

Preparing for systemic risks in the age of generative artificial intelligence

By [Thiago Fauvrelle](#), [Martin Rohár](#)

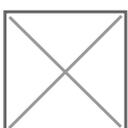
29 July 2025

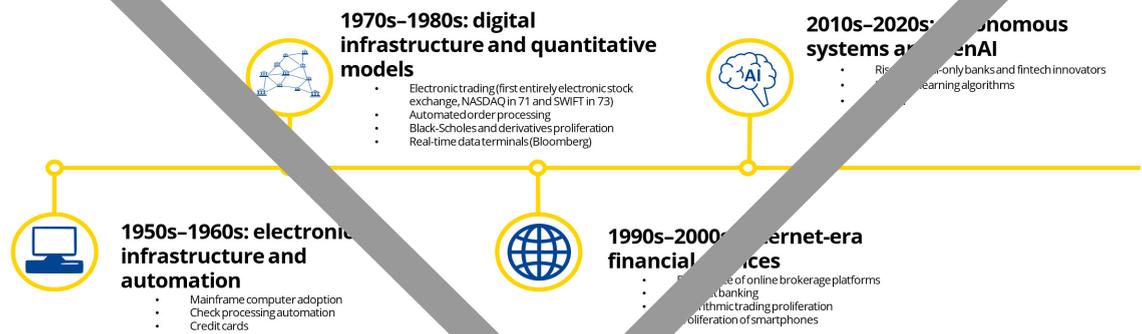
Large language models capable of generating human-like text have evolved rapidly over the past few years. With this leap in technology, it became possible to generate coherent text using artificial intelligence (AI), which, in turn, led to the proliferation of so-called generative AI (GenAI). Massive improvements in extracting insights from extensive domain-specific knowledge brought about by GenAI is widely expected to reshape parts of the economy that rely on processing large amounts of information, like the financial sector.

On a fundamental level, the financial sector operates through the continuous aggregation, analysis, and exchange of data. This intrinsic reliance on information processing has historically driven the rapid adoption of new technologies in finance, as seen in the shift from manual ledger systems to mainframe computers, open outcry to electronic trading platforms, and the telephone to the internet. Unlike previous technological developments, GenAI substantially improves both generation and synthesis of qualitative information. However, technological advancements in finance underscore a somewhat paradoxical truth: the very tools that improve the efficiency of processes, make the system more complex. Increased system complexity creates unpredictable interdependencies and could increase potential failure points.

At the ESM, we recognise that the systemic implications of widespread GenAI adoption demand vigilant monitoring and proactive risk management to maintain financial stability.

Figure 1: Adoption of technology in finance: a historical timeline





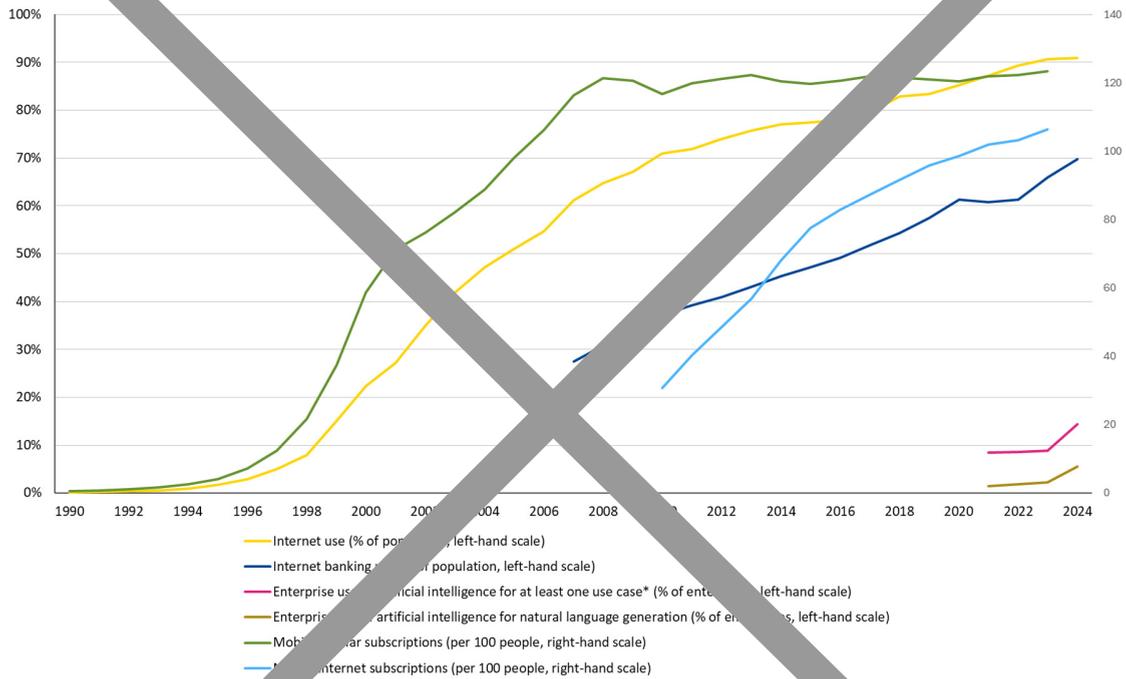
Source: ESM

Watch out for systemic risks

AI in the form of machine learning^[1] had some impact within the world of finance over the last decades. However, its adoption was often limited to specific use cases, as models were typically tailored to particular problems and their results were not easily explainable. GenAI can handle a wide range of complex tasks based on natural language input and it also offers some explainability of its outcomes via techniques such as chain-of-thought prompting^[2]. Moreover, it can absorb vast quantities of information and generate new content at unprecedented speeds. These advancements make it appealing for rapid mass market adoption, particularly by companies. (Figure 2).

Figure 2: Waves of selected technology adoptions in euro area





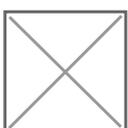
Notes: *AI use cases covered: text mining, speech recognition, natural language generation, image recognition, image processing, machine learning, AI-based software robotic process automation, and autonomous decisions based on observation of surroundings.

Source: IBM calculations based on International Telecommunication Union and Eurostat data.

Historically, the increasing velocity and volume of information have often coincided with increased systemic complexity and inherent risks. The capabilities unlocked by GenAI have the potential to accelerate the flow of information beyond expectations.

Past episodes offer cautionary insights. For instance, the 1987 stock market crash was exacerbated by portfolio insurance programmes that automatically triggered selling based on market conditions. The 1998 collapse of Long-Term Capital Management illustrated how the failure of complex quantitative models could rapidly become systemic threats. Similarly, the widespread use of mathematical models and sophisticated derivatives by the turn of the century promised better risk management but ultimately contributed to the 2008 financial crisis when these tools were poorly understood until the system was under severe stress.

So, despite its benefits, GenAI could create unpredictable interdependencies and vulnerabilities not apparent under normal market conditions – issues that need proper management to prevent cascading systemic failures.



Vibe trading and new market dynamics

GenAI has the potential to improve the accessibility of financial markets for retail investors but can also amplify risks. Two of the main obstacles for retail investors are barriers to entry and the relative complexity of financial products. GenAI-equipped platforms can help overcome both problems by providing personalised investment recommendations, simplifying complex financial data, and facilitating seamless execution of investment and trading strategies using natural language input. This way, GenAI can enable broader participation in complex trading activities using intuitive interfaces.

This democratisation improves financial inclusion. It can enhance market efficiency but also introduces new sources of volatility and systemic risk. Retail investors may not fully understand the underlying risks or limitations of the GenAI systems they are employing for investment advice. For instance, there is a potential for widespread adoption of “vibe trading” – an approach where investors describe their objectives in natural language as prompt to GenAI tuned for financial advice and subsequently leverage the technology to execute the strategies provided by GenAI without thorough critical evaluation, deferring largely to the GenAI’s implicit inference. As a result, retail investors might unknowingly adopt correlated strategies derived from similar model training data. This way, they could inadvertently mirror the coordinated behaviours as seen for instance in the 2007 GameStop frenzy, but at much faster speeds.

One of the main consequences of herd behaviour would be the amplification of liquidity risks. Traditional retail investment requires individual decision-making and manual execution, creating natural frictions that limit the speed and scale of market movements caused by retail investors. However, retail investors using similar decision making and trading algorithms may adopt correlated investment strategies. If these tools are widely used, the resulting trading volumes can exceed market absorption capacity and trigger liquidity crunches. Moreover, retail investors may be more prone to panic selling or speculative behaviour during market stress, lacking the sophisticated risk management and dynamic investment strategies, and regulatory oversight of institutional investors. These actions could amplify volatility and create feedback loops that destabilise markets. AI enabled systems can eliminate market frictions, allowing simultaneous mass participation in market movements by retail investors that could then overwhelm traditional market stabilisation mechanisms.

Unknown unknowns



Cutting-edge AI models may lead to the emergence of never-before-seen financial products. New GenAI models equipped with chain-of-thought reasoning capabilities can synthesise insights from vast datasets and identify patterns that escape human detection. As these systems become more sophisticated in analysing market data and identifying arbitrage opportunities, they may create increasingly complex financial instruments designed to exploit market inefficiencies or regulatory gaps. This, in turn, enables the creation of financial innovations and investment strategies that may outstrip the development of appropriate risk management frameworks, thus embedding risks that are not apparent to human analysts. History suggests that financial innovations often precede the regulatory and risk management structures needed to contain their systemic implications. For example, the proliferation of mortgage-backed securities and collateralised debt obligations in the early 2000s exemplifies how innovative financial instruments can create systemic vulnerabilities before their risks are widely understood or adequately regulated. Therefore, new generations of AI models have the potential to create hidden vulnerabilities that only become apparent under stress.

One particular avenue of GenAI propagation in the financial sector of potential concern is the introduction of fully autonomous AI investment systems. Current guidelines do not preclude companies from offering them as a service to retail investors. Some institutional investors are beginning to deploy AI-enabled systems capable of making investment decisions without human oversight. Traditional algorithmic trading systems operate according to predetermined rules with humans in the loop. New systems that integrate cutting-edge AI algorithms can adapt their strategies based on market conditions and learn from their outcomes without human oversight. This adaptability, while potentially beneficial for performance, makes it extremely difficult to predict how these systems will

provide limited guidance.

Continuous vigilance needed to maintain financial stability

The transformative impact of GenAI on financial markets demands an evolution in how financial stability risks are monitored, assessed, and managed. The uneven pace and scope of GenAI integration across the financial sector call for proactive surveillance mechanisms that can identify emerging vulnerabilities before they manifest as systemic threat.



At the ESM, we recognise that ensuring financial stability in the age of GenAI requires accumulating knowledge on impending AI integration in the financial sector and continuous enhancement of analytical tools. This includes critically assessing whether such tools are adequate for addressing a potentially AI-amplified crises and developing new approaches where existing frameworks prove insufficient. [Strong coordination](#) among all institutions tasked with safeguarding financial stability is essential. As transformative AI technologies become increasingly integrated into financial markets, the ESM strives to improve its capabilities to stay ahead of the curve and preserve financial stability across its Members.

Further reading

[Foucault, T., Gambacorta, L., Wang, W., Vives X. \(2025\), Artificial Intelligence in Finance. The Future of Banking 7. CEPR Press.](#)

[Gambacorta, L., Sabatini, F., Schiaffi, S. \(2025\), Artificial intelligence and relationship lending. BIS Working Paper No 1244.](#)

Acknowledgements

The authors would like to thank [Paolo Fioretti](#), [Josselin Hebert](#), [Nicoletta Mascher](#), [Golf Stauch](#) for valuable discussions and comments to this blog post, and Raquel Calero and [Karol Siskind](#) for their editorial review.

Footnotes

[1] Machine learning refers to the process by which computers can improve their ability (i.e. “learn”) to execute tasks by analysing patterns in data using statistical models, without the need for explicit instructions from humans.

[2] Chain-of-Thought prompting is a technique that enables AI models to improve the reasoning capabilities by breaking prompts down into step-by-step reasoning processes, something akin to asking the model to “think out loud” before providing an answer.

