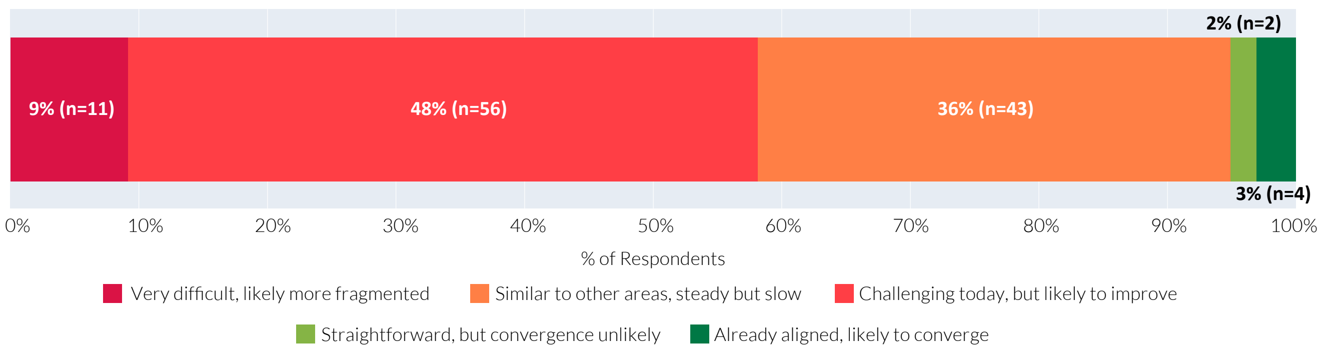

Only 9% hold a pessimistic view that cooperation will fragment further, while just 3% believe alignment is already in place. This suggests most regulators see international coordination on AI as achievable but not yet realised: a work in progress rather than a crisis. Perhaps reflecting the cross-border nature of AI risks creating a shared incentive for regulatory cooperation, even where formal alignment remains limited. Since there are already in place multiple bilateral and multilateral cooperation mechanisms between regulators, early cooperation has been possible without further arrangements. Nevertheless, further progress and learning may allow regulators to assess the need to have specific AI-related cooperation arrangements, as happened in the past with cyber security.

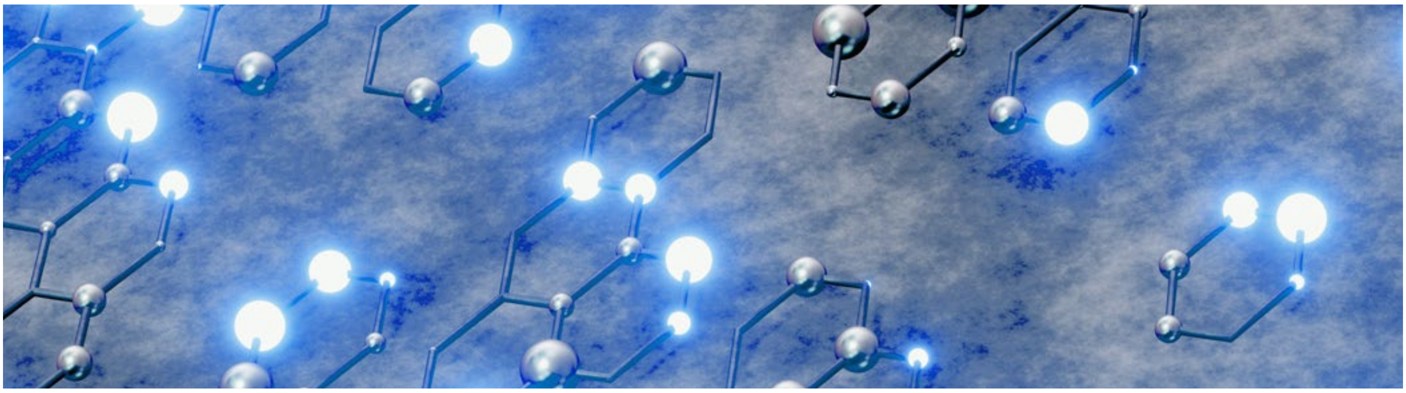

AI evolution stages: Today versus 2030

The aggregate view of 602 survey participants across industry (n=340), AI vendors (n=140) and regulators (n=122) show a dramatic expected shift in AI capabilities over the next five years. By 2030, data analytics is expected to reach near-universal deployment (88%).

However, the most dramatic growth is anticipated in agentic AI. The deployment of autonomous agents is expected to jump from 24% achieved today to a projected 81% by 2030 – the largest single increase across all categories. This signals a clear expectation that autonomous AI agents will become mainstream in financial services within five years. As documented in Chapter 2, this aspiration appears realistic given that 52% of industry respondents are already Piloting or Deploying agentic AI, with 17% of risk and compliance applications describing autonomous systems in production. The accountability framework for these systems, however, remains unresolved to some extent as examined in Chapter 5.

Figure 8.3: AI maturity: achieved today versus expected by 2030 all survey groups combined

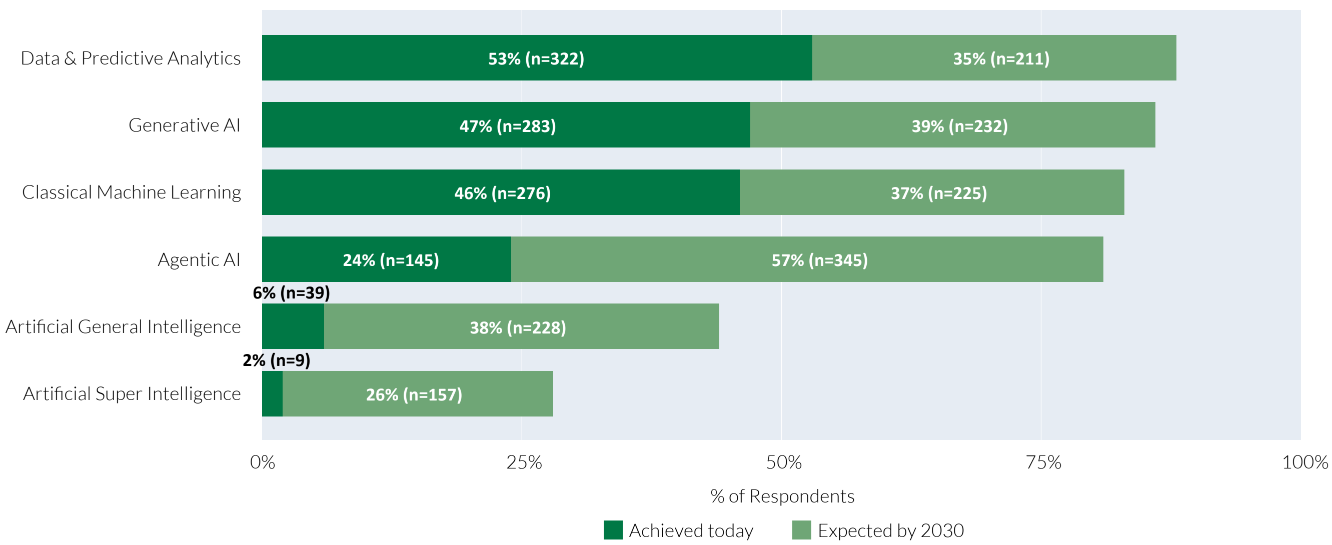

Expectations regarding AGI (where AI reaches parity with humans across all tasks) represents an interesting insight. 50% of industry respondents and 51% of AI vendors are expecting the achievement of AGI by 2030, which is quite surprising.

Figure 8.4: AGI expectation versus risk perception: expected by 2030 versus cited as top five

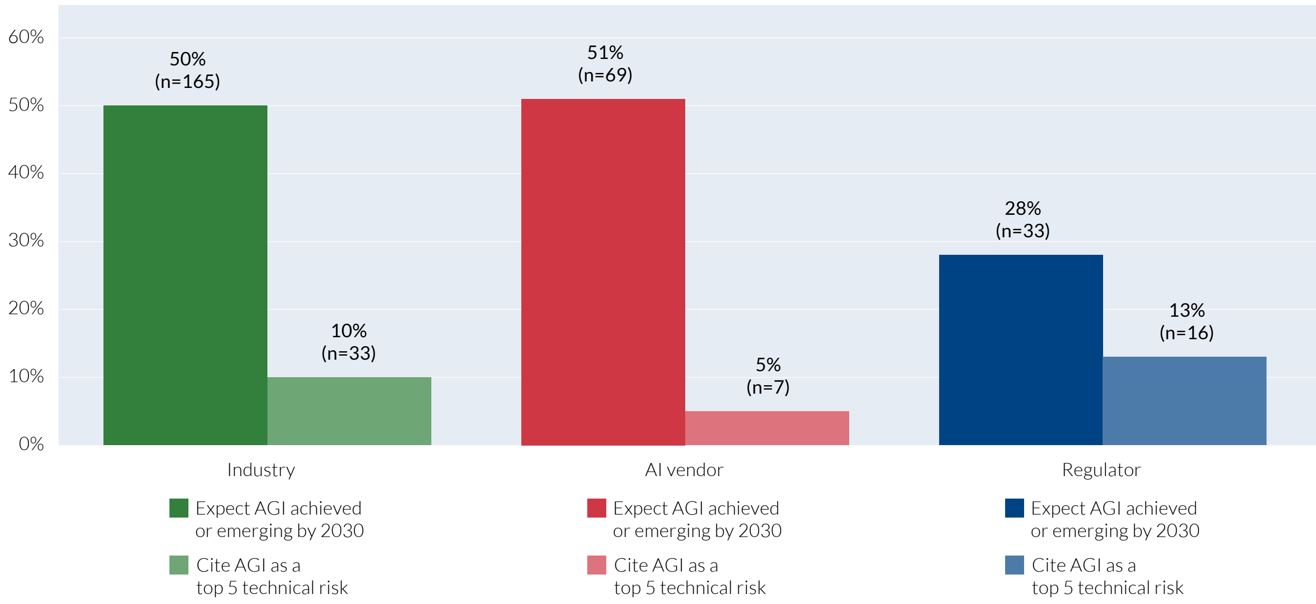

While half of industry and vendors expect AGI to be meaningfully achieved by 2030, fewer than 1 in 10 rank it among their top five risks today. This disconnect likely reflects pragmatism rather than complacency: stakeholders are focused on the immediate threats already materialising from current AI deployments, such as data privacy (80%), hallucinations (70%) and cyber threats (57%). AGI remains too abstract and uncertain to compete with these pressing operational concerns.

Figure 8.5: AI maturity: achieved today versus expected by 2030 by stakeholder group – industry (n=340), AI vendors (n=140), regulators (n=122)

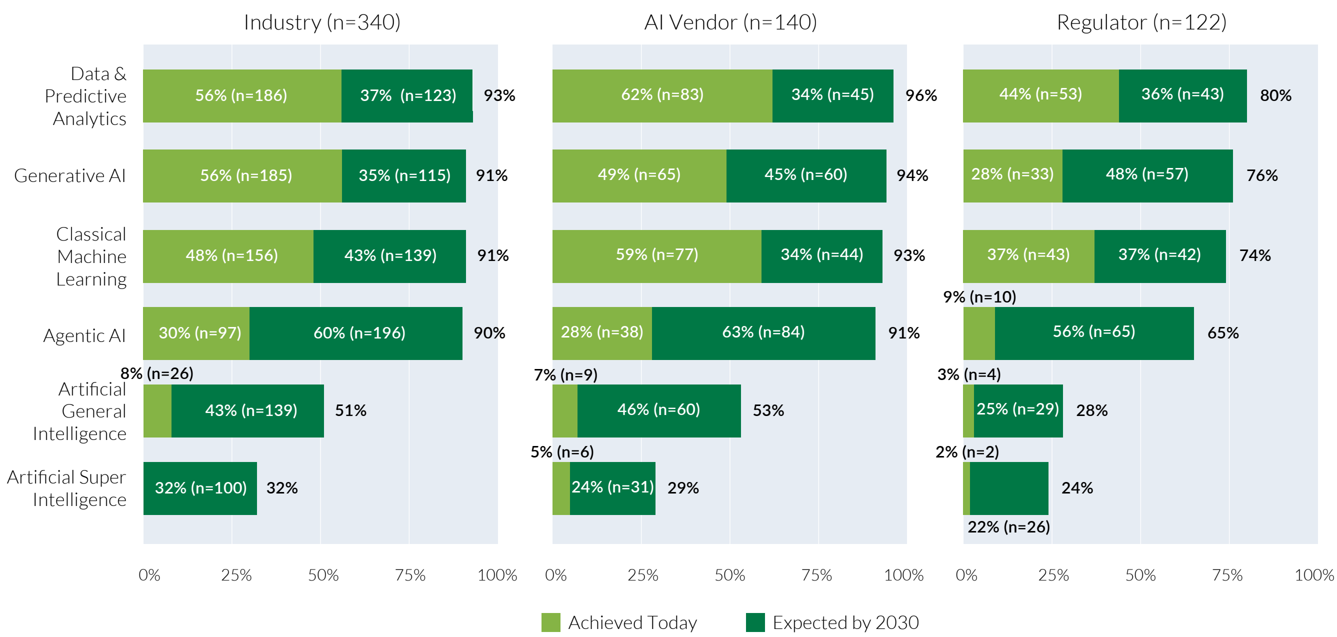

Concluding thoughts

This report documents broad but uneven AI deployment across the financial services sector. Around four in five respondent institutions now report AI in use at some level. Productivity gains are widely reported across technology, back-office, and operational functions, and a smaller but non-trivial share identify AI as having altered their organisation's strategic direction or competitive position. Analytical questions that remain concern the translation of reported productivity effects into measurable profitability, the gap between industry deployment and regulator capacity to observe it, and the adequacy of accountability and governance frameworks designed around earlier technologies for systems that now operate with significant autonomy.

Early signals that will shape the next five years: from intra-firm concerns to market-wide dynamics

The survey shows that within the broad deployment across the sector, a key gap remains between productivity gains and reported profitability increases, with 60-70% of firms still struggling to define and measure AI's value, and only around 40% reporting profitability gains on their AI investments. The highest profitability gains tend to be concentrated in specific functions, like FP&A, treasury and asset liability management (ALM), product creation, and trading and portfolio management.

Segments also differ in profitability metrics, with fintechs consistently outpacing traditional FIs. This is likely for a combination of reasons, including lower legacy infrastructure debt, different risk tolerance and lighter regulatory overhead. For established financial institutions, the question will likely be which of these constraints is most tractable, while for all firms, justifying the continued outlay for AI tooling will require increasingly strong bottom-line evidence. Another practical implication is that investment cases written in the period immediately ahead may need to carry a share of the measurement work themselves, given the absence of cross-industry benchmarks for AI-specific returns.

Agentic AI emerges as the clearest technological frontier in the survey data, with more than half of industry respondents already piloting agentic systems, and nearly nine in ten expecting meaningful deployment by 2030. Yet, where GenAI raises critical questions of explainability, agentic AI, particularly when deployed beyond a firm's own controlled environment, raises cybersecurity risk and accountability questions linked to digital identity, authorisation, and verification traceability. The survey records, however, that the appropriate framework to govern such autonomous processes has not been settled on, and views on accountability are fragmented between surveyed regulators and industry respondents. A shared framework for sound agentic practice, particularly for cross-boundary activity, is therefore a highly consequential item of governance work for the coming period. In the absence of such a framework, the gap between industry deployment and supervisory capacity is likely to continue widening.

That unresolved accountability picture has a direct parallel in how the sector views explainability. Approximately four in five surveyed regulators treat explainability as critical or important to their institutional objectives, whereas fewer than a quarter of respondent industry firms have adopted explainable AI methods and around two-thirds are not currently monitoring their AI systems for discrimination, exclusion, or systemic bias. This marks a clear gap with no obvious solution, as the technical difficulty of interpreting model outputs has significantly grown with the advent of foundation models. A possible path to resolution could entail recasting the supervisory expectation in outcome-oriented terms: behavioural testing, structured monitoring for discrimination, exclusion and systemic bias, standardised post-hoc attribution methods, and vendor-level attestation applied proportionately by use case.

Regarding expectations surrounding competitive market dynamics, the survey records a substantial revision of outlook when compared to 2020. The share of respondents expecting the market status quo to broadly prevail has fallen from nearly half to below 10%; around half of current industry respondents now anticipate some form of market consolidation by 2030, and around a fifth are concerned about winner-takes-all concentration dynamics. Lag on agentic and GenAI capability may carry future asymmetric weight in this setting: in a consolidating market, a delayed build-out may lock in disadvantage rather than defer returns, sharpening the build-versus-buy question for traditional institutions in particular.

The abovementioned expected shifts in market dynamics also map onto surveyed views on risk levels: the macro-level risks that matter most from a broader financial stability standpoint, such as algorithmic market abuse and cross-institutional concentration in AI supply chains, are precisely the areas where supervisory data collection is weakest. How well the future regulatory architecture keeps pace with market structure change, and whether it develops adequate instruments to manage systemic AI-driven risks across firm and jurisdictional boundaries, is likely to shape sector-level outcomes as much as any firm-level adoption decision.

Going beyond the data

Two factors drawn from evolutions in the broader online and AI ecosystems, while not yet showing up in the survey data, are also likely to shape the deployment of AI in financial services in the coming five years.

The first considers the impact of differing approaches to identity as finance becomes more embedded in the internet, and the second looks beyond the current frontier of AI to consider how what comes next will impact the financial services sector.

The internet has no base identity layer, having been built on anonymity or pseudonymity, whereas financial services has operated on the assumption of verified identity for several decades, reinforced through successive waves of regulatory KYC expectations. AI pulls these two infrastructures into closer contact, and the implications of that closer contact run in both directions. AI systems will make impersonation, synthetic identity, and automated fraud easier to produce at scale. However, the same technologies will also make verified-identity frameworks more capable, with the possibility of extending such frameworks to AI agents themselves. Financial services is not the only sector that will shape the resulting identity architecture, but decisions taken by financial regulators and institutions are likely to exert disproportionate influence on the digital identity standards that emerge for the wider economy, given the sector's established KYC infrastructure and regulatory weight.

In the meantime, research attention in AI is already shifting beyond generative and agentic systems. World models, which aim to build causal representations of how environments behave rather than relying primarily on pattern recognition over past examples, have plausible applications in risk management, scenario construction, and portfolio modelling. The current survey's findings bear on the sector's readiness to absorb the next wave of capability.

Workforce preparedness and data quality are the principal constraints on current AI deployment as surveyed, and these constraints will most likely also influence readiness for whichever AI paradigm may follow. Institutions that have treated current AI adoption primarily as a procurement exercise rather than a capability-building programme are likely to face the same structural constraints when the next generation of techniques reaches the sector.

Appendix

Some of the logos of participating institutions

We are incredibly grateful for the contributions of the 628 institutions that took the time to participate in this global study.

Survey respondents

Logos of association partners

Associations

List of financial authorities that participated in the survey

Afghanistan Payment System

Albanian Financial Supervisory Authority

Autoritat Financera Andorrana (AFA)

Capital Markets Commission

Banco Central de la República Argentina

Central Bank of Armenia

Ministry of Finance

Australian Competition and Consumer Commission

Oesterreichische Nationalbank

Central Bank of Bahrain

Financial Services Commission

National Bank of Belgium

Central Bank of Belize

Insurance Supervisory Agency of the Federation of Bosnia and Herzegovina

Central Bank of Botswana

Banco Central do Brasil

Commission de l'Informatique et des Libertés (CIL)

National Bank of Cambodia

Securities and Exchange Regulator of Cambodia

Canadian Investment Regulatory Organization

Ontario Securities Commission

Cayman Islands Monetary Authority

Guernsey Financial Services Commission

Banco Central de Chile

Financial Market Commission of Chile (CMF)

Banco Central de la República Colombia

Superintendencia Financiera de Colombia

Unidad de Regulación Financiera (URF)

Central Bank of Congo

Superintendencia General de Valores (SUGEVAL)

Superintendencia General de Entidades Financiera (Sugef)

Croatian National Bank (HNB)

Superintendencia del Mercado de Valores de la República Dominicana

Eastern Caribbean Central Bank

Banco Central del Ecuador

Superintendencia de Bancos

Central Bank of Egypt

Financial Regulatory Authority

Banco Central de Reserva de El Salvador

Superintendencia del Sistema Financiero

Estonian Financial Supervision and Resolution Authority (Finantsinspektsioon)

European Banking Authority (EBA)

European Central Bank (ECB)

Reserve Bank of Fiji

Bank of Finland

Insurance State Supervision Service

Data Protection Commission of Ghana

Securities and Exchange Commission

Bank of Greece

Hellenic Capital Market Commission

Bank of Guyana

Comisión Nacional de Bancos y Seguros

Hong Kong Monetary Authority

International Financial Services Centres Authority

Otoritas Jasa Keuangan (OJK)

Central Bank of Ireland

Israel Securities Authority

Banca d'Italia

Commissione Nazionale per le Societa e la Borsa (CONSOB)

Institute for the Supervision of Insurance (IVASS)

Bank of Jamaica

Central Bank of Jordan

Jordan Securities Commission

Astana Financial Services Authority

Capital Markets Authority

Central Bank of Kenya

Insurance Regulatory Authority

Central Bank of Kosovo

Central Bank of Kuwait

Banque du Liban

Capital Markets Authority

Commission de Surveillance du Secteur Financier (CSSF)

Bank Negara Malaysia (BNM)

Securities Commission

Bank of Mauritius

Financial Services Commission

La Comisión Nacional para la Protección y Defensa de los Usuarios de Servicios Financieros (CONDUSEF)

Bank of Mongolia

The Bank of Mongolia

Insurance Supervision Agency of Montenegro

Bank al Maghrib

Moroccan Capital Market Authority

Nepal Rastra Bank

The Dutch Authority for the Financial Markets (The Netherlands)

Financial Markets Authority

SIBOIF

Central Bank of Nigeria

Securities and Exchange Commission Nigeria

Insurance Supervision Agency

National Bank of the republic of North Macedonia

Securities and Exchange Commission of Pakistan

State Bank of Pakistan

Palestinian Capital Market Authority

Superintendencia de Bancos de Panamá

Superintendencia del Mercado de Valores

Superintendencia del Mercado de Valores de Panamá

Banco Central del Paraguay

Superintendencia de Banca, Seguros y AFP

Bangko Sentral ng Pilipinas

KNF

Banco de Portugal

Qatar Financial Centre Regulatory Authority

National Bank of Rwanda

Capital Market Authority

Saudi Central Bank

National Bank of Serbia

Securities Commission

Central Bank of Seychelles

Financial Sector Conduct Authority

South African Reserve Bank

Comisión Nacional del Mercado de Valores

Central Bank of Sudan

Centrale Bank van Suriname

Bank of Tanzania

Bank of Thailand

TTSEC

Central Bank of Tunisia

Conseil du Marché Financier

Turks & Caicos Islands Financial Services Commission

Central Bank of the United Arab Emirates

Dubai Financial Services Authority (DFSA)

Securities and Commodities Authority (SCA)

Bank of England

Financial Conduct Authority

Banco Central del Uruguay

The State Bank of Vietnam

Central Bank of Yemen

Bank of Zambia

Survey and cross tabulation analysis

Links to the three surveys can be found here:

1. **The Regulator Survey**: https://cambridge.eu.qualtrics.com/jfe/form/SV_0Aio1xTyA84p7Js
2. **The AI Vendor Survey**: https://cambridge.eu.qualtrics.com/jfe/form/SV_1YV0h35fKjXsuQm
3. **The Industry Survey**: https://cambridge.eu.qualtrics.com/jfe/form/SV_7O4Ou3RmktdLE10

Twelve binary segmentations were applied to the industry dataset. These cuts of the data provided sizable sample sizes, enabling systematic comparison of responses across key organisational characteristics. These segments included:

- Institution type (fintech versus traditional FIs)
- Economic development (advanced versus emerging and developing markets)
- Geographic regional base
- AI maturity (Transforming and Scaling versus Exploring and Piloting)
- AI spend (above versus below USD 100,000 in the most recent financial year)
- Firm size (above versus below 50 FTE in the most recent financial year)
- Revenue (above versus below USD 10 million in the most recent financial year)
- Profitability impact (increased versus no change)
- Productivity impact (positive versus negative or no change)
- Workforce preparedness (moderate/highly prepared versus limited/not prepared)
- Value measurement difficulty (challenging versus easy/neutral)
- AI offering type (tailored financial services versus general-purpose)

Equivalent segmentations were applied to the vendor and regulatory datasets where applicable, including economic development, firm size, revenue, AI staff, vendor offering type, regulator institution type (central bank versus non-central bank) and regulator AI maturity.

This study looks at the data from several angles to understand how AI is being adopted and experienced across the financial sector. Most firms are still in the early stages of their AI journey, with 60% classified as Limited, Exploring or Piloting. 40% are already Scaling or Transforming with AI. Overall, AI maturity is a useful way to analyse the data and this cut is used throughout this study.

Profitability offers another important point of comparison. While 40% of respondents report positive profitability, 43% see no change yet. Productivity gains are widely reported, with 69% in AEs and 61% in EMDEs seeing improvements even if these have not yet translated into profits. At the same time, 53% of firms struggle to measure AI's value, and larger firms are more heavily represented in the sample of the financial industry at 63%, while AI vendors are represented by smaller firms at 56%.

Investment in AI is another important way to compare the survey respondents. About 53% of firms are spending less than USD 100,000 in the most recent financial year on AI, while 45% are investing more. This broadly aligns with maturity levels as most firms are still making relatively small, exploratory investments, while higher spending tends to be linked to more advanced adoption and better outcomes.

The data segments are applied throughout this report to help make comparisons between these characteristics.

Figure A.1: Industry Data Segments

Abbreviations

AAI: AI Adoption Index – an IMF proposed index of AI Adoption

AE: advanced economies

AGI: artificial general intelligence

AI: artificial intelligence

AML/CFT: anti-Money laundering and counter-financing of terrorism

AMF: Arab Monetary Fund

APAC: Asia-Pacific regional grouping

AWS: Amazon Web Services

BIS: Bank for International Settlements

CCAF: Cambridge Centre for Alternative Finance at the University of Cambridge, Judge Business School

CNN: convolutional neural network

CRM: customer relationship management

EMDE: emerging markets and developing economies

EU: European Union

FCA: Financial Conduct Authority, UK

Fii: Financial Innovation for Impact

FIs: financial institutions

FM: foundation model

FP & A: Financial Planning and Analysis

FSB: Financial Stability Board

FTE: full-time employee

G20: Group of Twenty

GPUs: graphical processing units – needed as part of the AI supply chain

HKMA: Hong Kong Monetary Authority

HQ: headquarters

HR: human resources

IDB: Inter-American Development Bank

IMF: International Monetary Fund

KYC: Know Your Customer

LAC: Latin America and the Caribbean

LIME: Local Interpretable Model-agnostic Explanations

MAS: Monetary Authority of Singapore

MENA: Middle East and North Africa

MIT CISR: MIT Center for Information Systems Research

ML: machine learning

NLP: natural language processing

OECD: Organisation for Economic Cooperation and Development

RAG: retrieval-augmented generation

RDFE: Responsible Digital Finance Ecosystem – a CGAP framework for consumer protection

RNN: recurrent neural network – one of the examples of deep learning architectures

ROI: return on investment

SHAP: Shapley Additive Explanations

SSA: Sub-Saharan Africa

USD: United States dollars

WEF: World Economic Forum

Glossary and definitions

Advanced economies (AE): A macroeconomic classification based on World Bank income categories that refers to high-income economies.

Agentic AI: A set of autonomous, multi-step AI systems capable of executing complex workflows without continuous human intervention.

Artificial general intelligence (AGI): Defined as AI reaching parity with humans on all tasks or achieving human-level capability across domains.

AI: Artificial intelligence: Defined by the OECD as a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.

AI maturity stages: The study categorises organisational AI adoption into four progressive stages: Exploring, Piloting, Scaling and Transforming.

Back office: Financial services functions dedicated to execution, operational, and analysis-based activities.

Case-by-case liability: An accountability approach for AI-related harm where the allocation of responsibility is based on specific circumstances and is typically defined contractually.

Classical machine learning (ML): Encompasses well-established AI methods and applications including predictive models, such as fraud detection algorithms and credit scoring engines, as well as time-series forecasting and unsupervised learning.

Collective forgetting: A systemic risk tied to automation, where organisations can lose the institutional memory and capabilities to execute processes manually if required.

Emerging Markets and Developing Economies (EMDEs): A macroeconomic classification based on World Bank income categories encompassing low-income, lower-middle-income, and upper-middle-income economies.

Explainability/explainable AI (XAI): Deals with how stakeholders, including laypeople and non-technical experts, understand the results and reasoning outputs of AI systems. This often involves surrogate models (for example, SHAP and LIME) that estimate the behaviour of opaque black box models.

Fintechs: Digital-first, more agile adopters of new technologies within the financial sector.

Foundation model: A large general-purpose AI model trained on vast data that can be adapted for many downstream tasks.

Front office: Financial services functions where client interaction and commercial decisions take place.

Generative AI (GenAI): A category of AI that gained significant traction following the launch of ChatGPT. It is typically provided as packaged solutions from foundation model vendors, requiring little to no AI engineering capabilities.

Hallucinations: Unreliable results and outputs generated by AI models.

Human-in-the-loop: An AI autonomy level where the AI system acts but requires human approval.

Human-on-the-loop: An AI autonomy level where the AI system acts autonomously, but a human can intervene.

Human-out-of-the-loop: A fully autonomous level of AI operation.

Human support: An AI autonomy level where the AI recommends an action, but the human decides.

Interpretability: Refers to how well users can understand the inner workings of AI models.

Joint liability: An accountability framework where any party can be held fully responsible for the entire AI-related harm.

Open-weight models: AI models (such as those offered by DeepSeek and Meta) whose weights are openly accessible, as opposed to closed proprietary models offered by companies like OpenAI, Anthropic or Google. These cost-accessible models can lower cost barriers, particularly for lower-resource institutions and firms in emerging markets and developing economies (EMDEs).

Shared liability: An accountability framework where each party is responsible only for its specific share of the AI-related harm.

SupTech (Supervisory technology): The use of technology, including AI, by financial regulators and supervisors to support their oversight mandates.

Endnotes

1. Cambridge Centre for Alternative Finance (2020). Transforming Paradigms: A Global AI in Financial Services Survey. Available at: https://www.jbs.cam.ac.uk/faculty-research/centres/alternative-finance/publications/transforming-paradigms/
2. World Bank (2025). World Bank country classifications by income level for 2024-2025. Accessed at https://blogs.worldbank.org/en/opendata/world-bank-country-classifications-by-income-level-for-2024-2025#:~:text=The%20World%20Bank%20Group%20assigns%20the%20world%E2%80%99s%20economies-,GNI%20per%20capita%20of%20the%20previous%20calendar%20year.
3. See https://ccaf.io/atlas/methodology
4. NVIDIA (2026). State of AI in Financial Services: 2026 Trends. Available at: https://resources.nvidia.com/en-us-2026-fsi-survey/finance-state-of-ai-report
5. RAND Corporation (2025). Artificial Intelligence Adoption and Sectoral Transformation: Implications for Health Care, Financial Services, Climate and Energy, and Transportation. Available at https://www.rand.org/pubs/research_reports/RRA3888-1.html.
6. Khanfar, S., Al-Hajj, A., Kyle, Z., Nijim, N., & Abusaada, I. (2025). Investigation into the Factors Affecting the Integration of Artificial Intelligence (AI) in Human Resources Management (HRM) within Jordanian Banks. In 2025 International Conference on Metaverse and Current Trends in Computing (ICMCTC) (pp. 1-9). IEEE.
7. Change to World Economic Forum 2025, Artificial Intelligence in Financial Services. WHITE PAPER https://reports.weforum.org/docs/WEF_Artificial_Intelligence_in_Financial_Services_2025.pdf
8. Schönberger, M. (2023). Artificial intelligence for small and medium-sized enterprises: Identifying key applications and challenges. Journal of Business Management, 21, 89-112.
9. OECD (2022). Framework for the Classification of AI systems, 2022 (revised 2024). Available at https://www.oecd.org/en/publications/oecd-framework-for-the-classification-of-ai-systems_cb6d9eca-en.html
10. Anthropic (2026). Anthropic Economic Index report: Economic primitives. Available at https://www.anthropic.com/research/anthropic-economic-index-january-2026-report
11. FSB (2024). The Financial Stability Implications of Artificial Intelligence. Available at https://www.fsb.org/uploads/P14112024.pdf
12. Prepared by Vatsala Shreeti, Kumar Rishabh and Jon Frost. The views expressed here are those of the authors and do not necessarily reflect those of the Bank for International Settlements (BIS).
13. Gambacorta, L., Shreeti, V. (2025). The AI supply chain. Available at https://www.bis.org/publ/bppdf/bispap154.htm
14. Frost, J., Shreeti, V., Rishabh, K. (2026). Global giants in the AI supply chain, Available at https://www.bis.org/publ/bisbull122.htm
15. Valuations are in USD and rounded. Jurisdiction in the legend refers to headquarters location. If a firm both sells and uses a capability internally, it is coded as marketed. Valuations are whole-firm and are not allocated across layers. * = privately held firm. Sources: PitchBook (public market capitalisations, private company valuations); authors' coding of supply chain roles from publicly available information on firm products and services.
16. The supply chain is the network of individuals and organisations involved in creating hardware and software components of AI and delivering them to end users. The value chain can be defined as the sequence of interconnected stages – production, distribution, delivery and consumption – that are needed to produce AI solutions, ie the steps in which value is created in producing and applying AI.
17. See Franco, C., Frost, J., Mishra, S., Müller, C., and Velásquez, C. (2026), 'Artificial intelligence policy in emerging market and developing economies: growth opportunities in the AI value chain', mimeo; Gambacorta, L., Kharroubi, E., Mehrotra, A., and Oliviero, T. (2025), 'Artificial intelligence and growth in advanced and emerging economies: short-run impact', BIS Working Papers, no 1321, December.
18. McKinsey & Company (2023). The state of AI in 2023: Generative AI's breakout year. Available at: https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2023-generative-ais-breakout-year
19. OECD (2025). Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions. Available at: https://www.oecd.org/en/publications/governing-with-artificial-intelligence_795de142-en/full-report/implementation-challenges-that-hinder-the-strategic-use-of-ai-in-government_05cfe2bb.html?
20. WEF (2025). The future of jobs report 2025. Available at https://www.weforum.org/publications/the-future-of-jobs-report-2025/digest/
21. OECD (2024) Artificial intelligence and the changing demand for skills in the labour market. OECD Artificial Intelligence Papers No. 14. Paris: OECD Publishing.
22. Ferri de Camargo Paes, F., Propson, D., Suresh, K., Zhang, B., et al. (2025). Future of Global Fintech: From Rapid Expansion to Sustainable Growth – 2nd Edition. Available at SSRN: https://ssrn.com/abstract=5337491 or http://dx.doi.org/10.2139/ssrn.5337491
23. Boukherouaa, EB., Shabsigh,G., AlAjmi, K., Deodoro, J., Farias, A., Et al. (2021). Powering the Digital Economy: Opportunities and Risks of Artificial Intelligence in Finance. Available at https://doi.org/10.5089/9781589063952.087

24. ibid.
25. Anthropic, 2026, Mythos Preview and related cyber security issues https://red.anthropic.com/2026/mythos-preview/
26. Dib, L. (2024). Formal definition of interpretability and explainability in XAI. In Intelligent Systems Conference (pp. 133-151). Cham: Springer Nature Switzerland.
27. Evite, P. M., Svetlova, E., & Bucur, D. (2026). Trade-offs in financial AI: Explainability in a trilemma with accuracy and compliance. Available at https://doi.org/10.48550/arXiv.2602.01368
28. ibid.
29. Černevičienė, J., Kabašinskas, A. (2024). Explainable artificial intelligence (XAI) in finance: a systematic literature review. Artif Intell Rev 57, 216. Available at https://doi.org/10.1007/s10462-024-10854-8
30. Langer, M., et al. (2021). What do we want from Explainable AI? A survey of stakeholder perspectives. Communications of the ACM, 64(11), 42–51.
31. Project Noor: explaining AI models for financial supervision. Available at https://www.bis.org/about/bisih/topics/suptech_regtech/noor.htm
32. BIS – other AI-related projects By Bénédicte Nolens and Jack Lee, BIS Innovation Hub. The BIS Innovation Hub has developed a suite of projects leveraging artificial intelligence across use cases. In financial crime, Project Aurora applies AI, machine learning and privacy-enhancing technologies to enable cross-border collaborative analytics for improved money laundering detection, while Project Nadim explores AI-driven sharing of fraud patterns to strengthen international detection efforts. In climate and sustainability, Project Gaia uses AI and large language models to extract and standardise climate risk data from disclosures, and Project Symbiosis applies AI and big data to enhance supply chain sustainability insights. In monetary policy, Project Neo applies advanced data science and machine learning to novel high-frequency data for real-time economic insights, and Project Spectrum combines generative AI and machine learning to improve large-scale price classification for inflation nowcasting.
33. Rwanda National AI Policy. Available at https://www.minict.gov.rw/ai-policy
34. Irving Fisher Committee on Central Bank Statistics and Bank for International Settlements (2025). Governance and implementation of artificial intelligence in central banks (IFC Report No. 18). Available at https://www.bis.org/ifc/publ/ifc_report_18.pdf
35. FSB (2025). Monitoring Adoption of Artificial Intelligence and Related Vulnerabilities in the Financial Sector. Available at https://www.fsb.org/uploads/P101025.pdf
36. BIS (2024). Regulating AI in the financial sector: recent developments and main challenges. Available at https://www.bis.org/fsi/publ/insights63.pdf. However, it is also the most challenging to measure and achieve (pg 18, para 32, FSI Insights on policy implementation No 63)
37. OECD (2024). Regulatory approaches to Artificial Intelligence in finance. OECD Artificial Intelligence Papers, No. 24. Available at https://doi.org/10.1787/f1498c02-en.
38. BIS (2024). Regulating AI in the financial sector: recent developments and main challenges. Available at https://www.bis.org/fsi/publ/insights63.pdf?utm_
39. Bank for International Settlements (2025). Governance and implementation of artificial intelligence in central banks. Available at https://www.bis.org/ifc/publ/ifc_report_18.pdf
40. BIS (2025). The use of artificial intelligence for policy purposes, Available at https://www.bis.org/publ/othp100.pdf
41. ibid.
42. ibid.
43. Boeddu, G, Feyen, E, Martinez Jaramillo, S, Mesquita, S, Palta, Y, Sarkar, A, Sinha, S and Gutiérrez Traverso, A (2025), "Artificial intelligence for financial sector supervision: An emerging market and developing economies perspective", World Bank Prosperity Insight. https://documents1.worldbank.org/curated/en/099110525115015626/pdf/P180967-538291e2-656d-4682-adba-4f314244f6fd.pdf

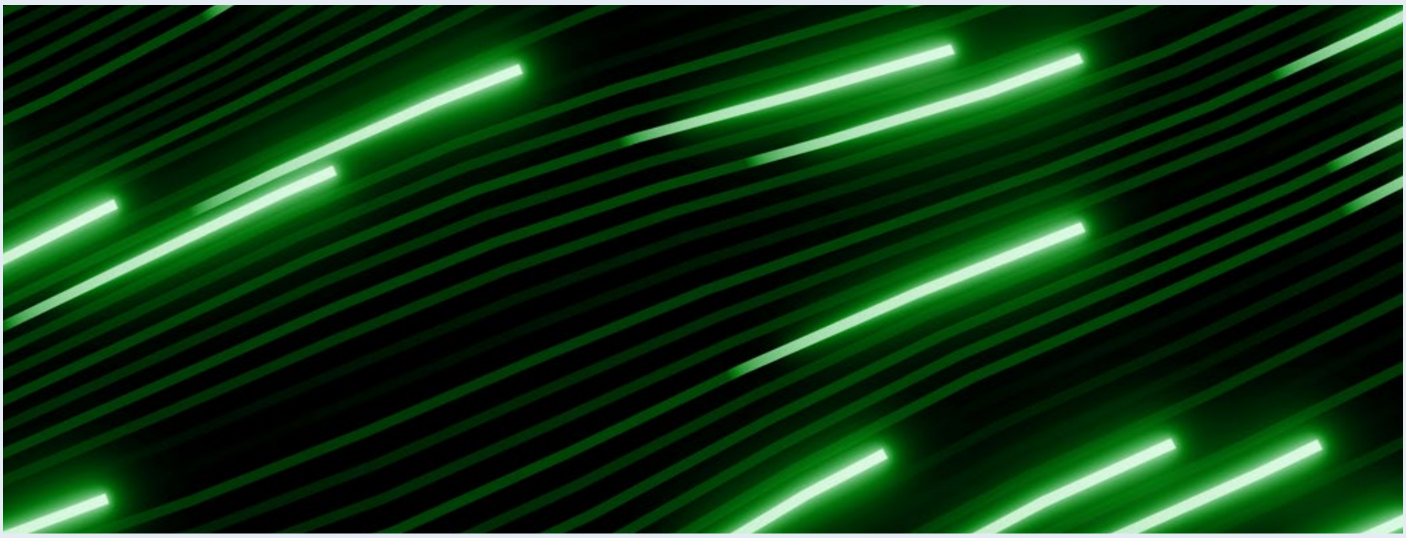

Cambridge
Centre
for Alternative
Finance

UNIVERSITY OF
CAMBRIDGE
Judge Business School

In partnership with

BIS

INTERNATIONAL MONETARY FUND

WORLD ECONOMIC FORUM

IDB

CGAP

صندوق النقد العربي
ARAB MONETARY FUND

In collaboration with

Financial Innovation for Impact

With the support of

Foreign, Commonwealth & Development Office