

FSI Briefs

No 26

Starting with the basics: a stocktake
of gen AI applications in supervision

Jermy Prenio

June 2025

FSI Briefs are written by staff members of the Financial Stability Institute (FSI) of the Bank for International Settlements (BIS), sometimes in cooperation with other experts. They are short notes on regulatory and supervisory subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS or the Basel-based standard setting bodies.

Authorised by the Chair of the FSI, Fernando Restoy.

This publication is available on the BIS website (www.bis.org). To contact the BIS Media and Public Relations team, please email media@bis.org. You can sign up for email alerts at www.bis.org/emailalerts.htm.

© *Bank for International Settlements 2025. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

ISSN 2708-1117 (online)
ISBN 978-92-9259-861-7 (online)

Starting with the basics: a stocktake of gen AI applications in supervision¹

Highlights

- *Many financial authorities are already experimenting with, developing or using generative artificial intelligence (gen AI) applications for supervision purposes.*
- *Financial authorities seek to leverage the new technology to find information more efficiently, but their gen AI activities are hampered by outdated information technology (IT) infrastructure, data security concerns and a lack of technical skills.*
- *Most of the reported gen AI applications in supervision can be grouped into three categories: (i) basic document processing; (ii) knowledge management; and (iii) document review. Most “in use” applications fall into the first category; development work is spread out across the three categories; and experiments are concentrated in the second and third categories.*
- *The main challenges identified in integrating gen AI applications in supervision are user acceptance and inaccuracies in information provided. These challenges will likely intensify as financial authorities move to more complex gen AI use cases.*

1. Introduction

The introduction of a generative artificial intelligence (gen AI) application in 2022 ushered in a flurry of exploration and experimentation to harness the potential benefits of the technology in the financial sector. This has been aided by the accessibility of the new technology, with cloud service providers expanding their offerings to include gen AI applications. While financial institutions’ gen AI use cases focus on enhancing operational efficiency, they are continuously exploring other potential use cases and are reportedly investing heavily in AI in general.²

This exploration and experimentation around the potential benefits of gen AI are not confined to financial institutions. The development of gen AI also provides fodder for experiments by financial sector supervisors to leverage innovative technologies to aid supervision work, or so-called suptech. In general, supervisory authorities have always been interested in exploring applications that can help them go through huge volumes of narrative reports and other relevant text documents. Thus, over the years, there have been many examples of tools using natural language processing. Gen AI has the potential to significantly improve the capabilities of these tools. Gen AI can also potentially make suptech applications more accessible to users, thus addressing one of the challenges of effectively integrating such applications in supervisory processes.³

This paper provides an overview of the state of gen AI applications in financial supervision. It is based on a survey of authorities represented in the Informal Suptech Network, which was established by

¹ Jermy Prenio (jermy.prenio@bis.org), Bank for International Settlements. The author is grateful to members of the Informal Suptech Network who responded to the survey; to Saleh Alqaryan, Samir Kiuhan and Nico Lauridsen for their helpful comments; and to Anna Henzmann for administrative support.

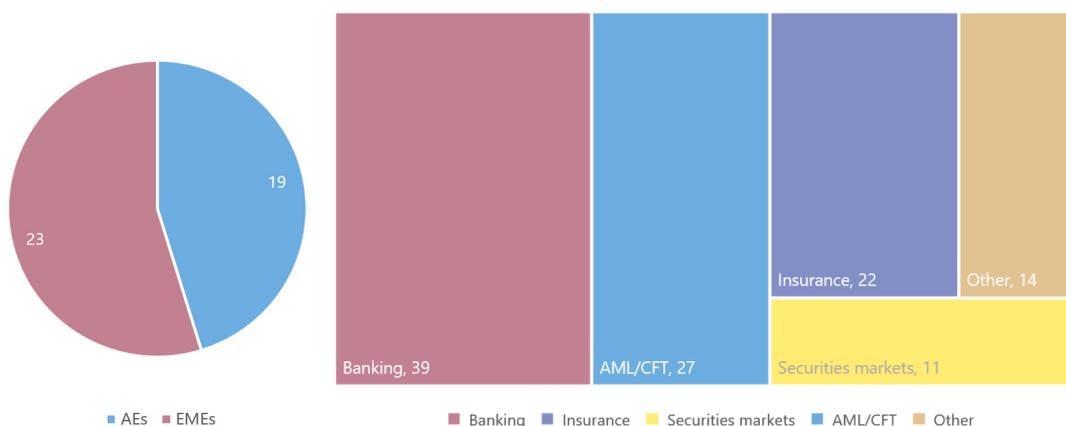
² See Crisanto et al (2024).

³ See Prenio (2024).

the Financial Stability Institute (FSI) in 2018 as a forum for information exchange and sharing of practices on how financial authorities integrate suptech in financial supervision. The survey was conducted from January to March 2025, and responses were received from 42 authorities,⁴ with almost balanced representation from advanced economies (AEs) and emerging market economies (EMEs). Respondent authorities have varied mandates and include banking supervision, insurance supervision, anti-money laundering/countering the financing of terrorism (AML/CFT) and securities markets oversight (Graph 1). Section 2 outlines the gen AI activities of financial sector supervisors—ie whether they are experimenting with, developing or already using gen AI applications. Section 3 describes the gen AI use cases in supervision, their perceived benefits, how they are deployed and the challenges faced in integrating them in supervisory processes. Section 4 concludes.

Profile of authorities that responded to the survey

Graph 1



Source: FSI.

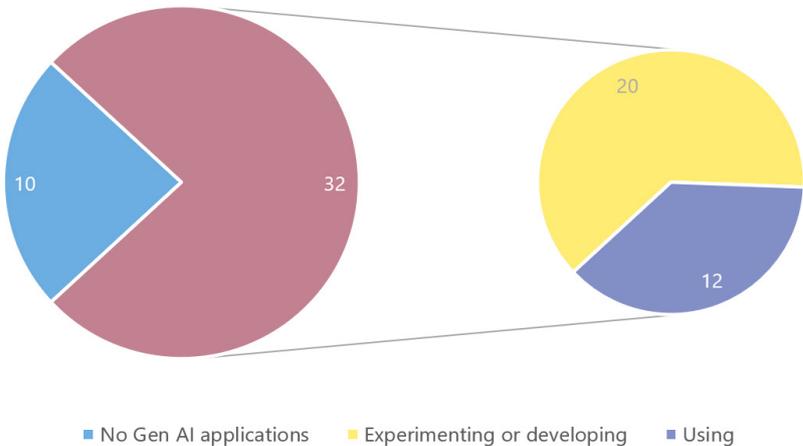
2. Authorities' gen AI activities

Of the 42 respondents, 32 reported that they are already either experimenting with, developing or using gen AI applications for supervision. Of these authorities, only 12 are already using such applications (Graph 2). Comparing gen AI activities between authorities from AEs and EMEs, the survey responses show no marked difference in the number of authorities in each type of jurisdiction that are experimenting, developing and using gen AI applications (Graph 3).

⁴ See Annex 1 for a list of respondent authorities.

Authorities with generative artificial intelligence (gen AI) applications in supervision

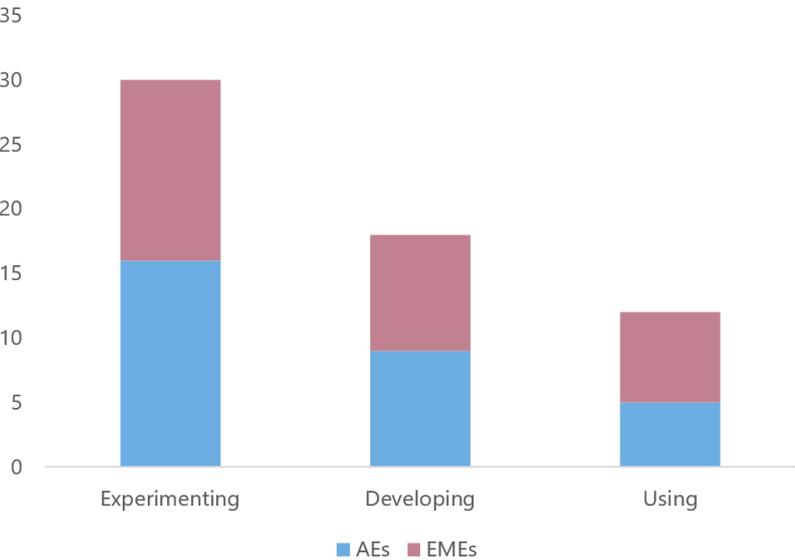
Graph 2



Source: FSI.

Authorities in advanced economies (AEs) and emerging market economies (EMEs) that are experimenting, developing or using generative AI

Graph 3



Source: FSI.

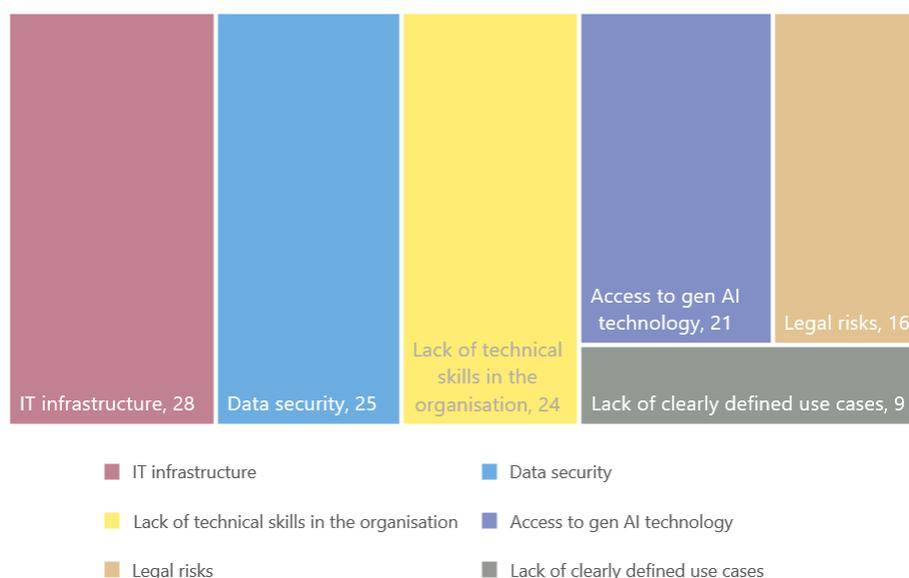
Respondents reported several challenges their organisations face in experimenting, developing, using or even starting to explore gen AI applications for supervision. The challenges respondents most often cited are outdated information technology (IT) infrastructure,⁵ data security and a lack of technical

⁵ The reluctance of many supervisory authorities to use a cloud infrastructure due to security concerns affects AI adoption and AI maturity. See for example Gambacorta et al (forthcoming).

skills. Many respondents also cited access to the gen AI technology as a challenge. This may relate to the cost of rolling out gen AI applications more widely to supervision staff. Legal risks, particularly when gen AI applications use data in the public domain, as well as legal risks that may arise from data breaches (eg when confidential data are inadvertently or advertently fed into a third-party application) are also a concern. In addition, some authorities are still not clear how they would use gen AI for supervision purposes.

Challenges in exploring generative artificial intelligence (gen AI) applications

Graph 4



Source: FSI.

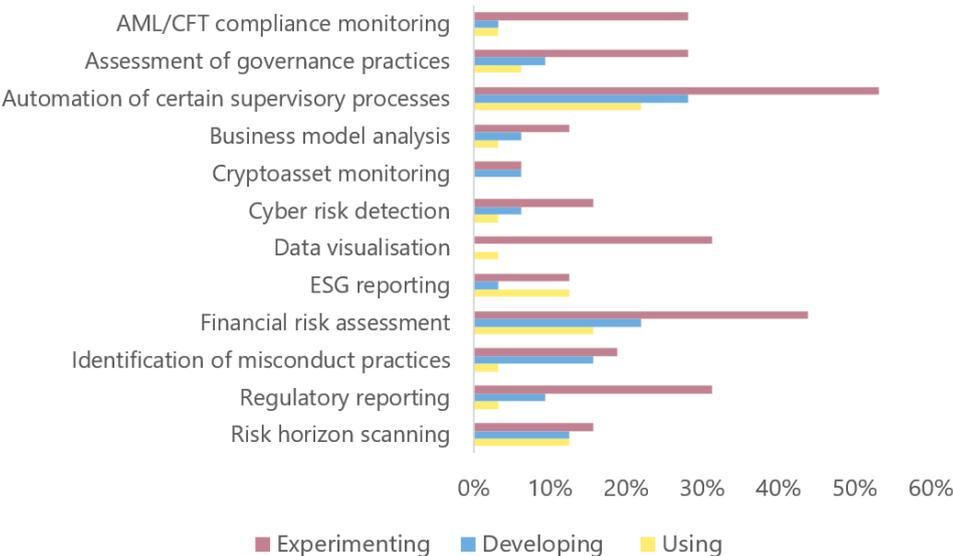
3. Gen AI use cases in supervision

Respondents were asked in which supervision areas they are experimenting, developing or using gen AI applications. The supervision areas are based on the classification used in Prenio (2024) and Prenio et al (2024). The supervision areas with gen AI applications are the same areas in which authorities are actively deploying supotech tools, but we see a notable concentration in automation of supervisory processes (Graph 5).

Like financial institutions, financial sector supervisors seem to be leveraging or wishing to leverage gen AI applications to enhance operational efficiency. Indeed, the survey responses show this, with most respondents indicating increased efficiency in finding information as the benefit gen AI applications bring to supervision. Given concerns about hallucination in gen AI models, the optimism of some respondents that gen AI applications can provide them with more accurate information is surprising and interesting (Graph 6).

Supervision areas with generative artificial intelligence applications

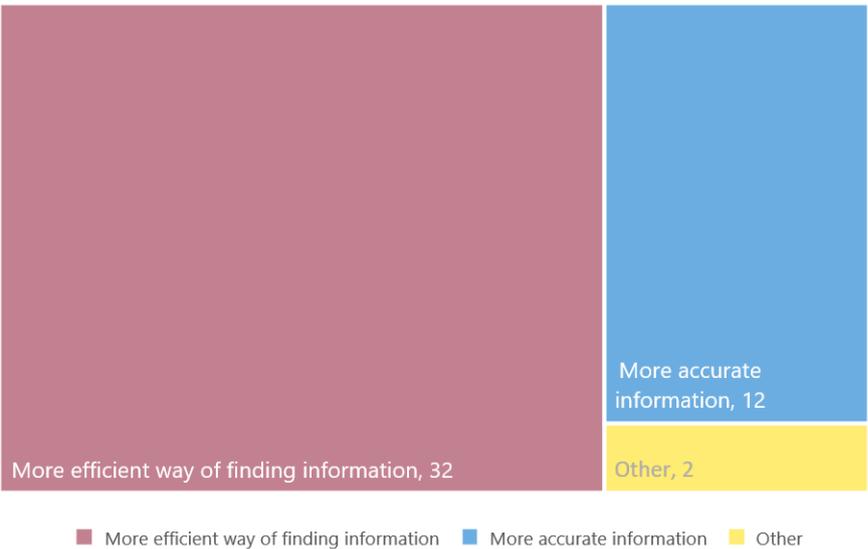
Graph 5



Source: FSI.

Benefits that generative artificial intelligence applications bring to supervision

Graph 6



Source: FSI.

Examples of specific use cases

Respondents were also asked to describe at least one gen AI application they are currently experimenting with, developing and using. Based on the responses, we can group most of the gen AI applications in supervision into three categories: (i) basic document processing; (ii) knowledge management; and (iii) document review. Most of the gen AI applications already being used fall into the first category;

applications in development are spread out quite evenly across the three categories; and those under experimental stage are notably concentrated in the second and third categories.

Basic document processing

Gen AI applications currently in use focus on aiding supervisors with basic document tasks, such as drafting and editing inspection reports, supervisory decision documents and communications to supervised entities. These applications are also used to summarise, translate or extract key information from relevant news articles or regulatory submissions. Here, publicly available gen AI tools like ChatGPT, Microsoft Copilot, DeepSeek and Perplexity are widely used.

Knowledge management

Gen AI applications for knowledge management are those that can help supervisors navigate the full body of regulations and other guidance issued by the authority. These applications can help supervisors respond to queries on specific regulations from supervised entities or from senior management, conduct regulatory research, review whether supervised entities' submissions comply with regulations and inform supervisory decisions or actions more broadly. Such applications are being trained on authorities' documents through retrieval augmented generation. Also included in this category are gen AI applications that can enable supervisors to query their regulatory database using natural language (ie without the need to learn programming language). In general, the applications are being designed as "co-pilots" or AI assistants or chatbots.

The following are examples of use cases in this category:

- Supervisors, particularly those new in the job, spend a great deal of time looking for relevant regulations and keeping abreast of regulatory changes. One supervisory authority is developing a gen AI application with a chat interface that allows supervisors to retrieve regulatory information easily and in full transparency. The chatbot responds to questions about regulations and provides references and citations to the document it used to generate the answer, thereby eliminating the "black-box" phenomenon commonly present in other chatbots. It also generates and suggests follow-up questions and can summarise and compare regulations. Users can also filter by national or regional regulations. This application will therefore streamline the process of finding relevant regulations and highlighting regulatory changes, thereby supporting informed supervisory decisions.
- Similar to the above, another supervisory authority is developing a regulatory chatbot using a large language model (LLM) that retrieves information from a knowledge base containing all relevant laws and the authority's regulations. The goal is to provide supervisors with quick and precise access to regulations of interest – including their details – by posing targeted questions. The authority is also evaluating the use of Graph RAG, which is an enhancement of the traditional RAG (retrieval augmented generation) technique through incorporating knowledge graphs. By using a knowledge graph of regulatory data, Graph RAG can potentially provide more nuanced and accurate context to the information generated. The authority is still comparing the long-term computational cost of Graph RAG with standard RAG.
- Another supervisory authority is also developing a "knowledge-bound" LLM chatbot that is trained on the authority's regulatory and supervisory requirements. The application will allow supervisors to retrieve relevant information that includes accurate citations. Users will be able to enquire about any regulatory policy or requirement without the need to know the precise keywords. The application will therefore facilitate supervisors' understanding of relevant regulations, enhancing the accuracy and efficiency of supervisory activities. To mitigate hallucination risks, the development process involves the creation of a robust knowledge base incorporating comprehensive semantic ranking techniques and acronym dictionaries. In addition, the adoption of a chain-of-thought reasoning method and implementation of

guardrails within the RAG workflow optimise the searches and help ensure outputs are properly aligned with validated sources.

- One supervisory authority provides support to the internal alternative dispute resolution (ADR) system, which decides on disputes between financial institutions and their customers. In deciding on cases, it is important to ensure consistency of treatment of similar cases and adherence to applicable laws. Doing so, however, is challenging due to the complexity and volume of cases and the diversity of subjects covered. The authority has therefore decided to develop a gen AI-powered chatbot that is being trained in past cases that have already been decided by the ADR. The chatbot, which is only available internally and accessible to the supervisory authority, will allow semantic search of previous cases, answer queries on specific cases based on the ADR's past decisions and draft a summary of specific cases.

Document review

Gen AI applications for document review are those that can go through supervised entities' documents (eg regulatory submissions, annual disclosures, insurance policies, recovery plans) and assess whether they comply with regulations. In a way, applications for document review are an extension of the applications for knowledge management. While the latter provides an AI assistant that will help identify relevant regulations the supervisor can use in assessing regulatory submissions, the former automates even the assessment part. Other applications in this category are those that can compare supervised entities' documents with other sources of information (eg whether relevant publicly available information is not mentioned in financial statements or whether submissions are consistent with other regulatory reports).

The following are examples of use cases in this category:

- During onsite inspections, supervisors must review voluminous documents submitted by financial institutions to extract insights and track issues. To help their supervisors in this task, an authority has built a proof-of-concept using gen AI that will serve as a smart assistant to help supervisors locate and navigate documents more effectively. In particular, the application aims to: (i) extract relevant information from a large set of documents; (ii) compare content across documents; and (iii) review qualitative information over time to identify trends and patterns. The application can also be used for benchmarking a financial institution with its peers and for finding similar issues from past inspections. The application will have several components, including for:
 - converting scanned documents or image files into machine readable text and looking through embedded documents
 - providing a preview of each document so users can choose relevant documents to assess
 - generating a list of references where insights have been extracted and explaining how the insights are relevant.
- As part of their evaluation of the effectiveness of governance in financial institutions, supervisors typically assess minutes of board meetings. One supervisory authority is developing an LLM-based application that will help with this assessment. The application will be able to provide supervisors with information such as:
 - the level of dominance of each director in the discussion
 - a mapping of topics discussed by the board and diversity of opinions during the discussion
 - an assessment of the institution's level of reputational risk based on these discussions.

The application will also be able to accumulate findings from previous board meeting minutes and compare changes over time in the level of engagement of each director. At a later stage, the plan is to use the application to assess other documents. However, the challenge for now is ensuring information security given the sensitive nature of information being assessed.

- Payment supervisors annually receive submission of compliance assurance reports from a large number of supervised payment service providers (ie money services businesses and non-bank e-money issuers). These assurance reports, which are prepared by their external auditors, are lengthy and have no consistent format (eg scanned PDF, PowerPoint slides, tabular format). Supervisors therefore face significant challenges in manually reviewing these reports and in identifying non-compliance issues. To address these challenges, one authority developed an application that integrates two core components:
 - an optical character recognition module that automatically extracts text from various report formats, including images and tables
 - an LLM module that summarises the extracted content by topic and categorises the level of compliance with regulations.

The output of the application is visualised in a centralised cloud-based dashboard platform for ease of monitoring by the supervisors. The application helps supervisors identify supervised entities to focus on in upcoming examinations and those that should be subject to enforcement actions.

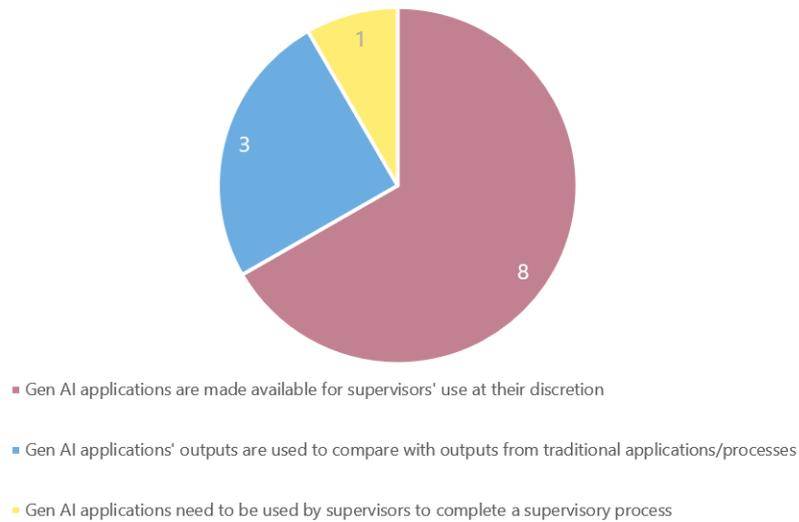
- Systemically important banks are required to maintain recovery plans that identify options to restore financial strength and viability when the bank comes under severe stress. In some jurisdictions, authorities have implemented standardised templates for assessing recovery plans. These standardised templates are especially valuable for guiding gen AI in evaluating recovery plans, particularly because the plans themselves are often unstructured and vary significantly in format and content. One authority is developing a gen AI application that will enhance this assessment process by automating the extraction of pertinent information and the comparison with relevant legal/regulatory requirements. In particular, the application:
 - verifies the presence of specific elements in recovery plans
 - classifies them based on predefined criteria
 - provides rationale for the findings
 - identifies aspects that need further supervisory review.

In this regard, supervisors make the final evaluation with inputs from the AI assessment. In its current state, the application is maintained and being run centrally without an interface for end users. The output is presented to end users in the form of an Excel file.

It should be noted that the above categories of gen AI applications in supervision are based on most often mentioned use cases in the survey responses. However, a few other reported applications do not fall into the above categories. These include use cases for coding assistance or providing alerts (eg scanning social media or websites for risks or other developments that are relevant to supervisors). While a wider and more forward-looking mapping of use cases into different categories would be helpful, the ones that the survey identified so far hopefully provide some useful examples for authorities that still lack clarity as to how gen AI could be used for supervision purposes.

Integration of gen AI applications in supervisory processes

For most authorities with applications already in use, supervisors use these applications at their discretion. Only one authority reported having a gen AI application that is necessary to complete a supervisory process (Graph 7). This suggests that gen AI may need to further develop before it gets integrated into supervision.

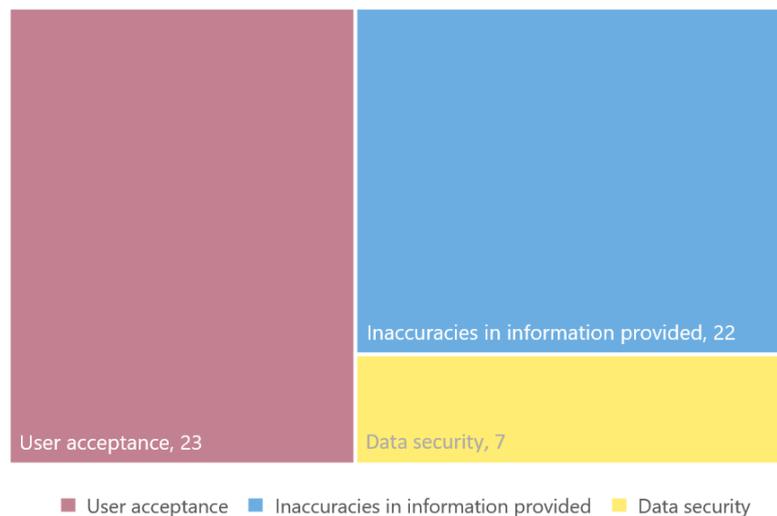


Source: FSI.

Respondents often cited two main challenges in integrating gen AI into supervision: user acceptance and inaccuracies in information provided (Graph 8). The latter seems to contradict the responses showed in Graph 6 that some authorities expect more accurate information from gen AI applications. It could be that authorities are less concerned about hallucination if gen AI applications are used only as reference tools (ie as providers of supplementary information at supervisors' discretion) than they would be if the applications are integrated into supervisory processes (ie having the applications as unavoidable or critical parts of supervisory processes).

The two main challenges are related. User acceptance will not really improve if the applications are providing inaccurate information. User acceptance may also be related to lack of explainability of gen AI, which also hinders financial institutions' wide-scale deployment of the technology in customer-facing use cases.⁶ After all, supervisory opinions, decisions and actions need to be supported by a clear and sound rationale. These issues highlight the importance of involving end users of gen AI applications throughout the development process, particularly in testing and subsequent improvements so they can validate the applications' outputs. This collaborative approach can enhance user trust, acceptance and adoption.

⁶ Crisanto et al (2024).



Source: FSI.

4. Conclusions

Financial sector supervisors are actively exploring the benefits of gen AI. Many supervisory authorities are already experimenting with or developing gen AI applications in supervision, while a few are already using the new technology. These activities come with some challenges. Outdated IT infrastructure, data security and a lack of technical skills make it hard for supervisory authorities to experiment with, develop and deploy gen AI applications. These challenges may have also prevented other authorities from even starting to explore the new technology.

Gen AI applications that are currently being used for supervision purposes are limited to basic document processing. Nevertheless, experiments and development work are already going into more complex use cases, such as for knowledge management and document review. Applications that are already in use have not really been integrated into supervisory processes – ie they are not necessary to complete supervisory processes – but are only made available to supervisors for them to use at their discretion. User acceptance and inaccuracies in information provided are the most often cited challenges in integrating gen AI into supervision. These challenges will likely intensify as financial authorities deploy more complex gen AI use cases.

References

Crisanto, J C, C B Leuterio, J Prenio and J Yong (2024): "Regulating AI in the financial sector: recent developments and main challenges", *FSI Insights on policy implementation*, no 63, December.

Gambacorta, L, S Kiuhan-Vásquez, N Lauridsen and J Prenio (forthcoming): "The drivers of supotech adoption: a new database".

Prenio, J (2024): "Peering through the hype: assessing supotech tools' transition from experimentation to supervision", *FSI Insights on policy implementation*, no 58, June.

Prenio, J, A Pustelnikov and J Yeo (2024): "Building a more diverse supotech ecosystem: findings from surveys of financial authorities and supotech vendors", *FSI Briefs*, no 23, May.

Annex – List of authorities that responded to the survey⁷

Jurisdiction	Authority
Belgium	National Bank of Belgium
Brazil	Central Bank of Brazil (BCB)
Canada	Office of Superintendent of Financial Institutions (OSFI)
Denmark	Danish Financial Services Authority (D-FSA)
European Union	European Central Bank (ECB)
	European Insurance and Occupational Pensions Authority (EIOPA)
France	French Prudential Supervision and Resolution Authority (ACPR)
Finland	Finnish Financial Services Authority (FIN-FSA)
Germany	Deutsche Bundesbank (DB)
Guernsey	Guernsey Financial Services Commission (GFSC)
Hong Kong SAR	Hong Kong Monetary Authority (HKMA)
India	Reserve Bank of India (RBI)
Israel	Bank of Israel (BoI)
Italy	Bank of Italy (BdI)
Kosovo	Central Bank of Kosovo (CBK)
Malaysia	Central Bank of Malaysia (BNM)
Mauritius	Bank of Mauritius (BM)
Montenegro	Central Bank of Montenegro (CBM)
Netherlands	De Nederlandsche Bank (DNB)
New Zealand	Reserve Bank of New Zealand (RBNZ)
Oman	Central Bank of Oman (CBO)
Panama	Superintendency of Banks of Panama (SBP)
Peru	Superintendency of Banking, Insurance and Private Pension Funds (SBS)
Philippines	Bangko Sentral ng Pilipinas (BSP)
Poland	Polish Financial Supervision Authority (KNF)
Qatar	Qatar Financial Centre Regulatory Authority (QFCRA)
Rwanda	National Bank of Rwanda (BNR)
Serbia	National Bank of Serbia (NBS)
Singapore	Monetary Authority of Singapore (MAS)
South Africa	Prudential Authority, South African Reserve Bank (PA – SARB)
Spain	Bank of Spain (BdE)
Sweden	Swedish Financial Services Authority (S-FSA)
Switzerland	Swiss Financial Market Supervisory Authority (FINMA)
Thailand	Bank of Thailand (BoT)
United Arab Emirates	Dubai Financial Services Authority
United Kingdom	Financial Conduct Authority (FCA)

⁷ Six respondent authorities did not indicate their institution names in their survey responses.