



BANK OF KOREA



WORLD BANK GROUP

AI Primer for Public Investors:

From AI Awareness to AI Readiness

AI Primer for Public Investors:

From AI Awareness to AI Readiness

© 2026 INTERNATIONAL BANK
FOR RECONSTRUCTION AND
DEVELOPMENT / THE WORLD
BANK GROUP

1818 H Street NW
Washington DC 20433
202-473-1000
www.worldbank.org

RIGHTS AND PERMISSIONS

The material in this work is subject to copyright. Because The World Bank Group encourages the dissemination of its knowledge, this work may be reproduced, in whole or in part, for noncommercial purposes as long as full attribution to this work is given.

Any queries on rights and licenses, including subsidiary rights, should be addressed to:

World Bank Group Publications
The World Bank
1818 H Street NW
Washington, DC 20433, USA
pubrights@worldbank.org

COVER PHOTO

© NeoLeo / Adobe Stock / stock.adobe.com.
Used with permission; further permission
required for reuse.



AI Primer for Public Investors: From AI Awareness to AI Readiness

Copyright © 2026

International Bank for Reconstruction and Development / The World Bank Group

If the Work is displayed in a digital format that does not use a traditional copyright page, the Parties will ensure that the copyright notice, disclaimers and CC-BY-NC 3.0 IGO license appear prominently in a manner appropriate to that format.

DISCLAIMER

This work is a product of the staff of The World Bank with external contributions. The findings, interpretations, and conclusions expressed in this work do not necessarily reflect the views of The World Bank, its Board of Executive Directors, or the governments they represent, or those of any external contributors.

The World Bank and any of its external contributors do not guarantee the accuracy, completeness, or currency of the data included in this work and does not assume responsibility for any errors, omissions, or discrepancies in the information, or liability with respect to the use of or failure to use the information, methods, processes, or conclusions set forth. The boundaries, colors, denominations, links/footnotes, and other information shown in this Work do not imply any judgment on the part of The World Bank or any external contributor concerning the legal status of any territory or the endorsement or acceptance of such boundaries.

Nothing herein shall constitute or be construed or considered to be a limitation upon, or waiver, of the privileges and immunities of The World Bank, all of which are specifically reserved.


AUTHORS

Parameswari Arunasalam, Eric Bouyé, Joey Hyun Joon Lee, and Ri Yu (all World Bank Group Treasury), and Byunghyun Choi (Bank of Korea).

ACKNOWLEDGEMENTS


This primer is a collaborative product of the Advisory & Partnership Department of the World Bank Group Treasury and Bank of Korea (BOK) Reserve Management Group.

The authors are grateful for valuable feedback from Katarzyna Zajdel-Kurowska, Stela Mocan, Uttandaraman Pugalumperumal, Patrick Cheng, Juan Carlos Quintero, and Mora Pearl Farhad. The authors also thank Zarina Dela Cruz for design and layout assistance.



CONTENTS

Foreword	1
Executive Summary	4
1. Introduction	7
2. AI Trends in Asset Management Industry	11
2.1. AI adoption in asset management	11
2.2. Areas of AI Applications	15
2.2.1. Productivity Gain	15
2.2.2. Investment Strategies	19
3. Responsible AI Adoption	27
3.1. Why Responsible AI Adoption	27
3.2. Vision and Leadership	29
3.3. Strategic Foundation	31
3.3.1. AI Principles	31
3.3.2. AI Institutional Policies and Standards	32
3.3.3. AI Legal and Regulatory Compliance	32
3.3.4. AI Governance	33
3.3.5. AI Risk Management	35
3.4. Enabling Environment for AI Adoption	37
3.4.1. Data and Infrastructure Foundation	42
3.4.2. Privacy and Cybersecurity	44
3.4.3. Talent and Cultural Readiness	44
3.4.4. Operationalizing AI Responsibly	46
4. Conclusions	49
References	51



FOREWORD

Public asset managers entrusted with safeguarding public wealth face an operating environment dramatically reshaped by macroeconomic uncertainty, accelerated digitalization, and the proliferation of unstructured data sources — all of which challenge traditional asset management business processes.

Artificial intelligence (AI) sits at the intersection of these changes. It is no longer a nascent technology on the horizon; it has moved from experimentation to production and become an integral part of the evolving financial ecosystem, reshaping asset management processes on information extraction, risk management, and decision-making. Automation and efficiency gain driven by this transformative technology also makes robust governance an essential requirement for public investors to ensure prudence and transparency when adapting AI.

The AI primer for public investors draws on a joint market outreach by the World Bank Group Treasury and Bank of Korea Reserve Management Group, based on a survey and interviews with global asset managers, and highlights key AI trends and real-world use cases. The paper also addresses the distinct challenges public institutions face in adopting AI and underscores the key building blocks and practical considerations that can help institutions at various stages of their AI journey.

The World Bank Group is the trusted partner for public asset managers and has shared its expertise in public asset management with central banks for over four decades. Today, the World Bank Group Treasury manages US \$300 billion in internal and external assets, making it the one of the largest external asset manager in the international development community. A key element of this work is the World Bank Group Reserve Advisory & Management Partnership, a partnership that connects over 100 public asset management institutions managing over US \$2 trillion in assets. Since 2001, RAMP has trained over 10,000 professionals and advised more than 100 institutions, helping central banks, sovereign wealth funds, public pension funds, and deposit insurance corporations to safeguard national and intergenerational wealth.

This AI primer represents an important contribution to our collective commitment to knowledge sharing and prepares public investors to harness AI and adapt to innovation. We are grateful to Bank of Korea for its partnership. Their deep institutional knowledge has greatly enriched the content and perspective presented here. Together, we continue to raise the standards of public asset management and to strengthen our stewardship of public assets while upholding the public trust at the core of our work.

Jorge Familiar
Vice President & Treasurer | WORLD BANK GROUP



FOREWORD

Global financial markets continue to be marked by heightened volatility and uncertainty, making it increasingly challenging to rely solely on conventional investment frameworks. In this environment, artificial intelligence (AI) is rapidly emerging as a core capability across the asset management process. Its role now extends well beyond task automation, with applications increasingly spanning information extraction, market analysis and forecasting, investment decision-making support, and risk management.

Against this backdrop, “AI Primer for Public Investors” offers a timely and practical reference for public asset managers preparing for the AI era. Prepared through collaboration between the World Bank Group Treasury and the Bank of Korea Reserve Management Group, this primer provides a balanced overview of the current state and limitations of AI adoption among global asset managers. It further addresses issues of particular importance to public institutions, including governance, risk management, and human oversight.

As Korea’s foreign reserves have grown, the scale and role of the Bank of Korea’s foreign reserves management have expanded significantly. The Bank of Korea has therefore come to perform not only the role of a monetary authority responsible for managing emergency funds for use in times of crisis, but also the function of an asset manager entrusted with managing large-scale foreign currency assets safely and efficiently.

In this context, the Bank of Korea has sought to diversify its foreign reserves across currencies, products, and investment strategies, while adapting to changes in global financial markets through expanded external fund management and ESG investing. In line with the AI-driven transformation of the financial environment, in January 2026 the Bank of Korea completed the development of Bank of Korea Intelligence (BOKI), a sovereign AI platform specialized in finance and economics, through a public-private collaboration. Building on this initiative, the Bank of Korea continues to explore ways to further strengthen its digital capabilities and advance innovation across its business processes.

The Bank of Korea has maintained close cooperation with the World Bank Group in the field of public asset management. The joint work on this primer has been a meaningful process for the two institutions, bringing together their respective experience and expertise in considering the future direction of public asset management in the AI era. In particular, the World Bank Group Treasury’s asset management know-how and global network contributed greatly to enhancing the depth and practical relevance of this report. As AI technologies and their application in asset management continue to evolve rapidly, we hope this primer will serve as a foundation for continued dialogue, learning, and knowledge sharing among public asset managers seeking to adopt AI responsibly.

Hee Sup Jung

Director General, Reserve Management Group | BANK OF KOREA





EXECUTIVE SUMMARY

Joint market outreach by the World Bank Group Treasury and Bank of Korea Reserve Management Group, which included surveys and interviews with global asset managers, revealed several notable industry trends. Most leading asset management firms have established AI governance frameworks over the past two to three years. However, the majority of asset managers and institutional investors remain at an early stage of AI adoption.


Current AI use focuses primarily on productivity gains through widely accessible generative AI tools, whereas applying AI to investment strategies remains more challenging due to the complexity of financial markets. Notably, firms consistently emphasize the importance of maintaining human oversight in AI implementation. Among asset managers, the most prevalent applications for productivity gain include automating data extraction from documents, improving reconciliation processes, and drafting emails, research papers, and RFP/DDQ (Request for Proposal/Due Diligence Questionnaire) responses. For these use cases, asset managers primarily leverage generative AI tools from third-party vendors rather than building their own.

Investment strategies are another key area of AI application for many asset managers, encompassing a broad range of functions, from market intelligence tasks such as sentiment analysis and clustering to market factors forecasting, scenario analysis with synthetic data, portfolio construction, manager selection, and trading execution. Firms apply techniques such as ensemble methods and explainable AI to address overfitting problems and challenges in the explainability of complex AI models in deep learning and reinforcement learning.

Responsible AI adoption entails an institutional transformation rather than a simple technology project. AI applications introduce new and distinct risks from biased data, black-box models, and model drift or manipulation, and they can undermine legal compliance, policy objectives, and public trust. Thus, public institutions must establish appropriate AI frameworks at an institutional level to begin their AI journey.

The framework for responsible AI adoption consists of three interrelated building blocks. First, a clear vision and strong leadership engagement are essential starting points, as they guide and support AI initiatives. Second, strategic foundations translate vision and leadership intent into actionable principles, policies, legal compliance, governance frameworks, and tailored risk management practices. Lastly, an enabling environment operationalizes these strategies through appropriate data and infrastructure, talent and cultural readiness, privacy and cybersecurity tools, and well-defined operational processes.

AI adoption is no longer a question of whether to proceed, but of how to implement AI effectively and prudently. While a wide range of potential AI applications is highlighted, institutions must assess their feasibility and potential impact and prioritize them accordingly. As institutions move from experimentation toward more systematic adoption, a robust framework should guide this journey, ensuring alignment with institutional mandates, accountability, lifecycle oversight, collaboration, and strong enabling environments. As AI evolves rapidly, continuous learning and peer-to-peer knowledge sharing are essential for navigating technological and regulatory change. We hope this primer serves as a useful reference for understanding the overall picture of AI adoption and for fostering ongoing dialogue and shared learning among peer institutions.



1. INTRODUCTION

Artificial Intelligence (AI) is no longer a futuristic concept. It is a transformative force reshaping industries, economies, and societies. In the asset management industry, AI is increasingly being applied across operations and investment processes, enhancing efficiency and strengthening decision-making. At the same time, its adaptation raises complex questions related to governance, risk management, ethics and regulatory compliance.

For public investors, these developments present both significant opportunities and important challenges. While AI has the potential to further improve productivity, strengthen analytical capabilities, and support investment objectives, it also introduces new risks that must be carefully managed within a public mandate. These include challenges and issues related to accountability, transparency, data governance and institutional capacity.

This primer provides public investors with an overview of how AI is being applied in the asset management industry by presenting key industry trends and exploring various use cases of AI and their implications for investment strategies and operations. Furthermore, the primer sets out a practical framework for its responsible adoption and outlines the foundational elements required to govern, implement, and scale AI in a responsible and sustainable manner.

To provide a foundation for the discussion that follows, this section first clarifies what is meant by artificial intelligence in the context of this primer and then explains the factors driving its recent surge in interest.

What is artificial intelligence?


AI refers to the science and engineering of creating machines that can perform tasks typically requiring human intelligence—such as learning, reasoning, and problem solving. As defined by John McCarthy, one of the field's pioneers:

“Artificial Intelligence is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.” (McCarthy 2007)

Traditionally AI is classified by technologies such as machine learning (ML), which enables systems to learn from data; deep learning (DL), which uses artificial neural networks with multiple layers; natural language processing (NLP), which allows machines to understand and generate human language; and large language models (LLMs), which are trained on massive datasets to perform tasks like summarization, translation, and question &

answering. (See Figure 1 and Box 1 for further details.) AI is often categorized according to its purposes; Predictive AI leverages historical data to forecast outcomes and support data-driven decision-making. Generative AI (Gen AI) refers to content creation, such as text, images, and audio, while Agentic AI denotes autonomous AI systems that can perceive their environment, make decisions, and take actions to achieve user-defined goals (GoSearch 2025).

FIGURE 1. AI Landscape by Technologies and Purposes



Source: World Bank Group

BOX 1: AI TECHNOLOGIES

AI and machine learning (ML) are used synonymously, but artificial intelligence has broader implications than does machine learning. AI encompasses several subfields:

- **Machine Learning (ML).** ML refers to algorithms that enable systems to learn patterns from data and apply these patterns to make predictions on new information without explicit programming. ML was introduced over half a century ago. However, it experienced a period of stagnation, often referred to as an AI winter, due to disappointing performance. It regained momentum in mid-1990s as computing power surged and the volume of available training data increased.
- **Deep Learning (DL).** DL computational models inspired by the structure and function of the element of the human brain known as neural networks. These models consist of interconnected layers of nodes that process information sequentially, enabling the extraction of increasingly abstract

features from raw input data. DL has transformed the field of artificial intelligence by allowing systems to learn complex patterns from large-scale datasets. The term “deep” refers to the substantial depth of these neural network architectures.

- **Natural Language Processing (NLP).** These techniques enable machines to understand, interpret, and generate human language. NLP integrates linguistics, statistics, and machine learning to analyze and model text data. These methods support a wide range of applications, including information extraction and sentiment analysis. Advances in deep learning have significantly improved NLP performance by allowing models to capture contextual and semantic relationships within language.
- **Large Language Models (LLMs).** Introduced in 2017, the Transformer architecture marked a pivotal milestone in AI by enabling attention-based, parallel computation at scale. This breakthrough laid the foundation for LLMs, models trained on massive texts to learn statistical patterns of language and generate context-aware responses. Over the past decade, transformer-based LLMs have grown rapidly in capacity and accuracy, driving major advances in generative AI, including summarization, translation, question answering, and content generation.

Why the surge in interest and implications for public asset managers

The recent surge in AI interest is driven by the emergence of Gen AI. Gen AI tools are accessible to nontechnical users through natural language, versatile across domains from document drafting to investment analysis, and capable of generating human-like content and insights. As of August 2025, ChatGPT had reached 700 million weekly active users, representing roughly 9 percent of the global population (CNBC 2025). This new wave of AI has sparked a global race for adoption.¹ A survey reveals that 74 percent of US CEOs believe they could lose their jobs within two years if they fail to deliver measurable business results from AI (Dataiku 2025). AI is now seen not just as a tool for innovation, but as a strategic imperative.

Extensive and fast AI adoption is now reshaping work across the financial sector and presents various considerations for public investors. First, staff roles are shifting from “doing” to “interpreting and validating” outputs from AI models (Yeyati 2025),

¹ The AI market, defined as revenue generated from AI-related products and services, reached nearly US \$189 billion in 2023 and is expected to grow to approximately US \$4.8 trillion by 2033 (UNCTAD 2025). Meanwhile, private investment in AI amounted to US \$151 billion in 2024, representing a seventeen-fold increase compared to 2015 (Stanford University 2025).

and upskilling focuses on AI model fluency alongside finance expertise. For example, portfolio managers make investment decisions using AI models, but they must understand how the model works and when to override the model output. Second, public institutions also must consider the AI's broader impacts, such as energy consumption and ethical use of AI. Data center electricity consumption is expected to more than double to 945 terawatt-hours by 2030, surpassing Japan's current total electricity use (IEA 2025). For public institutions, it is essential to consider energy-efficient architecture, green data centers, and sustainability metrics in procurement and deployment decisions. AI is not neutral. Global standards emphasize transparency, fairness, and human oversight.² Public investors should align AI adoption with these principles to ensure responsible and equitable outcomes.

To meet these demands along their AI journey, public investors at all levels need a foundational understanding of AI technologies and practical strategic frameworks for AI implementation. This primer serves as an introductory guide to support that need.

Structure of the primer

Following this introduction, this document is organized into three chapters.

Chapter 2 reviews trends in AI adaptation within the asset management industry. It draws on real-world examples gathered through the joint market outreach efforts of the authors from the World Bank Group Treasury and Bank of Korea Reserve Management Group, including surveys and interviews with global asset managers.³ Chapter 2 also introduces research papers that have been carefully selected from an extensive body of literature, prioritizing those published in academic journals or widely cited in scholarly work.

Chapter 3 presents a framework for responsible AI adoption tailored to public institutions. Drawing on institutional experience, empirical research, and emerging international guidance, the chapter identifies key building blocks for effective AI adoption. These components help organizations assess their readiness, identify capability gaps, and understand the institutional conditions needed to undertake and scale AI initiatives.

2 Recommendation on the Ethics of Artificial Intelligence (UNESCO 2022); EU Artificial Intelligence Act (EU 2024)

3 From August 2024 to June 2025, the World Bank and Bank of Korea conducted a survey of 15 asset management companies to examine their AI initiatives, governance frameworks, use cases, and implementation challenges. In addition, senior managers and practitioners from 12 firms were interviewed, both virtually and in person in New York, to observe demonstrations of AI applications and discuss their institutions' experiences and progress in adopting AI.

Recognizing that institutions differ in their levels of AI maturity and operating environments, the chapter does not prescribe a fixed sequence for implementation but offers instead practical insights and emerging good practices that institutions can use to strengthen the foundations for responsible AI adoption.

Chapter 4 concludes this Primer with key takeaways and future considerations.

2. AI TRENDS IN ASSET MANAGEMENT INDUSTRY

2.1. AI Adoption in Asset Management

The wave of AI adoption has not spared the asset management industry. Over the past decades, many firms, particularly those specializing in quantitative investments, have been actively applying ML to capture various investment signals and market insights. Additionally, some firms have attempted to automate manually intensive operational tasks, mostly relying on rule-based AI systems, such as Robotic Process Automation (RPA)⁴. These applications primarily used task-specific models at the individual business-unit level. However, with the introduction of LLMs and Gen AI, asset management firms have begun to consider AI applications across all investment processes at an enterprise level in an integrated fashion.


The survey indicated that many leading asset management firms have established AI governance frameworks within the past two to three years and are pursuing AI initiatives across their organizations. A high-level executive in a global asset management firm mentioned that investing in AI is now essential—not for generating alpha, but for staying competitive and surviving in the industry. This sense of urgency is also shared by public asset managers, as more than 90 percent of sovereign investors expect AI to become an essential tool in their investment process (Invesco 2024).

Despite its transformative potential, only a small fraction of asset managers and institutional investors were able to proceed beyond experimental stages of adopting

4 RPA is a software that performs routine, repetitive tasks by following the same steps a human would take on a computer.

emerging AI technology. According to an extensive survey of 300 asset managers across the US and Europe, half of the respondents reported that Gen AI adoption for day-to-day operations is either a low priority or not a priority at all (Index Industry Association 2024). Only 16 percent of global asset managers have a fully defined Gen AI strategy and are preparing to implement it across their businesses (BCG 2024). In the public investor space, only 12 percent of central banks are using AI tools in their reserves management (World Bank Group 2025). (See Figure 2.) These survey results highlight that while institutional investors acknowledge the significant potential of AI to transform asset management operations, they remain uncertain about how to start their AI journey.⁵

FIGURE 2. Central Banks' Use of AI tools within the Reserves Management Framework (N=134)



Source: The World Bank Reserve Management Survey 2025

The primary focus of current AI applications in the industry is productivity gain, driven by the innovative performance and relatively easy accessibility of GenAI applications through partnerships with third-party vendors. Many firms have started to see substantial productivity improvements by leveraging vast amounts of internal and external data. MIT research showed that using ChatGPT enabled a 37 percent faster completion of knowledge work with comparable quality (Noy and Zhang 2023). Another report projects that GenAI will empower portfolio managers to augment their investment research and risk analysis capabilities, potentially improving by up to 30 percent (Morgan Stanley and Oliver Wyman 2023).⁶ To allocate limited resources

5 "Many institutional investors, be it pensions, insurers, or sovereign wealth funds around the world, are struggling with how, when, and where to begin their technology transformations." (Baig, Sohoni and Lhuer 2024)

6 For the insurance asset managers, the BCG estimates that AI models enable to improve productivity across the value chain and achieve cost reduction by 5 percent to 15 percent (BCG 2024)

effectively, firms identify key pain points in the investment process and prioritize projects based on their feasibility and impact. Although metrics such as ROI are being used to gauge operational efficiency, many firms acknowledge that it is still too early to apply these metrics rigorously.

AI applications in investment strategies have been explored; however, they face greater challenges compared to productivity gains. It is well known that the inherent characteristics of financial markets—such as small data, low signal-to-noise ratios, and regime change—challenge the application of AI, including the prediction of asset returns (Israel, Kelly and Moskowitz 2020). A survey reveals that AI adoption for alpha generation is viewed as a distant prospect (Krause et al. 2024). Nevertheless, over the past decades, asset managers have leveraged advanced analytical techniques and alternative data (see Box 2 for further details) to uncover new market insights and investment signals, achieving information edges over competitors.

Another notable observation from the interviews is that the firms unanimously emphasized the need to keep humans in the loop when applying AI. This is primarily driven by concerns about regulatory uncertainties and the inherent randomness of LLM models. Fully automated investment decision-making models or use cases are not anticipated in the near future. A survey indicates that less than 10 percent of asset managers currently use automated models based on more stable ML models (Mercer 2024). A Swedish fintech company's failed attempt to replace 700 employees with AI tools also highlights that AI systems still need human oversight and judgment.⁷

BOX 2: HARNESSING UNSTRUCTURED DATA FOR AN INFORMATION EDGE


With advances in technologies for handling unstructured data, investors have increasingly sought to unlock these new information sources, as unstructured data accounts for approximately 80 to 90 percent of today's data (Harbert 2021). Unstructured data encompasses diverse formats, including text, images, audio, and video, unlike standardized tabular data organized in rows and columns. It is often contrasted with traditional data sources and is therefore frequently referred to as alternative data.

7 The Swedish fintech that separated from about 700 workers and replaced them with AI to cut costs in 2022 later reversed this decision and began rehiring human staff in 2025 due to poor customer service. Customers complained that automated responses lacked human interaction and were overly generic and unhelpful. (The Economic Times 2025).

In a broader sense, alternative data also includes structured data from nontraditional sources, such as logs, sensor data, and credit card transactions.

The value of unstructured data lies in its real-time and granular nature, enabling investors to monitor market developments in a timely manner and to capture multiple dimensions of an issue, thereby complementing traditional data analysis. Common datasets include news articles, social media content, earnings call transcripts, satellite imagery, job postings, corporate filings, and patent applications. In practice, most institutions use alternative data as one input among many, rather than in isolation.

Structured and Unstructured Data Today



Source: World Bank Group


Once used primarily by hedge funds, these data sources are now increasingly adopted by a broader range of institutional investors. According to a 2025 survey, 75 percent of buy-side institutional investors use alternative data in research and investment processes (Easthope 2025). Another survey that focused on equity investors reported that all 130 respondents indicated they use alternative data (Exabel 2025). Institutions access new data sources in various ways. Some rely on proprietary internal data or purchase raw external data for in-house analysis. Others, particularly those with limited internal capacity, acquire preprocessed datasets from vendors. The global alternative data market was valued at US \$11.7 billion in 2024 and is projected to grow at an estimated CAGR of 63 percent between 2025 and 2030 (Grand View Research 2025).

Many expect data to become a key differentiator as technologies are increasingly democratized. Recent advances in data-processing technologies and AI algorithms have significantly lowered the barriers to analyzing unstructured data. As a result, competitive advantage is likely to depend less on analytical techniques themselves and more on which data investors own or can access and how effectively they can translate those data into investment value.

2.2. Areas of AI Applications

AI applications in asset management can be broadly categorized into two areas: productivity gain and investment strategies. (See Figure 3.) Productivity gain encompasses enhancing the efficiency of the current manual-intensive workflow, while investment strategies involve generating new insights or investment signals.

FIGURE 3. Popular Use Cases for AI in Asset Management



Sources: World Bank Group-Bank of Korea joint market outreach; academic studies; World Bank Group compilations

The following use cases are neither mutually exclusive nor collectively exhaustive, as many AI approaches are often used concurrently. For instance, new market factors derived from sentiment scores and clustering information can serve as part of the input data for return forecasting or portfolio construction.

2.2.1. Productivity Gain

Information Extraction

Leveraging AI to extract relevant information from lengthy legal documents and financial reports can reduce manual effort and enhance productivity. While rule-based or traditional ML-based Named Entity Recognition (NER)⁸ technology has been widely applied in the past, the emergence of LLMs has significantly improved the accuracy and efficiency of information extraction from complex texts in various formats.

⁸ NER is a technique that automatically scans text to identify and categorize key information into predefined labels

Many institutions are exploring a wide range of use cases, including extracting key information from investment management agreements, loan agreements, financial statements, and corporate reports. To do this, retrieval-augmented generation (RAG)⁹ is used to constrain extraction to relevant target documents, mitigating any hallucinations from LLMs. For example, the BIS's project Gaia team extracted 20 climate-related indicators, such as net-zero commitment, green bond issuance, emissions, and so on, from over 2,000 corporate reports covering 187 financial institutions, storing them in the database for climate risk analysis (BIS Innovation Hub 2024). Similarly, the World Bank Group Pension Department has leveraged retrieval-augmented LLMs to automate the initial data collection process of key data from 10,000+ performance and audit reports submitted by external asset managers per year. Reconciliation is also a popular area for AI-based information extraction. One financial services company reduced manual effort by 20 to 30 percent through an exception-based AI reconciliation application, which automatically resolved most matching errors. Another use case in the World Bank Group Treasury is ASTRA (AI for Securities Terms Reconciliation and Analysis), which streamlined the manual reconciliation process. (See Box 3 for more details.)

BOX 3: AI FOR SECURITIES TERMS RECONCILIATION & ANALYSIS (ASTRA) SOLUTION – WORLD BANK GROUP TREASURY

[Pain point] Accurate cash projections are crucial for asset management operations, helping to minimize opportunity costs from underinvestment and losses due to overdraft. The integrity of these projections depends on the precise representation of securities data, such as payment dates, day-count conventions, and business-day conventions. As vendor-sourced data can contain errors, the World Bank Group Treasury extracts key information directly from the security Final Terms documents and reconciles it against the vendor data. However, manually reviewing hundreds of documents and identifying key data elements requires substantial staff time.

[Goal] ASTRA aims to strengthen cash forecasting by streamlining the manual, time-consuming, and error-prone reconciliation process. An AI-driven tool automatically extracts key data from security Final Terms and identifies discrepancies between the documents and the vendor data.

[Solution] Leveraging the Bond Data Taxonomy created by the International Capital Markets Association (ICMA), the team developed an extractive AI model using Document Intelligence from Microsoft's Azure Cognitive Studio.

9 RAG is an AI framework that combines information retrieval with text generation. It first retrieves relevant documents from an external knowledge source. The retrieved content is then provided to a LLM as context, guiding and constraining the final output.

This innovative solution reads and extracts key terms that affect interest calculations and payment schedules and then reconciles the outputs with the vendor data to ensure accuracy. The model has been trained on more than 150 Final Terms documents from different issuers. ASTRA identifies terms with 90 percent overall accuracy. The time required for reconciliation has been reduced from an hour to less than one minute per document.

Source: ASTRA. <https://treasury.worldbank.org/en/about/unit/treasury/impact/digitalization>

It is important to recognize, however, that achieving 100 percent accuracy with LLMs remains challenging due to their inherently probabilistic nature, which necessitates human supervision for final outputs. One practitioner noted that they deliberately avoid pushing information extraction accuracy to its theoretical maximum, as this encourages staff to remain alert to potential errors. For tasks requiring fully automated and highly precise outputs, rule-based NER or more robust machine learning models may still be the preferred choice.

Generating Documents

Since the emergence of Gen AI, asset management firms, like many other companies, have been actively exploring using AI for generating documents. It is widely applied in drafting emails, meeting notes, research reports, and marketing materials, as well as serving as a coding copilot. One of the most frequently mentioned use cases among asset managers is using Gen AI to draft responses for Requests for Proposals (RFPs) and Due Diligence Questionnaires (DDQs). These documents can contain thousands of questions, requiring weeks of effort to complete. By leveraging existing knowledge databases related to asset management, firms can efficiently generate appropriate responses, significantly reducing the time required. Previously, legacy RFP software relied on keyword matching, but with the introduction of LLMs, firms can now search across diverse resources and generate responses more effectively. Many companies are developing internal chatbot interfaces to facilitate these services, while more targeted applications are integrated into Excel and PowerPoint as add-ins to assist users. Additionally, Gen AI is being used to produce drafts that align with a company's unique writing style, as well as to support HR-related tasks, such as drafting documents that incorporate performance feedback requirements.

Workflow Enhancement

GenAI has enabled a range of workflow enhancements, often focused on quick wins. Common applications include document summarization, translation, coding assistance, and Q&A chatbots, typically deployed through vendor-provided solutions. These use cases are broadly consistent across industries and have delivered measurable efficiency gains. An asset management firm serving global clients reported significant productivity gains from AI translation tools. The Board of Governors of the Federal Reserve System has also published a list of internally deployed AI use cases, including virtual meeting summarization, coding support, and scheduling supported by various AI tools.¹⁰ Box 4 describes the next phase of workflow enhancement using agentic AI and explainable AI.

BOX 4: AGENTIC AI AND BEYOND

While Gen AI brings efficiency in single tasks involving content generation based on specific prompts, a more advanced form of AI called agentic AI is reshaping how complex work gets done, aiming beyond single tasks toward end-to-end process automation. Agentic AI is an AI system capable of autonomous decision-making and of executing complex sequential tasks that require ongoing judgment, coordination, and adaptation.

Some of the key features of agentic AI include (1) memory to retain context across steps, (2) capacity to orchestrate using different tools to fetch data, run code, or call on Application Programming Interfaces (APIs), (3) plan and structure tasks and execute in sequence, and (4) self-correct to adapt as new information arrives. It's important not to mix agentic AI and AI agents. Agentic AI is a framework that decomposes objectives into subtasks, invokes tools and data sources, and iterates on feedback to complete end-to-end multistep workflows. AI agents are just building blocks within the framework, often specialized in specific subtasks.

The landscape of agentic AI adoption is currently led by technology and software companies, followed by financial industries. Adapting agentic AI poses a significant challenge for institutions as it requires (1) a robust and dedicated governance framework due to a substantially different risk profile compared to GenAI, (2) an evaluation mechanism to constrain agents with guardrails and ensure the presence of fallback strategies, (3) stronger data provenance, and

10 <https://www.federalreserve.gov/AI-use-case-inventory.htm>

(4) integration with legacy systems and tool ecosystems to support multi-agent collaboration. Many companies are in the pilot stage trying to address these challenges as they assess their existing business processes and revamp them during the transition from static automation to adaptive operations that enable multi-agent collaboration. By early 2026, several major financial institutions will have moved beyond pilots into limited production deployments while evaluating different multi-agent orchestration frameworks (e.g., LangGraph, AutoGen, CrewAI) in the areas of compliance monitoring, trade reconciliation workflows, and client onboarding. Agentic AI has huge potential in different industries and framework. The rising importance of governance frameworks around AI has fueled a surge in related research. Recent publications in ArXiv's repository (<https://arxiv.org/list/cs.AI/recent>) in the artificial intelligence category show the emerging direction of and forthcoming developments in AI for finance.

In line with growing adoption of agentic AI, a heightened focus is being placed on studying domain-specialized LLMs and multi-agent coordination for complex finance workflows. Various types of AI agents specialized for financial tasks related to risk, compliance, and client service are being developed with the capacity to coordinate using shared memory and to operate within specific policy frameworks.

Another focus area of AI research is explainable AI. High-stakes decisions require explainability, auditability, and defensible governance. Under agentic AI, which involves autonomous decision-making, implementing a mechanism that provides the Agent with a rationale for its decisions indicates when and where it is reliable and supports identifying areas where it needs human oversight. Explainable AI also addresses the “black-box” problem of complex AI models, meeting regulatory demands for transparency

2.2.2. Investment Strategies

Market Intelligence: Sentiment Analysis

Sentiment analysis, a key application in finance, has been extensively studied for decades (Kong et al. 2024). By measuring sentiments expressed in textual data sources, including news articles, social media posts, corporate disclosures, and market research reports, asset managers can gain valuable market insights and identify new market factors. Over the past decades, this approach has evolved from simply counting the frequency of positive/negative words in the texts to incorporating contextual

understandings with pre-trained LLMs (Hansen and Kazinnik 2023). In macro research, the most popular use cases involve developing economic and market sentiment indices based on news articles.¹¹ Quantifying hawkish/dovish monetary policy stances based on central bank communication is another popular area of AI applications, as it is considered one of the most important macro factors for investments. Many firms regularly produce central banks' sentiment indices/scores.¹² For equity research, sentiment extraction from earnings call transcripts is widely applied to detect trading signals on equities.¹³ One research study reveals that firms with a less optimistic tone during conference calls experience higher stock price crash risk in the following year (Fu, Wu and Zhang 2021). Some firms have even taken further steps and analyzed the tonal sentiment of management during earnings calls by leveraging speech emotion recognition (SER) technology. Combining information on how something is said with what is said can enhance prediction accuracy. Sustainable investing also actively employs sentiment analysis methodologies, given that the majority of information is text-based. For example, MALENA developed by the IFC, predicts ESG sentiments from corporate reports and news articles. (For more details, see Box 5.)

BOX 5: IFC'S AI-POWERED SUSTAINABLE INVESTING ANALYSIS PLATFORM, MALENA

To accelerate and strengthen analytics for sustainable investment assessments in emerging markets, the World Bank Group's IFC has developed MALENA, an ESG domain-specific AI solution capable of extracting meaningful insights from unstructured text data at scale. MALENA is now publicly available for investors to upload and analyze documents of their own interest, such as annual reports, impact reports, et.¹⁴

ESG Risk Term Sentiment Analysis

- MALENA identifies 1,200 ESG risk terms within documents and predicts

11 For example, see the Daily News Sentiment Index, developed by the FRB of San Francisco, which measures the economic sentiments of the US. <https://www.frbsf.org/research-and-insights/data-and-indicators/daily-news-sentiment-index/>

12 For example, see "AI wants to know what dove is" <https://www.ft.com/content/4a6d60bd-62a2-4c98-a0c9-5d58814ce0bf> and "Sensing the Fed's Direction with the Help of AI", <https://www.morganstanley.com/articles/mnlpfeds-sentiment-index-federal-reserve>

13 For example, "Overall sentiment on REIT earnings calls improves" <https://www.spglobal.com/marketintelligence/en/news-insights/latest-news-headlines/overall-sentiment-on-reit-earnings-calls-improves-86620294> and "Earnings calls sentiments show a remarkable turn for the better" <https://www.bloomberg.com/professional/insights/trading/earnings-call-sentiments-show-a-remarkable-turn-for-the-better/>

14 <https://malena.ifc.org/>

sentiments using a pre-trained RoBERTa model, fine-tuned on 150,000 sentiment tags labeled by subject matter experts.

- Specifically, it leveraged a Term-Based Sentiment Analysis (TBSA) to learn ESG context based on a sustainability taxonomy for ESG aspect-oriented sentiment analysis. This ESG taxonomy is defined based on IFC's eight environment and social Performance Standards¹⁵ and six Corporate Governance Methodology¹⁶ parameters. (SandwidiBlaise, MukkolakalSuneer 2022)

Questionnaire Hub

- MALENA provides a due diligence assistant feature that enables users to select pre-defined questionnaire templates and extract answers directly from uploaded documents. This feature leverages LLMs, prompt engineering, and RAG technologies to identify relevant information and generate responses with accompanying source references for investors, saving time and increasing accuracy of document reviews.

ESG Risk Term Sentiment Analysis

Risk Terms			
Name	Pos(+)	Neg(-)	%Neg
Suppliers	19	0	
Employment	14	1	7%
Security	9	2	18%
Global Warming	8	6	43%
Renewable Energy	8	0	
⚠ Gender-based Violence	7	0	
Social Risk	7	0	
⚠ Vulnerable Workers	7	2	22%

Questionnaire Hub

Section 1: Disclosure and Transparency		
Question	MALENA Response	Source
1.1 Are any of the internationally accepted sustainability reporting standards like Carbon Disclosure Project, Climate Disclosure Standards Board, Global Reporting initiative, international.....	YES. The company's sustainability report has been prepared with reference to the Global Reporting initiative Universal Standards 2021, and the company has also referred to the CDP Climate Change program for its climate-related financial disclosures. This indicates that the company has utilized internationally accepted sustainability reporting standards to develop its sustainability report.	Sustainability Report_2022.pdf -p.3

Source: IFC

Market Intelligence: Clustering


Clustering is the process of grouping data points based on their features, where points within the same cluster are the most similar and distinctly different from those in other clusters. A popular case of this technique is company clustering, which is used to identify emerging market trends and themes or to facilitate peer group comparison.

¹⁵ <https://www.ifc.org/en/insights-reports/2012/ifc-performance-standards>

¹⁶ <https://www.ifc.org/en/what-we-do/sector-expertise/corporate-governance/cg-methodology-tools>

(See Figure 4) A study showed that company clustering using business descriptions from SEC 10-K¹⁷ filings enables the identification of closer peers, demonstrating higher return correlations within clusters compared to industry groupings defined by the Global Industry Classification Standard (GICS) (Vamvourellis et al. 2023). Another common application involves clustering news articles and extracting emerging-market themes using topic modeling techniques. Some asset managers process thousands of new articles daily to capture investment themes, which are then combined with additional information, such as sentiment scores and relevant stock lists, to offer more comprehensive market insights.¹⁸ One study applied this approach to credit risk management, demonstrating that news topics provide insights into the operational behavior of Chinese corporate bond issuers and support the prediction of company defaults within a three-month horizon (Tang et al. 2024).

FIGURE 4. Company Clustering: Finding related companies



Source: Robeco. 2024. The Current State of AI in Asset Management.

Note: Each small dot represents a company, color-coded based on the GICS (Global Industry Classification Standard). Dots that are close together indicate groups of companies with similar business descriptions.

17 SEC 10-K is an annual report that publicly traded companies are required to file with the US Securities and Exchange Commission (SEC); it provides comprehensive information on the company's business, risk factors, financial and operating results.

18 For example, see BlackRock, Thematic Investing, <https://www.blackrock.com/us/individual/insights/thematic-investing>.

Forecasting Market Factors

Forecasting stock markets has become the most popular AI application in academia and industry.¹⁹ A range of AI techniques, including regularized linear regression, Support Vector Machine, random forest, DL, and, more recently, LLMs, are employed for predictions. These methods leverage not only traditional time series data but also unstructured data such as news and social media (Mailagaha Kumbure et al. 2022). For example, an asset management company has developed a systematic investment platform that leverages a variety of machine learning models, based on more than 600 stock selection signals, to generate alpha forecasts (BlackRock 2024). Forecasting horizons are generally concentrated on daily or monthly predictions.²⁰ This short prediction horizon makes AI-based forecasting more suitable for the tactical asset allocation (TAA) process rather than for the strategic asset allocation (SAA) process, particularly for long-term investors such as pension funds and sovereign wealth funds.

Beyond stock price predictions, AI methodologies are also employed to forecast a range of other market factors, including volatility, bond returns, stock crash risk, and market regime switches. For example, a quant portfolio strategy team applied the Transformer model to forecast multiperiod one-month rolling volatility and constructed portfolios, showing outperformance against portfolios using historical estimates (LEZMI and XU 2023). One study finds that nonlinear machine learning techniques, such as extreme trees and neural networks, outperform traditional approaches in forecasting bond excess returns. It also highlights that incorporating macroeconomic variables improves predictive accuracy (Bianchi, Büchner, and Tamoni 2021). Regarding the market crash, an equity quant strategy team explored ML techniques to predict stocks with a high distress probability and showed the information could help outperform the market return by excluding them in a portfolio (Swinkels and Hoogteijling 2022). One study used random forest algorithms to estimate the probability of certain market regimes and construct regime-switching risk parity portfolios, which outperformed nominal risk parity portfolios (Uysal and Mulvey 2021).

The most common challenges in AI-driven prediction models are overfitting and a lack of explainability, issues that practitioners use various methodologies to mitigate. To enhance the robustness of prediction models, they often adopt an ensemble approach, combining prediction results from multiple submodels. Basic ensemble techniques include averaging, weighted averaging, and max voting, while advanced methods include stacking, bagging, and boosting (Nti, Adekoya, and Weyori 2020). One research study

19 A literature review of AI applications in finance reveals that predictive/forecasting systems for stock prices, performance, and volatility are the predominant research topics, followed by classification/detection/early warning systems (Bahoo, et al. 2024).

20 For example, academic research on DL applications for stock market prediction from 2017 to 2019 predominantly focus on daily predictions (Jiang 2021).

highlighted ensembles as the only free lunch in ML, similar to diversification impact in investing (Swinkels and Hoogteijling 2022). To address the “black box” issue of AI models, explainable AI techniques, such as SHapley Additive exPlanations (SHAP), are often adopted (Swinkels and Hoogteijling 2022). The SHAP value shows the feature’s marginal contribution to a model’s prediction. One study uses LLMs’ reasoning capabilities to forecast weekly and monthly stock returns with explanations (Yu et al. 2023).


Scenario Analysis with Synthetic Data

Gen AI is used not only for document and image creation but also for producing synthetic financial data, which is applied in market scenario analysis. DL models, such as GANs (Generative Adversarial Networks), VAEs (Variational Autoencoders), and diffusion models, learn the structure of nonlinear, sparse, high-dimensional financial data and generate synthetic datasets that retain key real-world patterns. One study demonstrated that VAEs generated stress scenarios for three-month TONA²¹ futures, capturing tail events more effectively than linear models such as PCA. These scenarios were then used to estimate portfolio Value at Risk (VaR), improving risk coverage (Aoki, Shah, and Viklund 2025). Boier (2024) presented a flexible approach to scenario generation by combining an LLM with generative models. The process involves the LLM converting a user’s query, such as “Generate yield curves similar to those in the year post-COVID,” into a historical period specified in JSON format. This input is then used by the generative models to produce synthetic data that is similar to that period.

When using synthetic data, robust and ongoing model validation is particularly important, as it may introduce risks related to generalizability, bias, and degraded real-world performance (Financial Conduct Authority 2025). The Synthetic Data Expert Group under the UK Financial Conduct Authority emphasizes that statistical similarity between real and synthetic data is insufficient to confirm a model’s fitness and introduces validation approaches such as Train-Synthetic-Test-Real (TSTR), in which models are trained on synthetic data and evaluated on real data. (See Figure 5.)

21 Tokyo Overnight Average rate (TONA) is JPY interest rate for uncollateralized overnight transactions in the interbank call money market.

FIGURE 5. Variational Auto Encoder (VAE) for Synthetic Data Generation



- **Training:** The VAE is trained to minimize the difference between X and X' , ensuring the encoder maps the data to a meaningful latent vector space ($X \rightarrow Z$) and the decoder accurately regenerates the data from those vectors ($Z \rightarrow X'$).
- **Synthetic Data Generation:** New data is produced by sampling random latent vectors from the standardized distribution of the latent space and feeding them to the trained decoder.

Source: World Bank Group

Portfolio Construction

AI technologies enhance portfolio construction mainly in two ways. First, as previously discussed, many researchers and practitioners leverage AI to improve the accuracy of return and risk forecasts. These are then used as critical inputs for traditional portfolio construction methodologies, such as mean-variance optimization (MVO) or risk parity. Second, AI can be used to develop alternative asset allocation methodologies, although empirical evidence is limited (Bartram, Branke, and Motahari 2020). Nonetheless, one notable approach is hierarchical risk parity (HRP), which does not require return estimates. HRP captures the hierarchical structure of the investment universe using clustering techniques. This approach aims to create a diversified portfolio with lower risk than traditional risk-parity methods (Prado 2016).

External Manager Selection

AI can also add value to the manager selection process as an alternative or complementary quantitative evaluation tool. Various ML methodologies are used to identify mutual funds or hedge funds with positive future returns, supporting fund selection (Fontanille et al. 2020). One study showed that ML models, incorporating fund characteristics data such as fund flow, total net assets, manager tenure, turnover ratio, etc., can help to predict ex-ante mutual funds return and generate a significant alpha net of costs (DeMiguel et al. 2023). The Japanese pension fund GPIF commissioned an

AI research institute to explore ways to use AI in selecting and monitoring funds. To avoid arbitrariness and subjectiveness in fund selection and to alleviate the workload of a small number of staff, they developed a DL-based model to assess the similarity between Japanese equity funds and monitor fund investment behaviors in a more systematic fashion (Tajiri et al. 2020).

Trading Execution

AI can enhance the development of optimal trading execution strategies, particularly in stocks and currencies, where electronic and algorithmic trading are prevalent. For example, traders can use AI to select the optimal trading execution algorithms²², where AI predicts and compares the transaction costs of each trading strategy based on order attributes (e.g., ticker, side, sector, size, and market cap), stock clustering, and real-time market conditions (e.g., relative volume, volatility, and spread) (Cao 2023). A large asset owner whose portfolio ultimately includes a bit of every stock in the market can leverage AI to minimize market impacts by adjusting order sequencing and timing, enhancing overall performance (Turner 2023). Another study showed that a reinforcement learning model derives the optimal trading actions based on limit order book data for nine US stocks, showing superior performance for seven of them compared to the VWAP strategy (Ning, Lin and Jaimungal 2022).

However, more complex and sophisticated models, often used for optimization, can exacerbate challenges related to explainability and robustness. As model complexity increases, practical value may decline if results become difficult to interpret, validate, or replicate. For example, a financial institution developed a reinforcement learning-based trading optimization strategy for currencies in 2019, but it reverted to a simpler but more robust trading strategy in 2023 (Parsons 2023). The institution found it difficult to explain the model's performance to clients and ultimately concluded that similar results could be achieved using less complex models. This experience is consistent with insights from quantitative investors we interviewed, who emphasized that the most advanced model is not always the most effective. Understanding the trade-off between model complexity and performance is therefore critical.

AI use cases in asset management continue to evolve and expand across functions. Having reviewed these developments, we now turn to the question of implementation. Translating promising AI applications into practice in a safe, effective, and sustainable manner requires a coherent institutional framework. The next section examines the conditions necessary for responsible AI adoption in the public sector.

22 The most common execution algorithms include POV (percentage of volume), VWAP (volume-weighted average price), TWAP (time-weighted average price), liquidity seeking, arrival price, and dark strategies. Details can be found at <https://www.cfainstitute.org/insights/professional-learning/refresher-readings/2025/trade-strategy-execution>.

3. RESPONSIBLE AI ADOPTION

Responsible AI adoption at an institutional level means evolving from early experimentation and basic AI awareness to scalable, trusted, and operationally effective AI use. Rather than approaching AI adoption as a technology deployment exercise, this chapter frames it as an institutional capability challenge, shaped by governance maturity, organizational readiness, risk management capacity, and enabling foundations in data, infrastructure, and human capital.

The AI adoption environment and associated implementation challenges may differ across public investors, including central banks, sovereign wealth funds, and public pension funds, reflecting variations in institutional mandates, governance structures, regulatory oversight, and risk tolerance.

The guidance and analysis were gathered predominantly from empirical results and institutional experience documented by findings from the Bank for International Settlements Irving Fisher Committee IFC survey (BIS IFC 2025) and the governance-focused analysis by the BIS Consultative Group on Risk Management (BIS CGRM 2025) on how AI adoption should be embedded within central bank risk and control frameworks. Together, these sources provide a robust empirical and policy-relevant foundation for understanding the opportunities and risks associated with AI in public financial institutions.

3.1. Why Responsible AI Adoption

For public institutions, AI is no longer a niche innovation. It is becoming a core capability, enabling faster analysis and automation and more data-driven decisions across functions. However, the very features that make AI powerful also introduce new and distinct risks when AI is not properly governed and controlled.

AI introduces risks that differ from those associated with traditional technology or digitalization. Models trained in biased or incomplete data can produce misleading or biased results. The “Black Box” systems of AI can make it difficult to explain why a particular decision was reached. Models can drift over time and gradually become unreliable. Cybersecurity risks expose AI models to the possibility of data or model manipulation. In the public sector context, these risks do not merely affect efficiency; they can undermine legal compliance, policy credibility, and public trust.

For Public investors operating under long-term fiduciary mandates, public accountability, and heightened regulatory scrutiny, AI cannot be treated as a competitive technology race. Unlike private firms, public investors must adopt AI sequentially and align with institutional risk appetite and statutory obligations. Responsible AI adoption in this context is inherently phased: strengthening foundations first, piloting cautiously, and scaling selectively.


For these reasons, responsible AI adoption is best understood as an institutional transformation rather than a technology project. In this context, “responsible” AI means ensuring clear governance, accountability, data stewardship, legal and ethical safeguards, and effective risk management embedded across the full AI lifecycle, from use case selection and design through deployment, monitoring, and retirement. Attaining this balance requires integrated governance, strong data and cybersecurity foundations, and sustainable investment in people, processes, and organizational culture, supported by continuous monitoring and transparency as technology and regulatory requirements evolve.

This transformation is a journey that begins with basic awareness and experimentation and progresses toward mature, ethical, and sustainable AI use. Many institutions today are AI-aware but not yet AI-ready. AI awareness refers to the stage at which institutions develop familiarity with AI concepts, experiment with tools, and build foundational literacy among staff. While this stage is valuable for learning and exploration, it does not necessarily indicate that an institution is prepared to operationalize AI in mission-critical environments. AI readiness, by contrast, requires institutional capability: a formally articulated AI strategy, defined decision rights, risk-tiered governance pathways, data governance maturity, infrastructure, skills, risk management capacity, and operational lifecycle oversight. Making this distinction helps institutions set realistic expectations, prioritize capability development, and avoid premature adoption of higher-risk use cases before appropriate safeguards are in place.

This journey is guided by three interrelated building blocks (illustrated in Figure 6):

- **Vision and Leadership**, which establishes purpose, strategic direction, and accountability for AI initiatives.
- **Strategic Foundation**, which translates leadership intent into actionable principles, policies, governance practices, regulatory alignment, and risk management;
- **Enabling Environment**, which ensures the necessary data, technology infrastructure, security, skills and organizational culture are in place to operationalize AI responsibly.

Together these building blocks enable institutions to harness the benefits of AI while upholding trust and institutional responsibility.

FIGURE 6. Framework for Responsible AI Adoption

Source: World Bank Group

3.2. Vision and Leadership

Strategic Direction

Without a clear institutional AI vision, adoption is often fragmented, leading to duplicated efforts, inconsistent risk handling, investment prioritization issues, and integration challenges for AI in decision-making. These problems are intensified given the rapidly evolving technology and rising expectations for AI, including Gen AI.

The BIS CGRM formally recognizes the absence of a clear AI strategy as a material strategic risk. It notes that without strategic clarity, AI applications may evolve in ways that are misaligned with institutional objectives, risk appetite, and accountability frameworks, thereby increasing legal, ethical, operational, and reputational exposure (BIS CGRM 2025).

Survey evidence from the IFC survey supported this fact when it found that many AI initiatives tend to originate at the departmental or team level, often driven by local experimentation rather than by an institution-wide strategy (BIS IFC 2025). While such experimentation can generate valuable learning, it also creates challenges for governance bodies responsible for oversight, prioritization, and risk assessment.

An institutional AI vision therefore functions as a foundational strategic control. It positions artificial intelligence not merely as a technology initiative but as a strategic capability that supports the institution's mission, statutory mandates, and core values.

A well-defined AI vision helps establish clarity across several key dimensions:

- **Purpose and institutional alignment:** Why AI is being adopted and which institutional objectives it supports.
- **Governance boundaries:** Where AI use is permitted, restricted, or prohibited.
- **Human oversight and accountability:** The expected role of human supervision and decision responsibility; and
- **Public-interest alignment:** How AI adoption aligns with legal obligations, ethical standards, and broader public-interest considerations.

In addition to defining purpose and governance boundaries, institutions may also benefit from articulating how AI is expected to create measurable institutional value. This often involves identifying priority domains where AI can deliver the greatest impact.

Leadership Commitment, Decision Rights and Stakeholder Engagement

Leadership engagement is a critical differentiator between institutions that remain in pilot mode and those that achieve controlled scaling of AI use. The IFC survey evidence indicates that institutions with explicit senior-level involvement are more likely to formalize AI governance processes and allocate resources for oversight (BIS IFC 2025).

Strong leadership commitment is therefore essential to translating AI strategy into action. This commitment is expressed through clearly defined decision rights, accountability, and executive sponsorship. A top-down framework signals institutional priority and clarifies who is authorized to approve, escalate, and assume responsibility for AI-related decisions, while cross-functional engagement ensures that AI initiatives remain operationally grounded rather than siloed. Effective stakeholder engagement should span

- Business and policy functions
- Risk management, compliance, and internal audit
- Legal, privacy, and cybersecurity
- Technology, data, and operations

Effective stakeholder engagement requires particularly strong alignment and partnership between business/investment functions and technology/IT. AI priorities must be codefined rather than handed off sequentially. Business and investment teams articulate value objectives, fiduciary considerations, and risk tolerance; technology defines architectural feasibility, data integration requirements, security constraints, and scalability implications. Fragmented ownership between these functions is a common cause of stalled AI initiatives, misaligned investments, and governance breakdowns. Sustainable AI adoption therefore depends on clear, shared accountability between business and technology leadership.

3.3. Strategic Foundation

A strategic foundation is required to translate leadership intent for responsible AI adoption into consistent and enforceable operational practice. In practice, it consists of six interdependent elements: AI principles, institutional AI policies and standards, legal and regulatory compliance, governance, lifecycle oversight, and risk management processes. The objective of this strategic foundation is to provide an institutional mechanism through which ethical values, legal obligations, governance, and risk consideration are embedded across the entire AI lifecycle, from use case selection to system design, deployment, monitoring, and decommissioning.

3.3.1 AI Principles

Institutional AI principles articulate the core values and boundaries for how artificial intelligence is used. It defines an organization's ethical expectations around AI deployment such as fairness, transparency, robustness, security, accountability, and human oversight. AI principles on their own do not impose enforceable requirements, rather they provide the essential value framework that informs policy design, governance choices, and risk tolerance.

In practice, human oversight means that AI-generated outputs are subject to review and validation by designated staff before being used to influence material decisions. For example, an AI system may generate market insights, detect anomalies in portfolio exposures, or flag potential breaches of risk limits. In such cases, human analysts or risk managers would review the outputs, assess their reliability and context, and determine whether any operational or policy action is warranted. This “human-in-the-loop” approach ensures that AI supports analytical processes while accountability and final decision-making authority remain with human decision-makers.

Internationally, the OECD Principles on Artificial Intelligence²³ emphasize that AI systems should be fair, transparent, robust, secure, and accountable, setting the benchmark for trustworthy and human-centric AI standards. Adopted by OECD members and partner countries in 2019, these principles have provided guidance on AI strategies and regulatory approaches.

23 Countries use the OECD AI Principles and related tools to shape policies and create AI risk frameworks, building a foundation for global interoperability between jurisdictions. The European Union, the United States, the United Nations, and other jurisdictions use the OECD's definition of an AI system and lifecycle in their legislative and regulatory frameworks and guidance. The principles, definition, and lifecycle are all part of the OECD Recommendation on Artificial Intelligence (OECD 2024).

3.3.2. AI Institutional policies and standards

AI principles must be translated into clear institutional requirements to shape day-to-day operations, decisions, and system designs. This is illustrated in the institutional AI policies and standards set to ensure enforceable requirements for both the staff and the organization. These requirements define what type of AI may be used, how data can be collected and processed, what documentation and explainability is required, how human oversight is to be applied, and how and what ethical and legal considerations are embedded into system procurement, development, and deployment. These policies and standards ensure that all AI deployments do not go through an “informal” experimental process or a vendor-driven adoption. Instead, it ensures that they are introduced and deployed in a fully auditable, controlled manner consistent with the institution’s obligations.

3.3.3. AI Legal and Regulatory Compliance

Responsible AI use also requires strict adherence to legal and regulatory obligations within public financial institutions; this encompasses the external laws, standards, and supervisory guidelines that govern the design and use of AI systems and enforces accountability.

Regulatory developments are now a central driver of AI strategy. The recent 2024 IIF–EY Annual Survey on AI/ML (IIF and EY 2025) reports that 73 percent of financial institutions now identify regulatory developments as key drivers of their AI strategy, up significantly from 57 percent in 2023. Key concerns cited include limited model explainability, regulatory constraints on advanced techniques, and growing complexity of AI systems.

From a compliance perspective, institutions stay informed and adhere to applicable emerging regulations and frameworks. Most notable is the European Union’s AI Act (EU 2024), enforced in 2024, introducing a risk-based framework with obligations to governance, transparency, documentation, and human oversight and monitoring. In parallel, the NIST AI Risk Management Framework (NIST 2024) and the OECD AI Principles (OECD 2024) provide structured reference points that help institutions interpret regulatory expectations and translate legal obligations into operational requirements.

To ensure compliance, legal and compliance teams should be engaged from the earliest stages of AI initiatives, particularly with response to data sourcing, model transparency, auditability, and documentation. Continuous dialogue with regulators is equally important, as both AI capabilities and supervisory expectations continue to evolve.

3.3.4. AI Governance

Legal and policy requirements are effective only when supported by robust AI governance that clearly defines responsibilities for approval, oversight, and enforcement across the AI lifecycle. AI governance provides the institutional mechanism through which leadership intent, ethical principles, and regulatory obligations are translated into operational practice, while managing key risks including bias, data protection and privacy concerns, limited explainability, model drift, and misuse, regardless of whether systems are developed internally or procured externally. Aligning AI investments, risk controls, and resource allocation with institutional objectives and risk tolerance leads to effective governance that enhances transparency and accountability while enabling innovation within well-defined boundaries.

Although the significance of AI governance is increasingly acknowledged, institutions often encounter practical challenges regarding its organization. Operationalizing governance principles necessitates deliberate decisions about authority and accountability assignment, proportional application of oversight across varying risk levels, and integration of AI governance into current institutional frameworks. These considerations are essential to ensure that governance remains effective as AI adoption transitions from pilot projects to core operational functions.

AI governance must remain dynamic, evolving in response to technological advancements, emerging risks, and shifting regulatory landscapes. It should be adaptive and iterative, anchored by sound regulatory foundations and facilitated by collaboration among business, risk, technology, legal, and policy stakeholders. Viewing AI governance as an ongoing process rather than a one-time initiative is crucial. Employing an adaptive, iterative methodology characterized by continuous learning, early risk identification, and multidisciplinary collaboration consistently yields better outcomes than rigid models.

Reflecting this adaptive approach, the World Economic Forum and the Personal Data Protection Commission (2020) (World Economic Forum 2020) emphasizes the importance of iterative, practice based governance supported by actual use case testing and continuous assessment. This enables institutions to refine AI applications through implementations experience, balancing innovation with safe guards and evolution of more mature governance frameworks.

Box 6 considers the translation of these governance principles into tangible institutional structures, with a focus on governance models, oversight bodies, and the integration of AI governance, rather than establishing parallel frameworks.

BOX 6: AI GOVERNANCE STRUCTURES AND INSTITUTIONAL INTEGRATION

There is no single AI governance structure that fits all institutions. Effective governance must reflect an institution's mandate, clear institutional risk appetite for AI use, regulatory environment, and level of AI maturity. As with other governance arrangements, institutions need clearly defined committees and decision-making bodies operating at strategic, tactical, and operational levels, with explicit roles, responsibilities, and escalation paths. Without this clarity, AI initiatives risk fragmentation, inconsistent oversight, and unclear accountability.

These choices should be addressed early, beginning at the AI use-case intake stage, before development or deployment. Early governance engagement ensures alignment with institutional objectives, clarifies risk expectations, and prevents costly rework or remediation later in the lifecycle.

1. Governance Models: Centralized, Decentralized, and Hybrid

AI governance models determine how decisions are made and risks are controlled by defining where authority and accountability for AI decisions reside.

- **Centralized models** concentrate oversight within a single enterprise-level body (e.g., an AI steering committee or center of excellence). These models promote consistency, uniform risk controls, and clear accountability, and are often well-suited to highly regulated environments or institutions at an early stage of AI adoption.
- **Decentralized models** distribute responsibility across business units, enabling flexibility and faster innovation. However, without strong coordination mechanisms, they can lead to fragmented oversight, uneven risk management, and inconsistent application of principles.
- **Hybrid models**, increasingly adopted in practice, combine centralized principles, standards, and risk classification with decentralized execution. This approach allows governance to scale proportionately with risk, applying stronger controls to higher-risk use cases—while preserving institutional agility and innovation capacity.

2. Oversight Bodies and Mandates Across the AI Lifecycle

Operationalizing AI governance requires clearly defined oversight bodies and mandates that specify which committees or functions review, approve, and monitor AI initiatives across institutional levels.

- **Strategic level:** senior committees or executive sponsors set direction, approve policies, and define risk appetite.

- **Tactical level:** cross-functional groups review proposed AI use cases, conduct initial risk screening, and determine appropriate governance pathways before resources are committed.
- **Operational level:** delivery and control functions implement governance decisions through concrete controls, monitoring, documentation, and reporting.

3. Embedding Governance into Existing Institutional Frameworks

Sustainable AI governance is most effective when embedded within existing institutional frameworks, including enterprise risk management, compliance, data governance, IT and cybersecurity, and internal audit. Embedding AI governance within these structures strengthens accountability, avoids duplication, and allows governance maturity to evolve incrementally alongside institutional capability. Analysis by the BIS consistently shows that integrating AI governance into established frameworks rather than creating parallel governance regimes enhances coherence, supervisory confidence, and long-term sustainability. (BIS 2023)

3.3.5. AI Risk Management

AI risk management operationalizes governance, principles, policies, and regulatory alignment by providing the structured processes and controls needed to identify, assess, mitigate, and monitor risks specific to AI systems, including bias, cybersecurity, vulnerability, model drift, and opaque decision-making.

The NIST AI Risk Management Framework (NIST 2024), offers a structured, lifecycle-based approach to managing these risks across the design, development, deployment, and monitoring phases. The framework emphasizes the importance of institutional capacity building, including governance bodies, impact assessments, and ongoing performance monitoring.

In practice, operationalizing AI risk management often involves structured processes such as intake screening for proposed AI use cases, risk-tiering frameworks, and ongoing monitoring and revalidation mechanisms. Intake screening allows institutions to evaluate proposed AI applications before development or deployment, assessing factors such as use-case criticality, potential impact, data sensitivity, and alignment with institutional mandates. Based on this assessment, AI applications may be classified into risk tiers, with higher-risk systems being subject to more stringent governance requirements, including model validation, explainability reviews, and approval by designated oversight bodies. Once deployed, AI systems typically require continuous

monitoring and periodic revalidation to detect issues such as model drift, data-quality deterioration, or unintended biases that may emerge as operating conditions or input data evolve.

In the public sector, where algorithmic decisions can materially affect individuals, markets or public outcomes, robust risk management is essential to ensure explainability, proportionality, and accountability, including preparedness for unintended consequences.

In addition to an AI risk management framework, it is critical for institutions to have the ability to identify AI-specific threats. Adversarial risks to AI systems go beyond the traditional cybersecurity concerns. They include data poisoning, model manipulation, prompt injection, model inversion, and misuse of generative capabilities. For example, the MITRE Adversarial Threat Landscape for Artificial Intelligence Systems (MITRE ATLAS 2023)²⁴ offers a structured classification of these threats, allowing organizations to convert general risk categories into specific threat scenarios and control measures throughout the AI lifecycle.

By using MITRE ATLAS, together with the NIST AI Risk Management Framework, organizations can systematically model threats, design better controls, and enhance their ability to anticipate, prevent, and address security and misuse risks unique to AI.

The distinct nature of AI-specific risks and their implications for governance and control design are summarized in Box 7.

BOX 7: WHY AI RISK IS NOT THE SAME AS TRADITIONAL CYBER RISK

Traditional cybersecurity frameworks focus on protecting systems, networks, and data from unauthorized access, disruption, or destruction. While these concerns remain relevant for AI systems, artificial intelligence introduces distinct and additional risk vectors that extend beyond conventional cyber threats.

1. AI systems are vulnerable not only at the infrastructure level but also at the data and model level. Attacks such as data poisoning, model inversion, and training-data manipulation can compromise system behavior without breaching network defenses.
2. AI systems may be deliberately misused or repurposed, particularly in the case of generative models, leading to reputational, legal, or societal harm even when systems function as designed.

24 MITRE is a not-for-profit operator of federally funded research and development centers (FFRDCs) for the U.S. government.

3. AI risks evolve over time; model performance can degrade, biases can emerge, and unintended behaviors can arise through feedback loops, changing data distributions, or interaction with users.

These characteristics mean that AI risk cannot be fully addressed through traditional cybersecurity controls alone. Effective AI risk management should be proportionate to use-case criticality and embedded within existing enterprise risk-management and cybersecurity frameworks. It should include clear escalation paths, defined accountability, and regular review cycles to ensure that controls remain effective as AI systems evolve. Treating AI risk management as a continuous process rather than a one-time compliance exercise is critical to sustaining trust and enabling responsible, scalable AI adoption.

Frameworks such as the NIST AI Risk Management Framework (NIST 2023) and the MITRE Adversarial Threat Landscape for Artificial-Intelligence Systems (MITRE ATLAS 2023) provide institutions with tools to address these AI-specific challenges in a structured and operational manner.

AI principles, institutional policies, regulation compliance, governance and risk management should function as a coherent and integrated framework. Principles define the values and boundaries for acceptable AI use, policies translate those values into operational requirements, regulation establishes binding legal obligations, governance assigns decision rights and accountability, and risk management implements these exceptions through controls, oversight, and continuous monitoring across the AI lifecycle. Treating these functions in silos leads to fragmentation, inefficiencies, and institutional exposure. An integrated approach enhances agility, maintains public trust, and enables responsible, scalable, and accountable AI adoption.

3.4. Enabling Environment for AI Adoption

While a strong strategic foundation establishes the principles, policies, governance structures, and risk-management arrangements necessary for responsible AI adoption, it is not sufficient on its own. These frameworks must be supported by an enabling environment that allows institutions to operationalize governance decisions, implement controls, and sustain AI use in practice. Experience across central banks and public financial institutions shows that AI initiatives primarily fail or stall not due to a lack of interest or use-case ideas, but because underlying institutional conditions, such as data readiness, infrastructure, skills, and operational processes, are insufficiently developed

The case studies in Boxes 8 and 9 illustrate this:

- Box 8 summarizes key findings from recent central-bank experience, drawing on survey evidence and operational lessons documented by the IFC survey (BIS IFC 2025). The findings highlight the most frequently cited constraints, risks, and trade-offs encountered by central banks. These findings and diagnoses illustrate why AI adoption remains uneven across institutions and why technical capability alone is insufficient for effective AI introduction and use.
- Box 9 presents the Bank of England’s institutional AI strategy (BoE 2025) as a practical illustration of how a public financial institution can respond to common adoption challenges by embedding AI within a broader data and analytics transformation, strengthening governance and ethical safeguards, and investing in people and capabilities alongside technology.

Together, these case studies reinforce the central message that responsible and scalable AI adoption depends not only on a sound strategic foundation but also on a coordinated enabling environment. This environment rests on a set of mutually reinforcing enablers, discussed in the sections that follow:

- Data and infrastructure foundations
- Privacy and cybersecurity
- Talent and cultural readiness
- Operationalizing AI responsibly

BOX 8: INSTITUTIONAL CONSTRAINTS SHAPING AI ADOPTION IN CENTRAL BANKS


Context and Objective: The Bank for International Settlements, through its Irving Fisher Committee (IFC), conducted a comprehensive survey to document how central banks use AI and ML in practice. The objective was to move beyond conceptual discussions to capture real institutional experience.

What was done: The survey covered nearly 60 central banks and examined AI use cases, governance arrangements, perceived benefits, and implementation challenges across statistical, supervisory, research, and operational functions.


Key Barriers: Across institutions, skills shortages and risk management capacity emerge as the most frequently cited barriers to AI adoption, reflecting constraints not only in technical expertise but also in the ability to oversee, validate, and govern AI systems responsibly. Additional barriers include system integration challenges, talent retention, and legacy technology constraints,

underscoring the difficulty of embedding AI into complex, mission-critical environments. Institutional factors such as limited organizational agility, leadership readiness, and unresolved ethical considerations also persist as structural hurdles.

A. Skills shortage and addressing risks are key barriers



B. Privacy, cyber security and biases are top concerns¹



¹ Normalised scores from 1 to 5 (1 = not sure; 2 = not impactful at all; 3 = slightly impactful; 4 = moderately impactful; 5 = highly impactful).

Source: BIS IFC 2025

Top Concerns: Central banks consistently rate privacy, cybersecurity, and bias as the most impactful risks they face. These concerns are closely followed by explainability challenges, institutional expertise gaps, and dependency on external technology providers, indicating that AI challenges extend well beyond implementation and into ongoing oversight, transparency, and accountability.

BIS experience to date suggests that AI and ML use cases have delivered considerable benefits, including more accurate forecasts, faster analytics, and improved supervisory effectiveness. The trade-offs between predictive accuracy and interpretability remain a defining challenge where highly complex models often function as “black boxes,” limiting transparency, complicating bias detection, and weakening communication with policy makers and the public. Furthermore, the emergence of Gen AI and LLMs introduces specific risks, such as “hallucinations” in which models generate confident but factually incorrect outputs. These characteristics reinforce the need for human oversight and structured governance throughout the AI lifecycle.

BOX 9: BANK OF ENGLAND: INSTITUTIONAL AI STRATEGY AND GOVERNANCE FRAMEWORK

In August 2025, the Bank of England published The Bank's artificial intelligence (AI) strategy as a standalone document outlining how it intends to apply AI internally across the institution. The strategy defines how AI adoption supports the Bank's mission to promote monetary and financial stability while ensuring accountability, transparency and ethical use of AI within a central bank operating environment. The Bank's AI strategy articulates a mission to enable the safe, ethical and effective use of AI across the organization, supported by an institutional vision that all staff, regardless of role, should be empowered to use AI tools and services to excel in their work. It is explicitly positioned as part of the Bank's overall Data and Analytics (D&A) strategy and cloud-based Enterprise Data Platform initiative, laying the foundation for broader analytical innovation.

The strategy outlines

- A set of AI guiding principles (the "TRUSTED" framework) to shape how AI systems are targeted, reliable, understood, secure, tested, ethical and durable.
- Five strategic goals including productivity enhancement, responsible and effective use of AI, collaboration and learning, and broad staff access to AI tools.
- Detailed priorities to operationalize those goals, such as enabling AI access for staff, promoting experimentation, developing governance frameworks, and investing in talent and skills.
- A formal AI governance framework under a Bank-wide AI Governance Committee that delineates responsibilities for both technical and non-technical staff, embeds risk management processes, and emphasizes ethical oversight.

AI principles in the 'TRUSTED' framework

T	Targeted	We focus on AI solutions that have a clearly defined business purpose that ties to our mission and generates benefits we can articulate.
R	Reliable	We focus on reliable AI solutions implemented with high performance standards and grounded to high-quality data.
U	Understood	We understand the behaviour of the AI solutions we implement in the context of their intended purpose, documenting relevant details including limitations that may apply.
S	Secure	We implement secure AI solutions that protect data, systems, and their users from risks such as unauthorised access, manipulation, or misuse.

T	Tested	We implement AI solutions that are thoroughly tested for their technology standards and for their impact on human behaviour and decision-making under relevant scenarios, including stressed conditions.
E	Ethical	We implement AI solutions that adhere to fundamental ethical principles, including being beneficial and scientifically rigorous, fair and inclusive, transparent and secure, and compliant and accountable.
D	Durable	We focus on AI solutions that are durable, remaining effective despite changing conditions and usage patterns.

Key challenges

- Legacy data architecture and integration constraints: Effective use of AI depends on a broader data transformation agenda to improve data accessibility, quality and interoperability across the Bank's functions to support advanced AI capabilities.
- Earning institutional trust in AI-driven insights: The strategy recognizes the need to build confidence among users across roles by emphasizing explainability, principles-based governance, and clear expectations for responsible use.
- Workforce and organizational change: Developing a comprehensive AI skills curriculum, tailoring training according to role requirements, and attracting and retaining AI talent are identified as ongoing challenges that require sustained investment and cultural adaptation.

Key lessons

- Strategy-led adoption is foundational: AI adoption is most effective when embedded in a broader D&A strategy, rather than pursued as isolated pilots or ad-hoc experimentation.
- Ethical and governance frameworks matter: Clear principles (e.g., the TRUSTED framework) and formal governance structures are indispensable to ensuring safe, transparent, and accountable use of AI in public institutions.
- Inclusive access and workforce development drive value: Providing AI tools and capabilities to all staff and embedding AI literacy and skills development across the institution helps democratize innovation and build internal capacity.

Source: BoE. 2025. The Bank's Artificial Intelligence (AI) Strategy. Bank of England, August 21, 2025.
<https://www.bankofengland.co.uk/about/governance-and-funding/staff-codes-and-policies/the-banks-ai-strategy>.

3.4.1. Data and Infrastructure Foundation

Data Governance

The application of AI starts with data. At the core of each model is a requirement for structured, reliable, and well-governed data. Without it, even advanced AI tools can produce biased, unreliable, or unexplainable results. For central banks and public-sector - investors, where credibility, transparency, and accountability are essential, data must be managed as a strategic asset.

Most public investors operate within complex data environments shaped by legacy systems, fragmented architectures, and multiple internal and external data sources. This has led to inconsistent definitions, parallel data views across front, middle, and back offices, and continued reliance on manual processes. These challenges already affect reporting and oversight and are amplified when AI is introduced.

AI-ready data goes beyond traditional data quality initiatives. It requires consistent definitions, integrated data across systems, controlled access, and full traceability through clear lineage and auditability. Strong data governance frameworks establish formal policies for ownership, access, quality controls, retention, and accountability.

Standards such as the DAMA Data Management Body of Knowledge (DAMA 2024) and ISO/IEC 38507 (ISO 2022b) define best practices in metadata management, auditability, and institutional alignment. The NIST AI Risk Management Framework similarly calls for embedding data governance into broader AI oversight to ensure traceability, explainability, and fairness (NIST 2024).

As a practical first step, institutions may benefit from conducting a data maturity assessment to evaluate data quality, accessibility, standardization, and integration across organizational silos. For public investors, this often involves assessing the consistency of investment data, market feeds, alternative asset documentation, and internally generated research and analytics.

Institutions should also prioritize data readiness for internal analytical use cases before relying heavily on external large language models or vendor platforms. This will help ensure that internal data governance and security controls are sufficiently mature.

Infrastructure

Beyond governance and security, AI adoption depends on robust, future-ready technology infrastructure. AI workloads are computationally intensive, requiring scalable computing environments, high-performance processing, and secure data pipelines capable of handling large volumes of structured and unstructured data.

A critical component of AI readiness is the availability of modern data infrastructure

that allows institutions to efficiently store, process, and govern large and diverse datasets. Many organizations are therefore investing in scalable data platforms such as cloud-native data architectures, data lakes, and data warehouse environments. These platforms enable integrated data access while supporting governance mechanisms for metadata management, data lineage, and controlled access to sensitive information. Importantly, infrastructure investments should remain aligned with clearly defined institutional use cases rather than being driven solely by technological ambition.

Computing capacity is another key enabler of AI adoption. Advanced AI models often require high-performance processing environments, including clusters of Graphic Processing Units (GPU) or Tensor Processing Units (TPU) capable of handling intensive model training and inference workloads. Real-time analytics platforms are also increasingly deployed to support applications such as payment monitoring, anomaly detection, and regulatory analytics. These capabilities allow institutions to process large data volumes and generate insights at the speed required for modern financial operations.

Cloud computing can play a significant role in providing scalable computing capacity and operational flexibility. However, considerations related to data sovereignty, regulatory compliance, and third-party risk often influence the pace of cloud adoption among central banks and public financial institutions. Many institutions, therefore, pursue hybrid architectures that combine public cloud scalability with on-premises control for sensitive data. Findings from the IFC survey indicate that over 40 percent of central banks prefer on-premises infrastructure, while approximately 39 percent are testing cloud-based solutions (BIS IFC 2025).

AI capabilities evolve rapidly, and infrastructure environments thus cannot remain static. Institutions must pursue continuous modernization of their technology platforms to ensure they remain capable of supporting AI workloads and evolving analytical capabilities. This includes upgrading IT platforms to support scalable AI processing, implementing secure Application Programming Interface (API) architectures and integration layers that allow AI systems to interact with existing operational platforms, and establishing controlled deployment environments for AI models. In practice, many public financial institutions deploy private instances of LLMs or hybrid environments that allow sensitive financial data to remain within controlled infrastructure boundaries.

While infrastructure design plays an important role in enabling secure and reliable AI deployment, broader considerations related to data protection, cybersecurity, and operational resilience are equally critical. These aspects are discussed in the following section.

3.4.2. Privacy and Cybersecurity

Central banks consistently report privacy and cybersecurity risks as among their top concerns regarding AI adoption. The IFC Survey identifies privacy breaches, cybersecurity vulnerabilities, and model misuse as among the most material risks constraining AI deployment (BIS IFC 2025). Implementing privacy protection from the outset, consistent with “privacy by design” principles, is both a legal requirement and an ethical obligation. The EU General Data Protection Regulation (GDPR)²⁵ mandates data minimization, lawful processing, and explainability, while the EU Artificial Intelligence Act introduces obligations for high-risk AI use cases.

Some organizations have adopted federated learning to facilitate collaborative model development without sharing raw data, thus supporting privacy and data sovereignty requirements.

Beyond privacy considerations, AI systems introduce new cybersecurity risks at both the data and the model levels. These include threats such as data poisoning, prompt injections, and model inversion posing threats to both operational continuity and institutional credibility.

Effective cybersecurity practices should span every stage of the AI lifecycle, from data ingestion to model deployment. Standards like ISO/IEC 27001 (ISO 2022a) and the NIST Cybersecurity Framework (NIST 2024) outline methods for securing infrastructure, controlling access, encrypting data, and responding to incidents.

AI-specific threat intelligence, including resources like the MITRE ATLAS framework (MITRE ATLAS 2023), documents adversarial tactics relevant to ML, providing guidance for defensive system design and threat modeling. Cybersecurity in this context involves proactively identifying and mitigating potential threats.

3.4.3. Talent and Cultural Readiness

Bridge the Talent Gap

Technological advancements in algorithms and enhanced digital infrastructure for AI are critical components, but technology alone is not sufficient for success. People and processes are both essential forces to ensure the successful adoption of AI.

Across institutions, human capital constraints consistently emerge as a dominant barrier to AI adoption.

²⁵ The European Union approved the General Data Protection Regulation (GDPR) in April 2016; it went into effect on May 25, 2018. See <https://gdpr-info.eu/>.

The IFC Survey identifies shortages in data science expertise, AI governance capacity, and model-risk oversight as more binding constraints than access to algorithms or computing resources (BIS IFC 2025).

This finding aligns with broader cross-sector evidence. A 2024 BCG study reports that 70 percent of successful AI transformations are attributable to institutional readiness in both people and processes, surpassing the impact of technology (20 percent) and algorithms (10 percent) (Loh et al. 2024). This observation has critical implications for central banks and other public-sector organizations aiming to leverage AI responsibly: without strategic investments in talent development, organizational culture, and internal coordination, even advanced AI systems are at risk of underutilization or misapplication.

To prepare, organizations must take a structured approach to workforce development that builds both broad institutional literacy and specialized technical expertise. Many institutions are adopting tiered AI literacy frameworks to support organization-wide capability development. The three tiers include:

- **Foundational level:** Broad AI awareness training for all staff to build a shared understanding of AI concepts, opportunities, and risks. Foundational literacy helps staff understand the basic mechanics of AI and machine learning, recognize limitations such as bias and model uncertainty, and appreciate the ethical and governance implications of AI use.
- **Applied level:** Covering practical AI tools and analytical skills tailored to business users, this level focuses on enabling professionals across research, portfolio management, risk management, compliance, and operations to incorporate AI tools into their workflows responsibly and effectively.
- **Advanced level:** This level encompasses the specialized technical training required for data scientists, engineers, and governance specialists responsible for developing, validating, and overseeing AI systems. This training includes expertise in machine learning engineering, model risk management, AI governance frameworks, and ethical oversight.

In addition to developing internal expertise, institutions often strengthen AI capabilities through partnership-based collaboration models, including engagement with academic institutions, collaboration with external technical experts, and participation in international knowledge-sharing networks among peer public-sector organizations.

This structured approach to capability development does more than build technical competence; it strengthens institutional agility, improves responsible use of AI tools, and enhances the organization's ability to adapt as AI technologies evolve.

Cultural Change

Advocating a responsible AI culture should involve encouraging innovation and ensuring that everyone has a role and is accountable in shaping how AI is used. Establishing foundational literacy supported by formal and mandatory training ensures staff have the requisite knowledge and skills to use AI responsibly. Trust and transparency are essential; AI should be positioned as a tool for human augmentation rather than replacement, and its capabilities, limits, and ethical principles should be embedded from the start.

3.4.4. Operationalizing AI Responsibly

Translating AI concepts into operational practice requires a disciplined and structured implementation approach. Organizations typically focus on carefully selecting early use cases that demonstrate clear value and present manageable levels of risk.

Institutions may also benefit from mapping potential AI applications across the investment and operational value chain, including research, portfolio construction, risk management, compliance monitoring, and reporting. Rather than pursuing broad experimentation, many successful organizations prioritize three to five high-value use cases where AI can deliver measurable improvements. Early applications often include productivity enhancements, such as summarizing earnings calls, analyzing market intelligence, or automating document processing.

Once potential applications are identified, responsible operationalization requires structured experimentation and governance controls before solutions are scaled across the organization. Typically, this involves the following:


- Controlled piloting and sandboxing, including conducting proofs-of-value in controlled environments to validate technical feasibility, governance controls, and measurable benefits before scaling solutions across the organization
- Clearly defined human oversight and escalation mechanisms to ensure that AI outputs remain subject to institutional accountability;
- Documentation of assumptions, model limitations, and design choices to support transparency and auditability; and
- Continuous monitoring for bias, performance drift, and unintended consequences throughout the lifecycle of the AI system. This is the stage at which responsible AI principles become observable in practice and where institutional trust is either reinforced or weakened.

Given that artificial intelligence technologies and regulatory expectations are evolving rapidly, responsible AI operationalization should be viewed as a continuous lifecycle rather than a one-time deployment exercise. Institutions should establish structured processes for ongoing monitoring, model validation, periodic retraining, and

performance review. Regular internal reviews and independent audits can help ensure models remain aligned with institutional objectives, ethical standards, and regulatory expectations. Because AI capabilities and risks will continue to evolve, institutions should periodically revisit their AI strategies and governance frameworks to remain adaptive and responsive to emerging developments.

Once AI systems transition into production environments, institutions must establish mechanisms for continuous monitoring and lifecycle management. Models should be regularly evaluated for performance drift, emerging bias, and changes in data distributions that may affect reliability or fairness. Feedback loops between operational users, data teams, and governance bodies help refine models and ensure systems remain aligned with institutional objectives.

Managing the AI model lifecycle, therefore, becomes a critical operational responsibility. Institutions should establish processes for model version control, performance monitoring thresholds that trigger reevaluation, retraining schedules, and documentation standards such as model cards or system fact sheets. These lifecycle management practices—often referred to as MLOps (Machine Learning Operations) help ensure that AI systems remain reliable, auditable, and aligned with institutional governance standards throughout their deployment.



4. CONCLUSIONS

AI adoption is no longer a question of whether to proceed, but of how to implement AI effectively to achieve the highest impact. Over the past few years, rapid technological advances, particularly in Gen AI, have generated both excitement and a sense of urgency in the asset management industry. Global asset managers have been racing to adopt AI and preserve a competitive edge. For private asset managers, the pressure to adopt AI has been intense, driven by competitive dynamics, cost reduction imperatives, and the expectation of alpha generation.

Public investors, even if not early adopters, are also required to act as early followers by closely tracking industry developments and accelerating institution-level AI initiatives. The same surge in capability comes with a fundamentally different set of institutional stakes. Their mandates extend beyond market outperformance to encompass long-term stability and accountability to citizens or beneficiaries.


The range of use cases in the asset management industry highlights both the breath of opportunity and diversity of potential impact. While AI can be applied across investment, risk, operations and reporting functions, the value derived from these applications differs across institutional contexts. Effective adoption therefore requires selecting and prioritizing use cases that are scalable, relevant and aligned with institutional needs, rather than pursuing initiatives driven by external demands.

Translating opportunity into sustainable value requires a clear connection between vision, ambition and execution. Governance, risk management and operational processes must evolve in parallel with AI adoption. As institutions move from growing AI awareness to building AI readiness, the ability to scale AI responsibly will depend on both technological advances and the AI maturity of the underlying institutional environment.

This includes establishing clear accountability, embedding lifecycle oversight, strengthening data and infrastructure, and fostering cross-functional collaboration. In many cases, meaningful progress will be driven by strengthening the conditions that enable initiatives to perform reliably and sustainably, alongside a disciplined approach to expansion.

As technology advances rapidly, the evolving regulatory expectations will continue to reshape the AI landscape. Institutions will need to keep up and remain adaptable, combining strategic direction with disciplined experimentation and continuous learning. This will require technical capabilities, coupled with sound judgement on navigating trade-offs between innovation, control, compliance and resilience.

Ultimately, the value depends on how effectively it is embedded within strong institutional foundations, while safeguarding the trust, resilience, and integrity that underpin institutional mandates.



REFERENCES

Aoki, Minoru, Hashim Shah, and Joel Viklund. 2025. "Hypothetical Market Data Scenario Generation Using Generative AI." Nomura Research Institute (NRI), August 4, 2025.

https://www.nri.com/en/knowledge/report/20250804_1.html.

Bahoo, Salman, Marco Cucculelli, Xhoana Goga, and Jasmine Mondolo. 2024. "Artificial Intelligence in Finance: A Comprehensive Review Through Bibliometric and Content Analysis." *SN Business & Economics* 4 (23). <https://doi.org/10.1007/s43546-023-00618-x>

Baig, Aamer, Vik Sohoni, and Xavier Lhuer. 2024. "Unlocking Value from Technology in Banking: An Investor Lens." McKinsey & Company, October 23, 2024. <https://www.mckinsey.com/industries/financial-services/our-insights/unlocking-value-from-technology-in-banking-an-investor-lens>.

Bartram, Söhnke M., Jürgen Branke, and Mehrshad Motahari. 2020. "Artificial Intelligence in Asset Management." CFA Institute Research Foundation. <https://www.cfainstitute.org/sites/default/files/-/media/documents/book/rf-lit-review/2020/rflr-artificial-intelligence-in-asset-management.pdf>.

BCG. 2024. "AI and the Next Wave of Transformation: Global Asset Management Report 2024." BCG (Boston Consulting Group). <https://www.bcg.com/publications/2024/ai-next-wave-of-transformation>.

Bianchi, Daniele, Matthias Büchner, and Andrea Tamoni. 2021. "Bond Risk Premiums with Machine Learning." *Review of Financial Studies* 34 (2): 1046–89.

BIS. 2024. "Artificial Intelligence in Central Banking. <https://www.bis.org/publ/bisbull84.htm>

BIS CGRM. 2025. Governance of AI Adoption in Central Banks. Bank for International Settlements (BIS) Consultative Group on Risk Management (CGRM). <https://www.bis.org/publ/othp90.pdf>.

BIS IFC. 2025. Governance and Implementation of Artificial Intelligence in Central Banks. IFC Report No. 18, Bank for International Settlements (BIS), Irving Fisher Committee (IFC) Survey. https://www.bis.org/ifc/publ/ifc_report_18.pdf.

BIS Innovation Hub. 2024. "Project Gaia: Enabling Climate Risk Analysis Using Generative AI." Bank for International Settlements (BIS) Innovation Hub. <https://www.bis.org/publ/othp84.htm>.

BlackRock. 2024. Augmented Investment Management. Accessed August 11, 2025.

<https://www.blackrock.com/ca/institutional/en/insights/investment-actions/augmented-investment-management>.

BoE. 2025. The Bank's Artificial Intelligence (AI) Strategy. Bank of England, August 21, 2025. <https://www.bankofengland.co.uk/about/governance-and-funding/staff-codes-and-policies/the-banks-ai-strategy>.

Boier, Ioana. 2024. "Generating Financial Market Scenarios Using NVIDIA NIM." NVIDIA Technical Blog. <https://developer.nvidia.com/blog/generating-financial-market-scenarios-using-nvidia-nim/>.

Cao, Larry. 2023. Handbook of Artificial Intelligence and Big Data Applications in Investments. CFA Institute Research Foundation, March 27, 2023.

CNBC. 2025. "OpenAI's ChatGPT Hits 700 Million Weekly Users as AI Arms Race Heats Up." CNBC, August 4, 2025. <https://www.cnn.com/2025/08/04/openai-chatgpt-700-million-users.html>.

DAMA. 2024. DAMA® Data Management Body of Knowledge (DAMA-DMBOK®). DAMA International. <https://dama.org/learning-resources/dama-data-management-body-of-knowledge-dmbok/>.

Dataiku. 2025. Global AI Confessions Report: CEO Edition. Dataiku. <https://pages.dataiku.com/global-ai-confessions-report-old>

DeMiguel, Victor, Javier Gil-Bazo, Francisco J. Nogales, and André A. P. Santos. 2023. "Machine Learning and Fund Characteristics Help to Select Mutual Funds with Positive Alpha." *Journal of Financial Economics* 150 (3). <https://doi.org/10.1016/j.jfineco.2023.103737>.

Easthope, David. 2025. AlternativeData 2025: Fueling the AI-driven Investment Revolution. Coalition Greenwich, December 17, 2025. <https://www.greenwich.com/market-structure-technology/alternative-data-2025-fueling-ai-driven-investment-revolution>.

Economic Times, Global Desk. 2025. "Company That Sacked 700 Workers with AI Now Regrets It; Scrambles to Rehire as Automation Goes Horribly Wrong." *Economic Times*, June 9, 2025. <https://economictimes.indiatimes.com/news/international/us/company-that-sacked-700-workers-with-ai-now-regrets-it-scrambles-to-rehire-as-automation-goes-horribly-wrong/articleshow/121732999.cms>.

EU. 2024. EU AI Act. European Union, June 13, 2025. <https://artificialintelligenceact.eu/ai-act-explorer/>.

Exabel. 2025. "Alternative Data Buy-side Insights & Trends 2025." 2025 Research Report, Exabel. [Exabel.com](https://www.exabel.com).

Financial Conduct Authority. 2025. "Generating and Using Synthetic Data for Models in Financial Services: Governance Considerations." Financial Conduct Authority, August 2025. <https://www.fca.org.uk/publication/corporate/synthetic-data-models-financial-services-governance.pdf>.

Fontanille, Cédric, Robert Kosowski, Ravi Ramakrishnan, Eugenio Rodriguez, and Giulio Trichilo. 2020. "Machine Learning for Fund Selection." UNIGESTION, October 2020. <https://www.unigestion.com/wp-content/uploads/2020/12/20201008-White-Paper-Machine-Learning-FINAL.pdf>.

Fu, Xi, Xiaoxi Wu, and Zhifang Zhang. 2021. "The Information Role of Earnings Conference Call Tone: Evidence from Stock Price Crash Risk." *Journal of Business Ethics* 73 (2021): 643–60. <https://doi.org/10.1007/s10551-019-04326-1>.

GoSearch. 2025. "The Ultimate Breakdown of Different AI Types and Models." GoSearch, November 25, 2025. <https://www.gosearch.ai/blog/breakdown-of-different-ai-types-and-models/#:~:text=Key%20Takeaways,how%20organizations%20evaluate%20AI%20capabilities>.

Grand View Research. 2025. Alternative Data Market (2025–2030). Grand View Research. <https://www.grandviewresearch.com/industry-analysis/alternative-data-market>.

Hansen, Anne Lundgaard, and Sophia Kazinnik. 2023. "Can ChatGPT Decipher Fedspeak?" April 10, 2023. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4399406.

Harbert, Tam. 2021. "Tapping the Power of Unstructured Data." MIT Sloan School of Management, February 1, 2021. <https://mitsloan.mit.edu/ideas-made-to-matter/tapping-power-unstructured-data>.

HSBC. 2024. "HSBC Reserve Management Trends." HSBC, April 15, 2024.

IEA. 2025. "Energy and AI." International Energy Agency, April 10, 2025. <https://www.iea.org/reports/energy-and-ai>.

IIF and EY. 2025. "IIF-EY Annual Survey Report on AI/ML Use in Financial Services." Institute of International Finance and EY, October 16, 2025.

Index Industry Association. 2024. "2024 Survey of Asset Managers." Index Industry Association. <https://www.indexindustry.org/wp-content/uploads/2024-IIA-Survey-of-Global-Asset-Managers.pdf>.

Invesco. 2024. "Invesco Global Sovereign Asset Management Study 2024." Invesco. <https://www.invesco.com/content/dam/invesco/igsams/en/docs/Invesco-global-sovereign-asset-management-study-2024.pdf>.

ISO. 2022a. ISO/IEC 27001: Information Security, Cybersecurity and Privacy Protection—Information Security Management Systems—Requirements. International Organization for Standardization (ISO), October 2022. <https://www.iso.org/standard/27001>.

ISO. 2022b. ISO/IEC 38507: Information Technology—Governance of IT—Governance Implications of the Use of Artificial Intelligence by Organizations. International Organization for Standardization (ISO), April 2022. <https://www.iso.org/standard/56641.html>.

Israel, Ronen, Bryan T. Kelly, and Tobias J. Moskowitz. 2020. "Can Machines 'Learn' Finance?" *Journal of Investment Management* 18 (2): 23–36.

Jiang, Weiwei. 2021. "Applications of Deep Learning in Stock Market: Recent Progress." *Expert Systems with Applications* 184 (Dec. 2021). <https://doi.org/10.48550/arXiv.2003.01859>.

Kong, Yaxuan, Yuqi Nie, Xiaowen Dong, John M. Mulvey, H. Vincent Poor, Qingsong Wen, and Stefan Zohren. 2024. "Large Language Models for Financial and Investment Management: Applications and Benchmarks." *Journal of Portfolio Management* 51 (2): 162–210. DOI: 10.3905/jpm.2024.1.645.

Krause, Helen H., Alex Miller, Savina Chahal, Lisa Finstrom, Adnan N. Memon, and Yirou Yu. 2024. *AI in Investment Management—The Pursuit of a Competitive Edge*. CITI, June 2024.

LEZMI, Edmond, and Jiali XU. 2023. "Time Series Forecasting with Transformer Models and Applications to Asset Management." Amundi Research Center. <https://research-center.amundi.com/article/time-series-forecasting-transformer-models-and-application-asset-management>.

Loh, Hean-Ho, Vinciane Beauchene, Vladimir Lukic, and Rajiv Shenoy. 2024. *Five Must-Haves for Effective AI Upskilling*. Boston Consulting Group, October 8, 2024. <https://www.bcg.com/publications/2024/five-must-haves-for-ai-upskilling>.

Mailagaha Kumbure, Mahinda, Christoph Lohrmann, Pasi Luukka, and Jari Porras. 2022. "Machine Learning Techniques and Data for Stock Market Forecasting: A Literature Review." *Expert Systems with Applications* 197 (C). <https://doi.org/10.1016/j.eswa.2022.11665>.

McCarthy, John. 2007. "What Is Artificial Intelligence?" Computer Science Department, Stanford University, November 12, 2007. <https://www-formal.stanford.edu/jmc/whatisai.pdf>.

Mercer. 2024. "AI Integration in Investment Management: 2024 Global Manager Survey." Mercer. <https://www.mercer.com/en-us/insights/investments/portfolio-strategies/ai-in-investment-management-survey/>.

MITRE ATLAS. 2023. MITRE ATLAS. MITRE. <https://atlas.mitre.org/matrices/ATLAS>.

Morgan Stanley and Oliver Wyman. 2023. The AI Tipping Point. Morgan Stanley and Oliver Wyman, October 2023. Accessed August 7, 2025. https://www.oliverwyman.com/content/dam/oliver-wyman/v2/publications/2023/october/Oliver_Wyman_Morgan_Stanley_Global_Wealth_and_Asset_Management_report_2023_The_Generative_AI_Tipping%20Point1.pdf.

Ning, Brian, Franco Ho Ting Lin, and Sebastian Jaimungal. 2022. "Double Deep Q-Learning for Optimal Execution." *Applied Mathematical Finance* 28 (4) 361–80.

NIST. 2023. AI Risk Management Framework. National Institute of Standards and Technology, January 26, 2023. <https://www.nist.gov/itl/ai-risk-management-framework>

NIST. 2024. Cybersecurity Framework. National Institute of Standards and Technology, February 26, 2024. <https://www.nist.gov/cyberframework>.

Noy, Shakked, and Whitney Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *MIT Economics*, March 2, 2023. https://economics.mit.edu/sites/default/files/inline-files/Noy_Zhang_1.pdf.

Nti, I. K., A. F. Adekoya, and B. A. Weyori. 2020. "A Comprehensive Evaluation of Ensemble Learning for Stock-Market Prediction." *Journal of Big Data* 7 (20). <https://doi.org/10.1186/s40537-020-00299-5>.

OECD. 2024. OECD AI Principles. Organisation for Economic Co-operation and Development, May 2024. <https://oecd.ai/en/ai-principles>.

Parsons, Joe. 2023. "JP Morgan Pulls Plug on Deep Learning Model for FX Algos." *FX Markets*, October 18, 2023. <https://www.fx-markets.com/tech-and-data/7949034/jp-morgan-pulls-plug-on-deep-learning-model-for-fx-algos>.

Prado, de, Marcos López. 2016. "Building Diversified Portfolios that Outperform Out-of-Sample." *Journal of Portfolio Management* (2016): 59–69. <https://doi.org/10.3905/jpm.2016.42.4.059>.

Robeco. 2024. The Current State of AI in Asset Management. Robeco, June 2016. <https://www.robeco.com/en-us/insights/2023/12/the-current-state-of-ai-in-asset-management>.

Sandwidi, Blaise, and Suneer P. Mukkolakal. 2022. "Transformer-Based Approach for a Sustainability Term-Based Sentiment Analysis (STBSA)." Proceedings of the Second Workshop on NLP for Positive Impact (NLP4PI): 157–70.

Stanford University. 2025. "AI Index Report 2025." Human-Centered Artificial Intelligence, Stanford University.

Swinkels, Laurens, and Tobias Hoogteijling. 2022. "Forecasting Stock Crash Risk With Machine Learning." Robeco, June 14, 2022. <https://www.robeco.com/en-us/insights/2022/06/forecasting-stock-crash-risk-with-machine-learning>.

Tajiri, Takao, Tomonari Murakami, Masanori Hashido, and Hiroaki Kitano. 2020. "A Study on the Use of Artificial Intelligence for Learning Characteristics of Funds' Behavior (Summary Report)." Government Pension Investment Fund, Government of Japan, September 2020. https://www.gpif.go.jp/en/investment/ai_report_summary_en.pdf.

Tang, Wenjin, Hui Bu, Yuan Zuo, and Junjie Wu. 2024. "Unlocking the Power of the Topic Content In News Headlines: BERTopic for Predicting Chinese Corporate Bond Defaults." Finance Research Letters 62, Part A (April 24). <https://doi.org/10.1016/j.frl.2024.105062>.

Turner, Matt. 2023. "How One of the Biggest Investors in the World Uses AI to Move Slower, Trade Less, and Make More Money." Business Insider, June 23, 2023. <https://www.businessinsider.com/ai-artificial-intelligence-investing-strategy-technique-norges-bank-investment-management-2023-6>.

UNCTAD. 2025. Technology and Innovation Report 2025: Inclusive Artificial Intelligence for Development. United Nations Conference on Trade and Development, April 7, 2025.

UNESCO. 2022. "Recommendation on the Ethics of Artificial Intelligence." Adopted November 23, 2021. United Nations Educational, Scientific and Cultural Organization (UNESCO).

Uysal, A. Sinem, and John M. Mulvey. 2021. "A Machine Learning Approach in Regime-Switching Risk Parity Portfolios." Journal of Financial Data Science (April 30, 2021): 87–108.

Vamvourellis, Dimitrios, Máté Toth, Snigdha Bhagat, Dhruv Desai, Dhagash Mehta, and Stefano Pasquali. 2023. "Company Similarity Using Large Language Models." arXiv, August 15, 2023. <https://arxiv.org/abs/2308.08031>.

World Bank Group. 2025. Reserve Management Survey Report 2025: Insights into Public Asset Management. The World Bank Group, November 11, 2025. <https://hdl.handle.net/10986/43953>.

World Economic Forum, Personal Data Protection Commission Singapore. 2020. "Companion to the Model AI Governance Framework." January 2020. <https://www.pdpc.gov.sg/-/media/files/pdpc/pdf-files/resource-for-organisation/ai/sgisago.pdf>.

Yadav, Vikas, and Steven Bethard. 2019. "A Survey on Recent Advances in Named Entity Recognition from Deep Learning Models." arXiv, October 25, 2019. <https://doi.org/10.48550/arXiv.1910.11470>.

Yeyati, Eduardo Levy. 2025. "Hybrid Jobs: How AI Is Rewriting Work in Finance." Brookings, July 10, 2025. <https://www.brookings.edu/articles/hybrid-jobs-how-ai-is-rewriting-work-in-finance/>.

Yu, Xinli, Zheng Chen, Yuan Ling, Shujing Dong, Zongyi Liu, and Yanbin Lu. 2023. "Temporal Data Meets LLM—Explainable Financial Time Series Forecasting." Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track, 739–53. Association for Computational Linguistics, June 19, 2023. <https://doi.org/10.48550/arXiv.2306.11025>.

AI Primer for Public Investors:

From AI Awareness to AI Readiness



BANK OF KOREA



WORLD BANK GROUP