

**Bank of England**

# Financial Stability in Focus: Artificial intelligence in the financial system

**Financial Policy Committee**

April 2025



# Financial Stability in Focus

April 2025

The primary responsibility of the Financial Policy Committee (FPC), a committee of the Bank of England, is to contribute to the Bank of England's financial stability objective. It does this primarily by identifying, monitoring and taking action to remove or reduce systemic risks, with a view to protecting and enhancing the resilience of the UK financial system. Subject to that, it supports the economic policy of His Majesty's Government, including its objectives for growth and employment.

The Financial Stability in Focus (FSiF) sets out the FPC's view on specific topics related to financial stability. It complements the Financial Stability Report, which is published twice a year, and sets out the FPC's overall view of the outlook for UK financial stability, including its assessment of the resilience of the UK financial system and the main risks to UK financial stability, and the action it is taking to remove or reduce those risks.

The Financial Policy Committee:

Andrew Bailey, Governor

Sarah Breeden, Deputy Governor responsible for financial stability

Clare Lombardelli, Deputy Governor responsible for monetary policy

Dave Ramsden, Deputy Governor responsible for markets and banking

Sam Woods, Deputy Governor responsible for prudential regulation

Nikhil Rathi, Chief Executive of the Financial Conduct Authority

Nathanaël Benjamin, Executive Director for financial stability strategy and risk

Colette Bowe

Jon Hall

Randall Kroszner

Liz Oakes

Carolyn Wilkins

Gwyneth Nurse attends as the Treasury member in a non-voting capacity.

This document, unless otherwise stated, uses data available as at 3 April 2025.

For the avoidance of doubt, the Financial Stability in Focus is not intended to satisfy the requirements of Section 9W of the Bank of England Act 1998.

---

# Contents

---

<b>Executive summary</b>	<b>3</b>
<b>1: Context: the potential benefits of AI and its growing role in the financial system</b>	<b>6</b>
<b>2: The financial stability implications of AI</b>	<b>11</b>
Greater use of AI in banks' and insurers' core financial decision making	14
Greater use of AI in financial markets	15
Operational risks in relation to AI service providers	17
Changing external cyber threat environment	18
<b>3: The FPC's approach to monitoring and mitigating risks from AI</b>	<b>21</b>
<b>Box A: AI adoption in principal trading firms and systematic hedge funds</b>	<b>26</b>
<b>Box B: AI in the Bank of England</b>	<b>28</b>

## Executive summary

---

The development and deployment of artificial intelligence (AI)<sup>[1]</sup> is likely to have a transformative impact across many sectors of the UK economy. AI has the potential to save workers time on a wide range of tasks, thus potentially boosting productivity. It can enhance firms' decision-making processes and help make products and services better and more tailored to customers' needs. At the cutting edge it can catalyse other scientific or technical breakthroughs, such as in computing or medicine. All of this has the potential to increase long-term productive economic growth.

Finance is among those sectors benefiting from this source of innovation. AI is already helping many financial institutions to automate and optimise their existing internal processes, such as code generation, as well as their interactions with customers. A likely area of development over the coming years is advanced forms of AI increasingly helping to inform firms' core financial decisions, such as credit and insurance underwriting, potentially shifting the allocation of capital. By enabling new sources of data to be used, the technology could ultimately enhance firms' offering to customers. However, in the context of the new and distinct features of advanced AI, and the rapid pace of its development, there is a high degree of uncertainty over how the technology and its use will evolve. Section 1 of this Financial Stability in Focus (FSiF) discusses this broader context to the Financial Policy Committee's (FPC's) consideration of AI.

As a macroprudential policymaker, the FPC is focused on financial stability risks, which can ultimately impact households and businesses. A stable financial system is one that has sufficient resilience to be able to facilitate and supply vital services by financial institutions, markets and market infrastructure to households and businesses in a manner that absorbs rather than amplifies shocks. Financial stability risks can arise even where risks to the safety and soundness of individual firms are well managed by microprudential authorities, for example arising as a result of the collective behaviour of firms.

Given the significant levels of uncertainty around how AI will evolve, the FPC is considering the potential macroprudential implications of more widespread, and changing, use of AI in the financial system. By doing so, the FPC can contribute to the safe adoption of the technology from the perspective of financial stability, which will support sustainable growth. In this context, the FPC is focused on the following areas:

- **Greater use of AI in banks' and insurers' core financial decision-making (bringing potential risks to systemic institutions).** While bringing potential benefits to both firms and their customers, such as greater choice and product availability, AI can introduce risks, especially in relation to models and data. Firm-level risk management and microprudential

regulation can help mitigate these risks by applying appropriate controls to the use of AI, including agentic AI (that is, systems which can take autonomous action to achieve specified goals by utilising tools, learning from feedback, and adapting to dynamic environments). But there is the potential for systemic consequences to emerge, for example if common weaknesses in widely used models cause many firms to misestimate certain risks and so misprice and misallocate credit as a result. Such common weaknesses could also lead to a loss of service provision for some households or businesses. More widely, a reliance on AI models for key decisions could lead to conduct-related risks, for example if certain decisions or processes were to be subject to legal challenge and financial redress.

- **Greater use of AI in financial markets (bringing potential risks to systemic markets).** Greater use of AI to inform trading and investment decisions could help increase market efficiency. But it could also lead market participants inadvertently to take actions collectively in such a way that reduces stability. For instance, the potential future use of more advanced AI-based trading strategies could lead to firms taking increasingly correlated positions and acting in a similar way during a stress, thereby amplifying shocks. Such market instability can then affect the availability and cost of funding for the real economy.
- **Operational risks in relation to AI service providers (bringing potential impacts on the operational delivery of vital services).** Financial institutions generally rely on providers outside the financial sector for AI-related services to develop and deploy AI (just as they do for various other IT services). Reliance on a small number of providers for a given service could lead to systemic risks in the event of disruptions to them, especially if it is not feasible to migrate rapidly to alternative providers.
- **Changing external cyber threat environment.** While AI might increase financial institutions' cyber defensive capabilities, it could also increase malicious actors' capabilities to carry out successful cyberattacks against the financial system. And financial institutions' own use of AI could create new vulnerabilities that actors could exploit.

A more indirect – but potentially significant – way in which AI could affect the financial system is through its adoption across the wider economy. For example, if it challenges established business models in certain economic sectors, this could impact some borrower firms' creditworthiness and thus increase credit risk for those lenders that have exposures to them. Section 2 of this FSiF sets out the FPC's view of the potential financial stability implications of AI in more detail.

The FPC's intends to build out its monitoring approach to enable it to track the development of AI-related risks to financial stability. The approach will need to be flexible and forward-looking given the uncertainties and potential pace of change in AI. To this end, the FPC plans – supported by the Bank, the Prudential Regulation Authority (PRA) and the Financial Conduct

---

Authority (FCA) – to make use of a blend of quantitative and qualitative information sources. These include the regular Bank and FCA [Survey on AI in UK financial services](#) (hereafter ‘the AI Survey’), the [AI Consortium](#), and targeted market and supervisory intelligence gathering. The FPC will continue to adapt and add to these tools in a flexible way as the risk environment evolves.

The effective monitoring of AI-related risks is essential to understand whether additional risk mitigations might be warranted in support of safe innovation, what they might be, and at what point they may become appropriate. The FPC will also continue to engage actively with domestic and international initiatives to monitor and mitigate AI-related risks, not least because many of these risks are likely to be cross-border in nature. Section 3 of this FSIF sets out the FPC’s planned approach to monitoring and mitigating AI-related risks in more detail.

# 1: Context: the potential benefits of AI and its growing role in the financial system

---

**AI represents a significant advance compared with previous modelling techniques, bringing transformative potential.**

Some forms of AI represent a discontinuity relative to previous modelling technologies.<sup>[2]</sup> For example, advanced AI models (including generative AI models) can be dynamic, learning automatically from new input data, meaning that their outputs can evolve over time. They can be used to produce complex outputs and make decisions autonomously. And they are trained on vast volumes of data, on a different scale to previous modelling tools. As a result, they have powerful capabilities across a wide range of use cases, bringing significant – and potentially transformative – benefits to their users. In the coming years, various parts of the UK economy, including financial services, may be reshaped as the use of this technology becomes more widespread and evolves.

**AI brings both benefits and potential risks, so is of relevance to the FPC's financial stability objective of protecting and enhancing the resilience of the UK financial system.**

While the distinct features of AI are the source of its unique benefits, they can also be additional sources of risk (as discussed by [Sarah Breeden \(2024\)](#)). For example, the complexity of some AI models – coupled with their ability to change dynamically – poses new challenges around the predictability, explainability and transparency of model outputs. And their use of very large amounts of data poses new challenges for users around ensuring the integrity of that data. The potential for market concentration in AI-related services, including vendor-provided models, is a further challenge.

The FPC fulfils its financial stability objective of contributing to protecting and enhancing the resilience of the UK financial system by identifying, monitoring and taking action to remove or reduce systemic risks to financial stability. This includes work on emerging systemic risks where there is considerable uncertainty over their potential timing and size. The FPC is considering the potential implications of AI in order to contribute to its adoption in a way that safeguards financial stability and so is conducive to sustainable growth. It is doing so in the context of a wider public discussion over the benefits and potential risks around AI, including in relation to financial stability.<sup>[3]</sup>

**The development and broad-based adoption of AI could lead to economic changes over the coming decades, increasing UK economic growth.**

As a general-purpose technology, AI can bring productivity gains to many economic sectors. For example, it can help reduce resources spent on routine administrative tasks, freeing up employees' time for higher value-added work. It can enhance institutional decision-making and help make products and services better and more tailored to customers' needs. And at the cutting edge it could catalyse other scientific or technical breakthroughs, such as in computing or medicine. All of this has the potential to support long-term productive economic growth.

As highlighted by the Government-commissioned [AI Opportunities Action Plan](#), effective and swift AI adoption has the potential to enhance the competitiveness of areas of UK economic strength, and to unlock new growth opportunities across the whole economy. As the third largest destination for AI investment globally, the UK is well placed to capitalise on these opportunities.

Subject to its objective in relation to financial stability (which is itself a vital foundation for sustainable growth), the FPC also has an objective to support the Government's economic policy. Supporting broad-based and resilient growth built on strong and secure foundations contributes to that objective. As such, when considering the implications for financial stability of emerging technologies such as AI, the FPC is also mindful of the significant economic opportunities presented by them.

**While AI adoption is currently happening at pace in many parts of the financial services sector, there is a high degree of uncertainty about its specific future impacts.**

Many AI-based analytical techniques, such as those which assist in statistical analysis, are not new and have been established for a decade or more. Meanwhile, some use cases for advanced forms of the technology, such as generative AI, are now in production. The AI market continues to change at pace, for example with the development of so-called agentic AI systems. Such systems, which are not widespread at present, can take autonomous action to achieve specified goals by utilising tools, learning from feedback, and adapting to dynamic environments. It is also possible that some applications will not in practice live up to their initial promise, while currently-unforeseen developments may have significant impacts.

Financial institutions' decisions about the pace, scale and nature of their AI adoption will depend on a complex combination of technical developments, commercial incentives and risk appetite, implying a wide range of possible future adoption scenarios. The regulatory environment will also be relevant to the pace and nature of industry adoption: the appropriate management of firm and system level risks can help create an environment in which firms are able to innovate safely and unlock the full benefits of the technology.

**AI in financial services is already helping to improve operational efficiency and effectiveness, especially through assisting employees with routine tasks.**

The 2024 AI Survey indicates that the top near-term (in the next three years) use cases for AI include optimising internal processes, enhancing customer support and combatting financial crime (Chart 1). Financial firms appear most willing to deploy AI in these types of operationally focused use cases, which are expected, according to the same survey, to be among those delivering the biggest benefits in three years' time.

**Chart 1: Current and planned use cases for AI deployment**

Percentage of firms currently using or planning to use AI, by use case (a)



Source: Bank of England.

(a) AML/CFT is 'anti-money laundering and combating the financing of terrorism'.

A key breakthrough has been the rapid development of generative AI, such as large language models that can generate natural language text output. These models are often pre-trained and provided as cloud-based services by third-party providers. Such models are being used to help streamline various internal functions within financial institutions, such as code generation and information search and retrieval, helping improve productivity. And AI-based analytics are being used to enhance customer interactions, for example helping payment firms better predict a customer's preferred payment option.

While there is considerable uncertainty about the longer-term economic impacts of advanced AI, some analysis suggests that these could be very significant, with one study estimating that, over the next 15 years, generative AI could bring productivity gains of up to 30% to the banking and insurance sectors, and to firms operating in capital markets.<sup>[4]</sup> AI can also help public authorities to achieve their objectives more efficiently and effectively. Box B discusses how the Bank of England is adopting AI in its work, including in support of the FPC's financial stability objective.

**AI can also assist in financial institutions' core business decisions, such as lending, insurance underwriting and trading.**

The broad capabilities of AI point to its likely spread across use cases over time, including in areas central to financial institutions' business models. Greater use of AI could ultimately help firms to enhance their offering to customers. For example, the potential ability for some lenders, including non-bank lenders, to leverage a wider range of structured and unstructured data could, in principle, widen choice and access to finance for creditworthy companies, including small and medium-sized enterprises.

Lending decisions are at the heart of banks' financial risk management. Supervisory intelligence suggests that while in aggregate the use of AI in credit risk management is still in its infancy, some firms are using AI-based techniques (such as established gradient boosting decisions tree models) at various stages of the lending process. This includes in the pre-screening, application scoring, pricing and provisioning steps, and applies across lending classes. Among insurers, AI-based models are currently widely used to support pricing and underwriting decisions. The Organisation for Economic Co-operation and Development (OECD) has highlighted how the use of telematics – which can generate large quantities of data to feed into AI-based models – may in the future increasingly be integrated into insurance products, beyond its current widespread use in motor insurance.<sup>[5]</sup> Such innovations have the potential to enable more tailored insurance products or pricing, as well as potentially helping with insurers' risk management.

Firms participating in financial markets seek to make the best possible use of available data to optimise their trading strategies. Among institutions undertaking algorithmic (rules-based) trading in highly liquid markets, established AI techniques (such as decision trees) are already

---

deployed to help refine the predictive power of models that feed into their trading strategies. And there is active innovation in this space, with a recent International Monetary Fund (IMF) report highlighting that over half of all patents filed by high frequency or algorithmic trading firms now relate to AI.<sup>[6]</sup> And while autonomous AI-based trading models do not yet appear to be in widespread production, it is plausible that in the future such approaches will be employed (Box A).

Investment managers similarly stand to benefit from the rapid development of AI techniques. For example, the same IMF report identified the use of generative AI by investment managers to help them better use alternative data sets, such as social media content, to uncover previously unknown relationships between economic or financial variables, and hence to generate new investment strategies. As well as benefiting end-investors, the exploitation of novel sources of data could help increase market efficiency, with new sources of information incorporated into pricing faster and more accurately than was previously possible.

## 2: The financial stability implications of AI

---

### **There are various ways in which AI-related developments might impact financial stability.**

This section explores several key ways in which AI might interact with vulnerabilities at the firm and system level, and could – especially in the absence of sufficient mitigations – lead to financial stability risks. These risks to financial stability could transmit to the real economy via their effect on systemically important institutions, systemically important markets or by affecting the operational delivery of vital services. Specifically, it explores the following four AI-related areas of focus for the FPC:

- Greater use of AI in banks' and insurers' core financial decision-making.
- Greater use of AI in financial markets.
- Operational risks in relation to AI service providers.
- Changing external cyber threat environment.

The distinct features of advanced AI models relative to other modelling technologies are relevant to each of these areas, and feed into potential risks to financial firms and the financial system. For example, the potential for dynamism in complex AI models (updating as new data is available) and a lack of predictability and explainability of their outputs could, other things equal, make it harder for firms to manage risks related to their use. In the context of widespread use of vendor-provided AI models, around half of the respondents to the 2024 AI Survey report having only a 'partial understanding' of the AI technologies they use. More generally, the quality of AI model output relies on the quality of the input data. The use of very large-scale training data can add to existing challenges for firms to ensure that data used is relevant, of sufficient quality, and does not introduce bias.

The potential for market concentration in AI-related services is also of relevance. In the generative AI market in particular there are various factors that could increase concentration over time.<sup>[7]</sup> These include the cost and complexity of the models and vertical integration of parts of the 'AI stack'. However, there are also various factors that could have the opposite effect on market concentration, including the widespread availability of open-source models.

Given these considerations, a scenario in which AI models are increasingly deployed in an autonomous manner (as opposed to being used largely as an assistive tool) could potentially pose significant additional risks to financial stability in the future.<sup>[8]</sup>

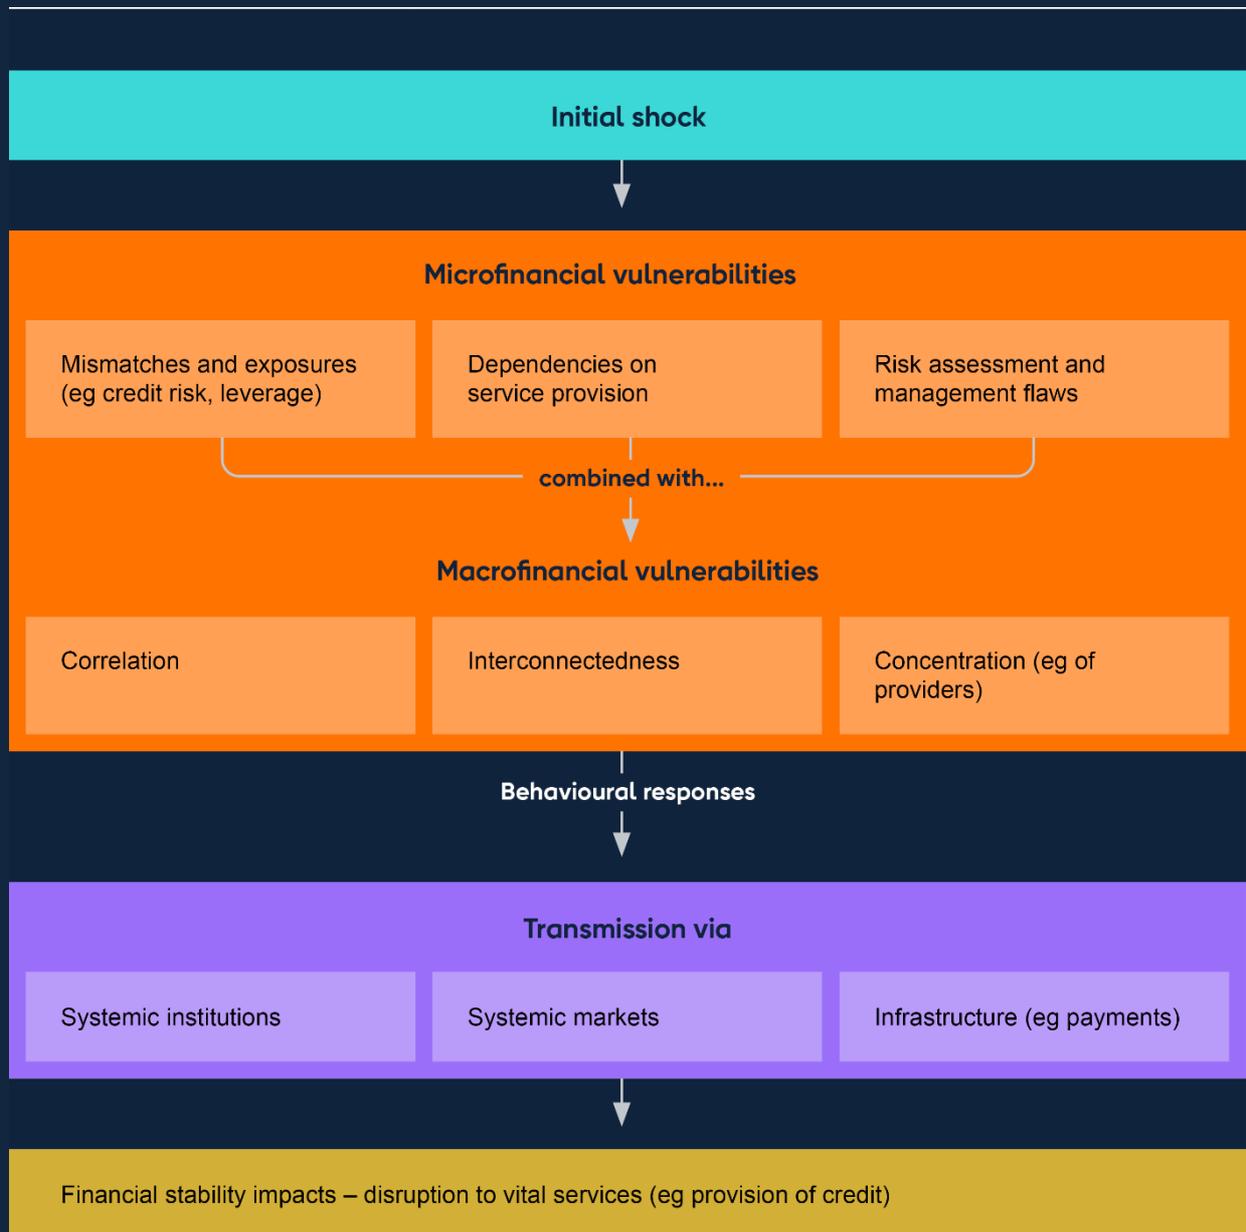
This section does not seek to present a comprehensive overview of all possible AI-related risks to financial stability, which would also encompass potential effects on business models and market structures. These could be significant, especially over the longer term. For example, if AI challenges established business models in certain economic sectors this could, in principle, impact some borrower firms' creditworthiness and thus credit risk for those lenders that are exposed to them. AI could also increase the relative footprint of non-bank financial institutions (NBFIs) in certain markets, for example as a result of a greater use of algorithmic trading. Given the high level of uncertainty around the future trajectory of AI, it is challenging at present to assess all such potential longer-term impacts.

### **| Risks can stem from vulnerabilities at both the firm and system level.**

The FPC identifies and assesses risks by considering vulnerabilities arising at both the institution level (microfinancial) and system level (macrofinancial) and the transmission channels through which they can impact financial stability (Figure 1). Actions by both financial institutions and public authorities can help build resilience to systemic risks. Microfinancial vulnerabilities often relate to risks that can impact individual firms' safety and soundness, such as model risk, and microprudential regulation helps to mitigate such risks.

But even where risks are well managed from the perspective of individual firms, macrofinancial vulnerabilities can mean that the collective behaviour of firms in response to a shock can have implications for financial stability. In particular, this may be the case when firms do not have sufficient information or incentives to take account of system level outcomes in deciding their actions – in other words when they are outcome agnostic from a system perspective. Such risks to the system are the focus of the FPC's macroprudential work.

Figure 1: How shocks can interact with vulnerabilities to affect financial stability (a) (b)



Source: Bank of England.

- (a) 'Microfinancial vulnerabilities' are entity-level vulnerabilities that determine their financial and operational susceptibility to shocks.
- (b) 'Macrofinancial vulnerabilities' are features of the system that help determine the impact of shocks on the system and real economy, given microfinancial vulnerabilities. More precisely, they are features that – conditional on the existence of microfinancial vulnerabilities – can increase the impact of shocks across the system and real economy.

The FPC’s analysis of AI-related risks will continue to be updated as the external risk environment evolves and as more information becomes available, including through its ongoing monitoring work (Section 3).

## Greater use of AI in banks' and insurers' core financial decision making

**While bringing various potential benefits to both firms and customers, AI can introduce new risks for individual firms, especially in relation to data and models.**

As described in Section 1, it is likely that banks and insurers will increasingly integrate the use of AI into their core business decisions around the provision of credit and insurance, respectively. Doing so could help enhance their product and service offering to customers, and it could also improve the accuracy of their financial risk management. At the same time, it is important to be alert to potential risks that could arise from the deployment of AI in business functions that have a direct impact on the financial position of the firm and outcomes for customers. For example, the lack of explainability and potential autonomy of advanced AI models could – if deployed without appropriate testing, governance and risk controls – lead to a level of financial risk-taking that is not properly understood at the time.

**Microprudential regulation can help mitigate risks from AI...**

A range of existing microprudential principles, regulation and guidance is of relevance to firm level risks, notably measures in relation to model risk management, data and governance, and conduct.<sup>[9]</sup> And the Senior Managers and Certification Regime (SM&CR) is a supervisory tool that can be used to ensure appropriate individual accountability for conduct and competence in relation to these issues. In the context of the changing risk landscape around AI, a number of aspects where existing regulatory regimes might need to evolve were highlighted in responses to the FCA, Bank and PRA discussion paper [DP5/22 – Artificial Intelligence and Machine Learning](#) (summarised in feedback statement [FS2/23 – Artificial Intelligence and Machine Learning](#)).

It will be important to ensure that existing regulatory frameworks, alongside firm level controls, mitigate microfinancial risks from AI sufficiently, especially as AI models are increasingly used in agent functions. The FPC will continue to engage other regulatory authorities on relevant frameworks, to help inform its assessment and monitoring of systemic risks from AI (Section 3).

**... but it is also important to consider system level implications.**

At the system level, common weaknesses in model and data risk management across firms would represent a macrofinancial vulnerability. For example, in the event that large numbers of firms rely on the same open-source model components or data libraries, a significant unknown error or bias could cause many firms to misestimate certain risks and so misprice and misallocate credit as a result. The eventual crystallisation of such a weakness could generate losses for a number of systemic firms, leading to a tightening of credit supply to the real economy, or broader financial contagion through a loss of confidence. This type of scenario was seen in the 2008 Global Financial Crisis, where a debt bubble was partly fuelled

by the collective mispricing of risk (as transformed by innovations around securitisation). More widely, a high level of reliance on AI models for key risk management decisions, could, in principle, impact other areas of firms' resilience, such as liquidity preparedness.

Under a scenario in which core decisions on the availability and pricing of services are underpinned by AI models, biased or wrongly calibrated data or models could directly affect outcomes for consumers, such as their access to products. This could in turn give rise to conduct-related risks, for example if certain decisions or processes were to be subject to legal challenge and financial redress. This could be amplified by practical issues related to establishing who is ultimately liable for decisions made by AI models.

## Greater use of AI in financial markets

**AI could be used to inform more trading and investment decisions, and that may be associated with greater market efficiency but also require appropriate risk management.**

Market participants appear likely to integrate more advanced AI-based analysis into their core trading and investment activities, although the speed and scope of AI deployment is uncertain and could vary significantly across institutions and asset classes. In particular, institutions undertaking algorithmic trading already widely use established AI techniques (such as decision trees) to calibrate their algorithms, with scope for further innovation in this space (Box A). And some investment managers are turning to AI to help generate profitable insights.

Greater use of AI by market participants in their core business processes could help increase market efficiency (for example through the faster incorporation of new information), while also being beneficial for end-investors through increased returns. At the same time, the deployment of increasingly complex AI models in this way raises various potential firm level risk management challenges, in common with those already discussed under the previous section. Unknown data or model flaws might mean that a company's exposures turn out to have been incorrectly measured or interpreted, leading to it having insufficient financial resilience to market stress events. And it may be particularly challenging for AI models to respond to extreme events and situations of radical uncertainty, such as historically unprecedented shocks.

**AI-driven trading and investment strategies could increase the tendency for market participants to take correlated positions.**

Greater use of AI-driven trading strategies could lead to various potential outcomes for markets and the practical implications are uncertain. From a systemic risk perspective, the potential for AI-based participants to take increasingly correlated positions is an important consideration. This could be driven by the widespread use of a small number of open-source or vendor-provided models or underlying data sets, or a more general convergence on very

similar model designs across the market. Herding and market concentration was the top risk cited in recent IMF outreach when stakeholders were asked about risks that could result from wider adoption of generative AI in capital markets.[10]

As explored by [Jonathan Hall \(2024\)](#), a potential future market with widespread use of autonomous AI-based trading might be more informationally efficient than a market shaped by human traders, but it could also be less resilient to shocks. For example, increasingly correlated positioning and strategies could exacerbate the impact of fire-sales in response to a stress event (where firms could be forced to unwind leveraged positions). The potential for this type of correlated deleveraging was explored in the [system-wide exploratory scenario](#). It arises, in part, because individual institutions may not factor in the collective impact of their actions on the market. Systemic markets, such as core bond markets, are central to the flow of finance to the real economy. Their effective functioning is therefore an important aspect of financial stability.

There are also ways in which the greater use of AI could, in principle, contribute to the improvement of market resilience. For instance, AI could enhance risk management by enabling better use of available data, meaning that the sort of fire-sale scenario described above – where leveraged firms are caught out by price moves – becomes less likely or has less of an impact. And, as noted by the Financial Stability Board (FSB), the ability of investment managers to offer increasingly customised options to their clients might have the effect, other things equal, of reducing market correlations.[11]

**Advanced AI models could rationally exploit profit-making opportunities in a destabilising way or engage in other adverse behaviours.**

Under a scenario of advanced AI trading models being deployed to act with more autonomy, these models might identify and exploit weaknesses in the trading strategies of other firms in a way that triggers or amplifies price movements. For example, models might learn that stress events increase their opportunity to make profit and so take actions actively to increase the likelihood of such events.

Another source of risk is the potential for such models to facilitate collusion or other forms of market manipulation. Given the ability of some AI models to learn dynamically in multi-agent environments, and challenges around the explainability of model outputs, such adverse behaviours might emerge without the human manager's intention or awareness.

**Existing market monitoring and oversight measures are relevant to AI risks, and the FPC will continue to follow closely the implications of AI for systemic markets.**

There are market conduct regulations to guard against market manipulation, alongside the SM&CR to help ensure appropriate individual accountability. And risks around herding in markets, potentially leading to procyclical fire-sales, are not new. The FPC already considers

the resilience of systemic markets to such sources of disruption, including through assessment tools such as the system-wide exploratory scenario. Internationally, the FSB is consulting on enhancing public disclosure on aggregate market positioning and liquidity.<sup>[12]</sup>

The rapid pace of change in AI technology could lead to correspondingly fast and significant shifts in risks to systemic markets. As such, AI-related risks in this area merit an ongoing focus from a macroprudential perspective, and the FPC intends to monitor relevant developments closely (Section 3).

## Operational risks in relation to AI service providers

**In order to capitalise on the productivity benefits of AI, financial institutions generally rely on service providers outside the financial sector.**

For many AI implementations, financial institutions rely on vendor-provided AI models. This is particularly so for very complex and powerful models, such as the most recent large language models, where significant scale is required to motivate the large amounts of capital investment needed for their development. As discussed in Section 1, these models are used in various ways to increase companies' productivity, such as assisting with code generation. For other use cases, financial institutions build AI models in-house. But even so, they may rely on cloud computing to develop and operate these models, and on external data aggregators to obtain the large data sets on which the models are trained.

**Growing concentration in the supply of AI-related services could increase risks to the financial system.**

Evidence from the AI Survey supports the view that third-party exposure will continue to increase as the complexity of models increases and outsourcing costs decrease. A further increase in interconnectedness between nodes in the financial system driven by AI has the potential to heighten existing vulnerabilities in this regard. A reliance on a small number of providers for a given service could also generate systemic risks in the event of disruptions to them, especially if it is not feasible to migrate rapidly to alternative providers.

A severe operational disruption related to external service providers could, if its impact on the financial system is severe, lead to financial stability issues. For example, under a scenario in which customer-facing functions have become heavily reliant on vendor-provided AI models, a widespread outage of one or several key models could leave many firms unable to deliver vital services such as time-critical payments. The potential for disruption from such operational risks has been underscored by several temporary outages of important banking and payment services (for example those triggered by the July 2024 worldwide IT outage caused by a flawed update distributed by the cyber security technology firm CrowdStrike).

**This highlights the importance of building and maintaining operational resilience, an existing area of focus for the FPC.**

In March 2024, the FPC set out its [macroprudential approach to operational resilience](#). The FPC noted that firm level operational resilience provides the essential foundation for operational resilience across the system. Firms can mitigate operational risks through effective control frameworks and investment in operational resilience, and both microprudential and macroprudential policies are in place to help manage certain existing risks in this space.<sup>[13]</sup>

Public-private sector collaboration (between the Bank, other authorities and a range of financial and non-financial firms) is supporting the development of a 'shared responsibility model' for AI. The output will be guidance on a structure for managing implementation risks, such as whether the third-party provider or the client firm is responsible for managing the data within different kinds of AI deployment. This should help minimise the potential for divergences in approach leading to firm or sector level operational impacts.

Additionally, the [Financial Services and Markets Act 2023](#) established a new regulatory regime for critical third parties, in response to the FPC's 2021 recommendation that additional policy measures were likely to be needed to mitigate the financial stability risks stemming from concentration in the provision of services to UK firms and FMIs. The Bank, the PRA and the FCA jointly published rules for the new regime in November 2024 ([supervisory statement 6/24](#)). Certain third parties providing data and AI models could also emerge as potential future critical third parties as a result of increasing use of them by the financial sector. There are ways in which AI services might differ from other types of third-party usage, including the complexity of the most computationally powerful foundation models, and also potential challenges around identifying specialised and niche providers.

Given the rapid evolution in how financial institutions are interacting with the AI market, and in the structure of that market, the FPC will monitor developments in this space closely (Section 3).

## Changing external cyber threat environment

**AI could be a new tool for malicious actors that already pose risks to financial companies...**

Cyberattacks are a significant source of risk faced by many financial firms, and in the Bank's most recent [Systemic Risk Survey](#), they remain near the top of the list of the perceived key sources of risk to financial system. The FPC has [previously noted](#) that higher geopolitical tensions create an environment of heightened risk of cyberattacks. And in the 2024 AI Survey, cybersecurity came near the top of perceived current AI-related risks, and respondents expected this risk to grow over the next three years.

The use of AI by threat actors could increase their capability to carry out successful cyberattacks against the financial system, with potentially greater sophistication and scale than was previously possible.<sup>[14]</sup> Financial institutions' own use of AI could also open up new opportunities for malicious actors to exploit, for example via any vulnerabilities around the software or hardware of third-party providers. The model development stage could also be a potential target, for example via the malicious manipulation of model training data (so-called data poisoning). Longer term, the potential for cyberattackers to combine AI with possible future developments in quantum computing will need to be monitored as both technologies evolve over the coming years.

In addition to cyberattackers seeking to cause disruption to the financial system, AI might also increase the capabilities or opportunities of other types of malicious actor. For example, those engaged in illicit financing (money laundering or terrorism financing) could seek to use AI models to circumvent institutional controls. And the use of public or customer-facing AI models by financial institutions creates new risks such as 'prompt injection', whereby attackers seek to manipulate models to extract confidential information. Further, cyberattackers perpetrating fraud schemes (against financial institutions directly or against retail customers) could be rendered more effective and harder to detect as a result of generative AI models. For example, the capability of these models to produce so-called 'deepfakes', as well as highly personalised text, could increase the ability of those intent on committing fraud to manipulate employees or retail customers.

### **| ... exacerbating the risks they pose to the wider financial system.**

As well as posing material risks to individual institutions, AI-related cyberattacks could have systemic implications. For example, the widespread deployment of common AI models with shared cyber vulnerabilities across systemic firms would represent a system-wide vulnerability. This might increase the impact of large-scale cyberattacks, which could spread to other parts of the financial system through operational contagion or a general loss of confidence. Financial stability might ultimately be affected, if, for example, systemic markets or the operational delivery of vital services were to be materially disrupted as a result. While recent ransomware attacks at several financial firms and third-party providers did not impact financial stability, they showed how such incidents have the potential to amplify risks across the financial system, as disruption at one firm can cause disruption at others.

A more general increase in cyberthreats could have impacts on the real economy and so indirectly on financial institutions. For example, AI-based disinformation tools such as deepfakes could be used to exacerbate existing geopolitical tensions, increasing economic uncertainty.

### **| However, AI could improve firms' ability to combat threat actors.**

The use of AI by financial institutions could improve their ability to combat malicious actors. For example, it could assist with the detection of cyberthreats by improving the automated identification of malware or illicit finance activity. Indeed, respondents to the 2024 AI Survey expected the benefits of AI for cybersecurity and anti-money laundering to grow significantly over the next three years.

The impact of AI in this area is therefore bi-directional, raising the prospect of a technological arms race between financial companies and malicious actors, and making the overall impact of AI from a financial stability perspective uncertain.

**Public-private sector collaboration will be important for helping address AI-related cyberthreats.**

UK authorities and industry have established approaches to analysing and mitigating cyber risks. This includes public-private sector collaboration through the Cross Market Operational Resilience Group (CMORG).

In 2024, CMORG established an AI Taskforce to identify and mitigate potential emerging operational risks to the sector through the widescale adoption of AI. The Taskforce is developing scenarios exploring how malicious actors could utilise generative AI to enhance their ability to conduct attacks against individual financial services firms and the sector more widely. This includes consideration of the way in which generative AI could be used to circumvent established security and authentication controls at scale. The scenarios developed through this work will inform sector-wide collaboration on the development of proactive mitigation measures, while also informing firm and sector level exercises.

Tracking the evolution of AI-related cyber risks will be an important part of the FPC's approach to monitoring AI (Section 3).

### 3: The FPC's approach to monitoring and mitigating risks from AI

---

#### The FPC has a range of information sources to help it track AI-related developments in the financial system...

There are various sources of quantitative and qualitative information which the FPC is using to build up its monitoring of AI-related implications for financial stability. This includes drawing on work by the Bank, the PRA and the FCA and using existing mechanisms, such as supervisory intelligence gathering, to focus on AI-related developments. In particular:

- **The survey on AI in UK financial services.** The regular Bank and FCA AI Survey provides an overview of the current and expected future extent of AI usage across sectors, including respondents' perspectives on the associated benefits and risks.
- **The AI Consortium.** The [Artificial Intelligence Consortium](#) is being established to provide a platform for public-private engagement to gather input from stakeholders on the capabilities, development, deployment and use of AI in UK financial services.
- **Market intelligence.** Discussions with financial market participants provide detailed and timely insights into industry trends, including from those market participants most advanced with the adoption of the technology.
- **Supervisory intelligence.** Supervisory engagement with firms can provide a detailed insight into the development and use of AI, including across PRA and FCA regulated firms and Bank regulated FMIs.
- **Regulatory and commercial data sources.** As detailed in the [FPC's approach to assessing risks in market-based finance](#), the FPC routinely draws on a range of regulatory data and commercial data sources, including financial market data, to help it track risks to financial stability.

#### ...and the FPC will continue to adapt and add to these tools in a flexible way as the risk environment evolves.

Given the pace of technological change, rapid and unforeseen shifts in the implications of AI for financial stability are plausible. As such, the FPC's approach to monitoring AI risks will need to be flexible and forward-looking. And there are features of AI that make it particularly challenging to monitor. AI is used as an input by firms across a number of internal processes and the impacts of this may not be immediately apparent in the sort of information the FPC would typically look at to assess other types of risks, such as market data or metrics on firms' resilience.

---

In the next iteration of the AI Survey, the FPC intends to work with the Bank and the FCA to increase responses from currently underrepresented sectors, and to ensure it continues to provide relevant insights into potential financial stability risks as the risk environment changes. More generally, the FPC will consider the future need for any additional sources of information or data to better understand specific risks. For example, the system-wide implications of AI-related developments might be explored via future system-wide exercises.

**| This will enable the FPC to track developments in AI use cases and implementation...**

**Table A: Information sources of relevance to monitoring AI-related systemic risks**

Potential sources of systemic risk	Examples of current information sources	Potential future information sources (depending on how the risk environment evolves)
<b>Greater use of AI in banks' and insurers' core financial decision-making</b>	<ul style="list-style-type: none"> <li>• AI Survey evidence on firms' use of AI as key input or driver of decisions (eg use of AI in underwriting).</li> <li>• Supervisory intelligence gathering.</li> <li>• Cross-border and cross-authority intelligence sharing.</li> </ul>	<ul style="list-style-type: none"> <li>• Adapted AI Survey.</li> <li>• AI-related incident reporting.</li> <li>• Structured AI-related market intelligence gathering.</li> <li>• Increased thematic supervisory activity on AI-related risks (eg use of AI models in various types of lending).</li> <li>• Potential future system-wide exercises to explore AI-related risks (such exercises could themselves use AI, including to help better understand existing risks in the financial system).</li> </ul>
<b>Greater use of AI in financial markets</b>	<ul style="list-style-type: none"> <li>• AI Survey evidence on firms' use of AI to inform key trading decisions.</li> <li>• Market and supervisory intelligence gathering.</li> <li>• Market data reporting (eg volatility, positioning trends).</li> <li>• External research and commentary.</li> </ul>	<ul style="list-style-type: none"> <li>• Future operational resilience stress testing focusing on AI enabled threats.</li> </ul>
<b>Operational risks in relation to AI service providers</b>	<ul style="list-style-type: none"> <li>• AI Survey evidence on third-party implementation.</li> <li>• Market share metrics for common AI service providers.</li> <li>• Supervisory intelligence.</li> </ul>	
<b>Changing external cyber threat environment</b>	<ul style="list-style-type: none"> <li>• CMORG AI Taskforce on operational risks to the sector from the adoption of AI.</li> <li>• Supervisory intelligence.</li> <li>• Participation in G7 Cyber Experts Group on cyber risk.</li> </ul>	

...and to follow developments that could lead it to identify additional emerging risks.

The FPC will also continue to monitor the development of AI-related technologies and their capabilities beyond the financial system. Such developments can have direct implications for some of the potential sources of risks to financial stability. In this context, the FPC will engage with AI-related expertise from across other public bodies, such as the [AI Security Institute](#), including on cyber-related risks, as well as drawing more widely on academic and industry-led research on the development of AI.

**Stakeholder engagement will help ensure the FPC’s monitoring approach is well-targeted and incorporates relevant external sources of information.**

The Bank will gather feedback on its planned monitoring approach through various stakeholder engagement channels, including supervisory engagement, AI Survey responses and the AI Consortium.

**The FPC supports international initiatives to monitor and mitigate AI risks.**

The global financial system is highly interconnected, meaning that risks arising in one part of it can quickly have implications elsewhere. As such, global developments at the intersection of financial services and AI (which is also a highly globalised market) could have implications for UK financial stability, and understanding them is an important aspect of the FPC’s monitoring approach.

The Bank, along with the PRA and the FCA, is actively engaged with various strands of international work to monitor the adoption of AI across jurisdictions, and to assess the associated benefits and potential financial stability risks, as well as any resulting gaps in the regulatory framework. This includes recent and planned work by the FSB on AI and financial stability (notably the [2024 report](#), which was informed by outreach with both member jurisdictions and industry), as well as ongoing work in the G7 and the G20. In addition, the [IMF](#) and [IOSCO](#) have both recently undertaken work on the use of AI in capital markets, contributing to a wider understanding of the potential risks to financial stability globally.

**Monitoring will allow the FPC to understand if any systemic risks develop and to ensure that any risk mitigations are calibrated appropriately to support the safe adoption of AI.**

The effective monitoring of AI-related risks is essential to understand whether additional risk mitigations might be warranted in support of safe innovation, what they might be, and at what point they may become appropriate. In the near term, a key area of focus for the Bank will be working with industry, including through the AI Consortium, to understand the changing ways in which AI is being deployed, and to identify and share good practice for managing AI-related risks. It will also be important to understand where AI adoption may have led to incidents (or ‘near misses’).

The Bank will also be mindful of the potential need for regulators to evolve existing guidance and regulation to support the safe adoption of AI across the industry. In principle, even where microprudential measures are sufficient to mitigate risks to the safety and soundness of individual firms, macrofinancial vulnerabilities (such as those discussed in Section 2) could indicate that further macroprudential measures are needed to safeguard the financial system as a whole. For example, should AI-related developments significantly increase correlations or procyclical decision-making in markets, then existing FPC work to mitigate risks in market-based finance, such as work on NBF1 leverage, may need to be adjusted. And potential changes to AI market structure that lead to a greater reliance on common models or providers may require different forms of oversight or response.

## Box A: AI adoption in principal trading firms and systematic hedge funds

Firms active in capital markets are already users of AI, as recently explored by both the [IMF](#) and [IOSCO](#). According to the IMF, a rise in the use of AI for trading and investment and a higher degree of autonomy of AI-based decisions is expected in the future. Informed by recent engagement with market contacts, this box discusses the use of AI by principal trading firms (PTFs), including high frequency trading firms, and particular hedge funds using systematic strategies. It covers the adoption of a range of AI models, including more advanced ‘black box’ models such as neural networks.

### **A range of quantitative techniques – including AI models – are being used in some markets.**

PTFs and hedge funds have used quantitative techniques for several decades. Over time, these have evolved from well-known statistical techniques such as generalised linear models to AI models such as tree-based models or neural networks. These models tend to be developed in-house.

These quantitative techniques are well-suited to traditional statistical prediction problems where there is a clear goal – such as predicting the likely future price of a security – and an abundance of high-quality historical data on which to train models.<sup>[15]</sup> Equities, foreign exchange, and exchange-traded derivatives markets are examples of markets with ample high-quality historical data. Many trading strategies, broadly defined as a view on whether firms should buy, sell or hold specific financial assets, are based on some form of a price prediction model, with additional layers such as risk management and time horizon preferences, which add constraints.

The current level of experimentation by firms suggests there is generally no single technique that outperforms another, even in comparison to legacy strategies, on a range of tasks to which they are applied. However, over the longer term, there may be greater adoption of neural networks by these firms. Currently, where such techniques are deployed, they are implemented largely as rules-based systems with human-in-the-loop.<sup>[16]</sup>

### **Autonomous trading strategies are a future possibility, but there are current challenges to their development and deployment...**

Firms are experimenting with using more complex models, such as neural networks, to develop trading strategies that can operate autonomously, with minimal human involvement. A challenge is that there are many scenarios and factors to consider

when deciding whether to buy or sell an asset. These include determining whether an asset is over or undervalued, when to buy or sell, how much to buy or sell, adapting to how other market participants are behaving in real-time. A firm's risk tolerance must also be considered, particularly in relation to diversifying its risk as part of a portfolio, in addition to legal or regulatory requirements. As a result, these models require significantly higher volumes of high-quality training data than other types of models to learn the 'optimal' solution in specific scenarios.

#### **| ...including in relation to data...**

In many markets, however, there can be a scarcity of high-quality data. This scarcity of data can lead neural networks to perform in unexpected ways, particularly when exposed to market conditions unlike those on which they have been trained, such as in different volatility regimes or market paradigms.

#### **| ...and unpredictable behaviour.**

Deploying these models in live environments in an unconstrained way could result in them taking positions in the market that expose the firm to substantial losses. Risk management of these positions is made more challenging by the lack of interpretability of neural networks, as actions may be unpredictable and the reasons for the positions may not be well understood by human risk managers at the firm. In addition, models with sufficient autonomy could act in ways that are detrimental to the overall stability or integrity of markets, for example by ignoring regulatory or legal guardrails such as market abuse regulations. Human managers would also need to manage such regulatory risks.

## Box B: AI in the Bank of England

This box sets out how the Bank of England is approaching the adoption of AI in its operations.

### **The Bank is embracing the transformative potential of AI to enhance its internal operations...**

Data are at the heart of the AI transformation. Having data relevant to the question or task and being able to easily find, access and interpret them, supported by high-quality meta-data is a critical foundation for all models. The Bank is accelerating efforts to ensure decision-making is informed by the best available data, analysis and intelligence and recently published a refreshed [data and analytics strategy](#). Setting up a new Enterprise Data Platform on the Cloud is core to this.

The latest advancements in AI, particularly generative AI, introduce fundamental changes to how data is used. This includes a significant increase in the size and complexity of the models the Bank uses; the inclusion of all data types in training and real-time applications; and the possibility for all staff to interact with AI solutions through natural language. These changes are placing additional requirements on the Bank's enterprise data platform, broadening the scope of data management and governance, and adding a new dimension to staff training. Strengthening these foundations is critical to enable the Bank to harness the potential of AI to help support its mission.

In common with approaches taken across industry, many internal applications of generative AI focus on supporting the personal productivity of staff across the Bank via off the shelf AI assistants. Internal trials of some of these tools have shown significant productivity benefits, particularly in tasks like document summarisation, meeting notes, and code generation. These tools are now being rolled out across the Bank.

### **...including those which directly support the delivery of its monetary and financial stability objectives.**

Other internal applications of generative AI require building more bespoke solutions. For example, work is underway to enable the PRA to use the latest cloud AI technology to gain supervisory insights into large volumes of unstructured data (text-based documents). Another example resides in the Bank's Agency network where a generative AI solution is being built to support the Bank's Agents when extracting insights collected at company visits. The tools for both these use cases will have wider applicability across the Bank.

In the Bank's financial stability work, given the availability and importance of quantitative data, established AI methods will continue to play an important role, for example in analysing regulatory or market data. Generative AI methods can also support the Bank's work in this space, for example with the production of synthetic data or through using large language models to undertake tasks such as summarising external reports.

**The application of AI within the Bank will need policies and a governance framework to be safe and effective.**

The Bank is using seven principles to guide its approach so that AI solutions are 'TRUSTED' ([James Benford \(2024\)](#)). They must be **Targeted**, **Reliable**, and **Secure**, clearly **Understood** and supported with **Ethical** guidance. They must also be stress-**Tested** to understand how models might behave under a range of scenarios and **Durable** so that the benefits of new AI solutions exceed the costs of acquiring or developing and maintaining them.

The Bank has set up an Artificial Intelligence and Data Ethics Governance Committee that is accountable for the internal use and build of AI solutions, including the resourcing strategy. The approach puts responsibilities on teams to make sure their use of AI tools is in line with policies.

The Bank's work to use AI internally in a safe and effective way can both help inform – and be informed by – the FPC's ongoing work to identify, monitor and mitigate AI-related risks to financial stability.

1. In this FSiF, 'AI' refers to the simulation of human intelligence by machines, including the use of computer systems, which have the ability to perform tasks that demonstrate learning, decision making, problem solving, and other tasks which previously required human intelligence (in line with the definition of 'AI' used in the Bank and FCA [Survey on AI in UK financial services](#)). The term encompasses established AI techniques (sometimes referred to as 'machine learning') as well as more advanced forms of AI, such as generative AI.
2. In this FSiF, 'AI' refers to the simulation of human intelligence by machines, including the use of computer systems, which have the ability to perform tasks that demonstrate learning, decision-making, problem solving, and other tasks which previously required human intelligence (in line with the definition of 'AI' used in the Bank and FCA [Survey on AI in UK financial services](#)). The term encompasses established AI techniques (sometimes referred to as 'machine learning') as well as more advanced forms of AI, such as generative AI.
3. For example, see Gensler, G and Bailey, L (2020), [Deep Learning and Financial Stability](#), and Danielsson, J et al (2021), [Artificial intelligence and systemic risk](#).
4. Accenture (2024), [Generating growth: How generative AI can power the UK's reinvention](#), Figure 4, page 14.
5. OECD (2020), [The Impact of Big Data and Artificial Intelligence \(AI\) in the Insurance Sector](#), Box 1, Page 12.

6. IMF (2024), [Global Financial Stability Report](#), Chapter 3, Figure 3.6.
7. For a discussion of market structure dynamics across the AI stack, see BIS (2025), [The AI Supply chain](#).
8. The latest Bank and FCA [Survey on AI in UK financial services](#) shows 55% of respondents' AI use cases have some form of autonomous decision-making, although currently only 2% are described as fully autonomous.
9. For a full summary, see discussion paper [DP5/22 – Artificial Intelligence and Machine Learning](#).
10. IMF (2024), [Global Financial Stability Report \(October 2024\)](#), Chapter 3, Figure 3.12.
11. FSB (2024), [The Financial Stability Implications of Artificial Intelligence](#), pages 21–22.
12. FSB (2024), Recommendation 3, [Leverage in Non-Bank Financial Intermediation: Consultation report](#).
13. There are policies set by the Bank, the PRA and the FCA to require regulated firms to deliver important business services and to set impact tolerances, even under severe but plausible disruption. The FPC has taken action with a view to further mitigating systemic risk in this area, setting an 'impact tolerance' for critical payments.
14. For a more general discussion of the impact of AI on near-term cyberthreat to the UK, see: National Cyber Security Centre (2024), [The near-term impact of AI on the cyber threat](#).
15. In more technical terms, the goal is the objective function of a technique or model. In optimisation problems, such as price prediction problems, the objective function aims to minimise some function of the error (generally the mean-squared error) between the price predicted by a technique and the actual price in the training data. Techniques that too accurately predict the training data may not generalise well to unseen data – a problem known as 'overfitting'. This can result in a model that performs well in backtesting, which evaluates a model's performance using historical data, but fails to perform as well in live implementation.
16. Human-in-the-loop is a concept that captures building AI systems with human input.