

III. Artificial intelligence and the economy: implications for central banks

Key takeaways

- Machine learning models excel at harnessing massive computing power to impose structure on unstructured data, giving rise to artificial intelligence (AI) applications that have seen rapid and widespread adoption in many fields.
- The rise of AI has implications for the financial system and its stability, as well as for macroeconomic outcomes via changes in aggregate supply (through productivity) and demand (through investment, consumption and wages).
- Central banks are directly affected by AI's impact, both in their role as stewards of monetary and financial stability and as users of AI tools. To address emerging challenges, they need to anticipate AI's effects across the economy and harness AI in their own operations.
- Data availability and data governance are key enabling factors for central banks' use of AI, and both rely on cooperation along several fronts. Central banks need to come together and foster a "community of practice" to share knowledge, data, best practices and AI tools.

Introduction

The advent of large language models (LLMs) has catapulted generative artificial intelligence (gen AI) into popular discourse. LLMs have transformed the way people interact with computers – away from code and programming interfaces to ordinary text and speech. This ability to converse through ordinary language as well as gen AI's human-like capabilities in creating content have captured our collective imagination.

Below the surface, the underlying mathematics of the latest AI models follow basic principles that would be familiar to earlier generations of computer scientists. Words or sentences are converted into arrays of numbers, making them amenable to arithmetic operations and geometric manipulations that computers excel at.

What is new is the ability to bring mathematical order at scale to everyday unstructured data, whether they be text, images, videos or music. Recent AI developments have been enabled by two factors. First is the accumulation of vast reservoirs of data. The latest LLMs draw on the totality of textual and audiovisual information available on the internet. Second is the massive computing power of the latest generation of hardware. These elements turn AI models into highly refined prediction machines, possessing a remarkable ability to detect patterns in data and fill in gaps.

There is an active debate on whether enhanced pattern recognition is sufficient to approximate "artificial general intelligence" (AGI), rendering AI with full human-like cognitive capabilities. Irrespective of whether AGI can be attained, the ability to impose structure on unstructured data has already unlocked new capabilities in many tasks that eluded earlier generations of AI tools.¹ The new generation of AI models could be a game changer for many activities and have a profound impact on the broader economy and the financial system. Not least, these same capabilities

can be harnessed by central banks in pursuit of their policy objectives, potentially transforming key areas of their operations.

The economic potential of AI has set off a gold rush across the economy. The adoption of LLMs and gen AI tools is proceeding at such breathtaking speed that it easily outpaces previous waves of technology adoption (Graph 1.A). For example, ChatGPT alone reached one million users in less than a week and nearly half of US households have used gen AI tools in the past 12 months. Mirroring rapid adoption by users, firms are already integrating AI in their daily operations: global survey evidence suggests firms in all industries use gen AI tools (Graph 1.B). To do so, they are investing heavily in AI technology to tailor it to their specific needs and have embarked on a hiring spree of workers with AI-related skills (Graph 1.C). Most firms expect these trends to only accelerate.²

This chapter lays out the implications of these developments for central banks, which impinge on them in two important ways.


First, AI will influence central banks' core activities as stewards of the economy. Central bank mandates revolve around price and financial stability. AI will affect financial systems as well as productivity, consumption, investment and labour markets, which themselves have direct effects on price and financial stability. Widespread adoption of AI could also enhance firms' ability to quickly adjust prices in response to macroeconomic changes, with repercussions for inflation dynamics. These developments are therefore of paramount concern to central banks.

Second, the use of AI will have a direct bearing on the operations of central banks through its impact on the financial system. For one, financial institutions such as commercial banks increasingly employ AI tools, which will change how they interact with and are supervised by central banks. Moreover, central banks and other authorities are likely to increasingly use AI in pursuing their missions in monetary policy, supervision and financial stability.


The adoption of AI¹

Graph 1


A. The adoption of AI is happening fast...



B. ...and in all sectors...



C. ...while investments in AI companies and job openings soar



¹ See technical annex for details.

Sources: Allcot (2023); Comin and Hobijn (2004); Maslej et al (2024); McKinsey & Company (2023); IMF, *World Economic Outlook*; US Census Bureau, *Current Population Survey*; International Telecommunication Union (ITU); PitchBook Data Inc; Our World in Data; Statista, *Digital Market Insights*; BIS.

Overall, the rapid and widespread adoption of AI implies that there is an urgent need for central banks to raise their game. To address the new challenges, central banks need to upgrade their capabilities both as *informed observers* of the effects of technological advancements as well as *users* of the technology itself. As observers, central banks need to stay ahead of the impact of AI on economic activity through its effects on aggregate supply and demand. As users, they need to build expertise in incorporating AI and non-traditional data in their own analytical tools. Central banks will face important trade-offs in using external vs internal AI models, as well as in collecting and providing in-house data vs purchasing them from external providers. Together with the centrality of data, the rise of AI will require a rethink of central banks' traditional roles as compilers, users and providers of data. To harness the benefits of AI, collaboration and the sharing of experiences emerge as key avenues for central banks to mitigate these trade-offs, in particular by reducing the demands on information technology (IT) infrastructure and human capital. Central banks need to come together to form a "community of practice" to share knowledge, data, best practices and AI tools.

The chapter starts with an overview of developments in AI, providing a deep dive into the underlying technology. It then examines the implications of the rise of AI for the financial sector. The discussion includes current use cases of AI by financial institutions and implications for financial stability. It also outlines the emerging opportunities and challenges and the implications for central banks, including how they can harness AI to fulfil their policy objectives. The chapter then discusses how AI affects firms' productive capacity and investment, as well as labour markets and household consumption, and how these changes in aggregate demand and supply affect inflation dynamics. The chapter concludes by examining the trade-offs arising from the use of AI and the centrality of data for central banks and regulatory authorities. In doing so, it highlights the urgent need for central banks to cooperate.

Developments in artificial intelligence

Artificial intelligence is a broad term, referring to computer systems performing tasks that require human-like intelligence. While the roots of AI can be traced back to the late 1950s, the advances in the field of **machine learning** in the 1990s laid the foundations of the current generation of AI models. Machine learning is a collective term referring to techniques designed to detect patterns in the data and use them in prediction or to aid decision-making.³

The development of **deep learning** in the 2010s constituted the next big leap. Deep learning uses neural networks, perhaps the most important technique in machine learning, underpinning everyday applications such as facial recognition or voice assistants. The main building block of neural networks is *artificial neurons*, which take multiple input values and transform them to output as a set of numbers that can be readily analysed. The artificial neurons are organised to form a sequence of *layers* that can be stacked: the neurons of the first layer take the input data and output an activation value. Subsequent layers then take the output of the previous layer as input, transform it and output another value, and so forth. A network's *depth* refers to the number of layers. More layers allow neural networks to capture increasingly complex relationships in the data. The weights determining the strength of connections between different neurons and layers are collectively called *parameters*, which are improved (known as *learning*) iteratively during training. Deeper networks with more parameters require more training data but predict more accurately.

A key advantage of deep learning models is their ability to work with unstructured data. They achieve this by "**embedding**" qualitative, categorical or visual data, such

as words, sentences, proteins or images, into arrays of numbers – an approach pioneered at scale by the Word2Vec model (see Box A). These arrays of numbers (ie vectors) are interpreted as points in a vector space. The distance between vectors conveys some dimension of similarity, enabling algebraic manipulations on what is originally qualitative data. For example, the vector linking the embeddings of the words “big” and “biggest” is very similar to that between “small” and “smallest”. Word2Vec predicts a word based on the surrounding words in a sentence. The body of text used for the embedding exercise is drawn from the open internet through the “common crawl” database. The concept of embedding can be taken further into mapping the space of economic ideas, uncovering latent viewpoints or methodological approaches of individual economists or institutions (“**personas**”). The space of ideas can be linked to concrete policy actions, including monetary policy decisions.⁴

The advent of **LLMs** allows neural networks to access the whole context of a word rather than just its neighbour in the sentence. Unlike Word2Vec, LLMs can now capture the nuances of translating uncommon languages, answer ambiguous questions or analyse the sentiment of texts. LLMs are based on the **transformer** model (see Box B). Transformers rely on “multi-headed attention” and “positional encoding” mechanisms to efficiently evaluate the context of any word in the document. The context influences how words with multiple meanings map into arrays of numbers. For example, “bond” could refer to a fixed income security, a connection or link, or a famous espionage character. Depending on the context, the “bond” embedding vector lies geometrically closer to words such as “treasury”, “unconventional” and “policy”; to “family” and “cultural”; or to “spy” and “martini”. These developments have enabled AI to move from narrow systems that solve one specific task to more general systems that deal with a wide range of tasks.

LLMs are a leading example of **gen AI applications** because of their capacity to understand and generate accurate responses with minimal or even no prior examples (so-called **few-shot** or **zero-shot** learning abilities). Gen AI refers to AIs capable of generating content, including text, images or music, from a natural language prompt. The prompts contain instructions in plain language or examples of what users want from the model. Before LLMs, machine learning models were trained to solve one task (eg image classification, sentiment analysis or translating from French to English). It required the user to code, train and roll out the model into production after acquiring sufficient training data. This procedure was possible for only selected companies with researchers and engineers with specific skills. An LLM has few-shot learning abilities in that it can be given a task in plain language. There is no need for coding, training or acquiring training data. Moreover, it displays considerable versatility in the range of tasks it can take on. It can be used to first classify an image, then analyse the sentiment of a paragraph and finally translate it into any language. Therefore, LLMs and gen AI have enabled people using ordinary language to automate tasks that were previously performed by highly specialised models.

The capabilities of the most recent crop of AI models are underpinned by advances in **data** and **computing power**. The increasing availability of data plays a key role in training and improving models. The more data a model is trained on, the more capable it usually becomes. Furthermore, machine learning models with more parameters improve predictions when trained with sufficient data. In contrast to the previous conventional wisdom that “over-parameterisation” degrades the forecasting ability of models, more recent evidence points to a remarkable resilience of machine learning models to over-parameterisation. As a consequence, LLMs with well designed learning mechanisms can provide more accurate predictions than traditional parametric models in diverse scenarios such as computer vision, signal processing and natural language processing (NLP).⁵

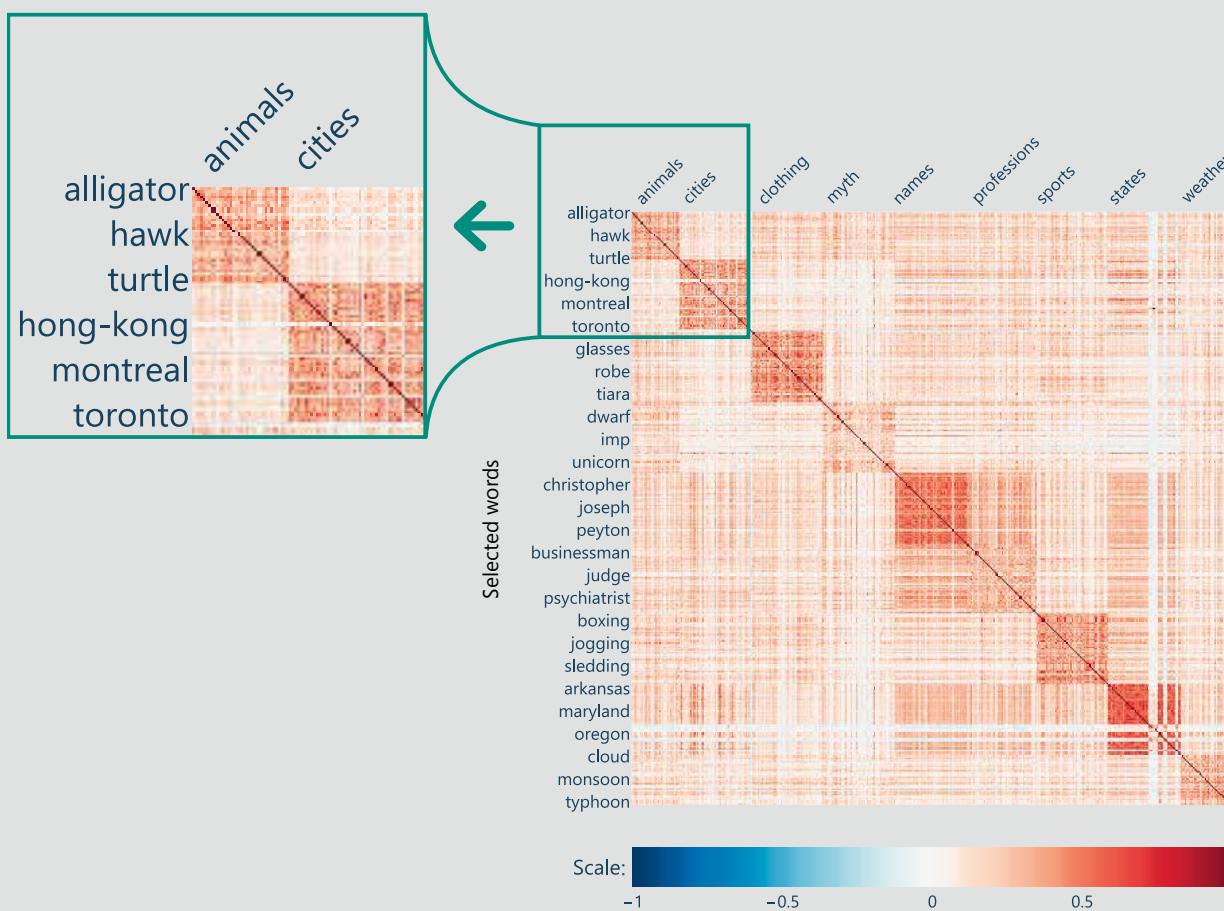
Words as vectors: a primer on embeddings

Modern machine learning methods excel at imposing mathematical structure on unstructured data, allowing massive computing power to be unleashed in processing information. The mapping that imposes such structure is known as an “embedding”, and the canonical example is the embedding of words as points in a vector space, so that each word is associated with an array of numbers.

An early example of word embedding is Word2Vec,¹ which maps a word to an embedding vector of a few hundred dimensions that is learned by a neural network. The neural network is refined by being asked to predict the centre word in a short window of text (typically four to eight words before and after the centre word) and being scored by its success rate. This procedure is known as the “Continuous Bag of Words” method because all surrounding words are first added into a single vector. The Word2Vec learning algorithm computes the prediction error over all the words in a corpus (which can be trillions of words) and iteratively adjusts the embedding vector for each word to reduce this classification error and optimise prediction.

Embedding distances between 420 words in nine categories of words¹

Graph A1



¹ Cosine similarity matrix between 420 words. The value ranges from -1 (completely dissimilar) to 1 (completely similar), with 0 indicating orthogonality (no similarity). The y-axis labels correspond to selected 420 words; the axis labels indicate the categories to which these words belong.

Source: Adapted from Grand et al (2022).


These procedures result in similar embeddings for words with similar meaning, in the sense that the distance between the vectors representing the two words is mathematically close. For example, the embedding of the word “cat” is close to that of the word “mouse”, and that of “Mexico” close to “Indonesia”. Graph A1 illustrates the “cosine similarity” between 420 words in nine different word categories (animals, cities etc).

Cosine similarity measures the cosine of the angle between two non-zero vectors, reflecting how similar their directions are. It calculates the dot product of the vectors divided by the product of their norms. The value ranges from -1 (completely dissimilar) to 1 (completely similar), with 0 indicating orthogonality (no similarity). In Graph A1, the colour scheme indicates the degree of similarity between word pairs. The diagonal of this matrix consists of 1 everywhere, as the diagonal measures each word's similarity with itself. Darker red indicates high cosine similarity, while lighter red indicates low similarity. Graph A1 shows that words from the same category (eg animals) have a high cosine similarity, while they have low cosine similarity with words from other categories (eg cities or sports). The resulting vectors give rise to embeddings that can be used in various natural language processing tasks such as text classification, sentiment analysis and machine translation with minimal or no human-labelled data.

The embeddings uncover the mathematical relationships between words. Not only are similar words placed closer together in the vector space, but the semantic connections are also captured through the mathematical relationships between the vector embedding of each word. For instance, analogies like "man is to woman as king is to?" can be solved directly from vector addition and subtraction operations: $\vec{queen} = \vec{woman} + \vec{king} - \vec{man}$. These embedding relationships also apply to the link between countries and their capitals ($\vec{Quito} = \vec{Ecuador} + \vec{Oslo} - \vec{Norway}$), opposites ($\vec{unethical} = \vec{ethical} + \vec{impossibly} - \vec{possibly}$), and the tense of words ($\vec{swam} = \vec{swimming} + \vec{walked} - \vec{walking}$). Semantic relationships between words can also be projected to concepts. Graph A2 illustrates how by projecting the word embeddings of animals to the vector representing variation in size (ie the difference between the word embedding for "large" and "small"), the animals are mostly sorted according to their sizes.

Embedding projection of animal words onto size concept vector¹

Graph A2



¹ Two-dimensional illustration, as the embeddings are in a 300-dimensional vector space.

Source: Adapted from Grand et al (2022).

Word2Vec has subsequently been superseded by other methods that achieve more meaningful embedding, such as GloVe, ELMo, BERT and GPT,² by employing more sophisticated learning of concepts with more complex neural network architectures. The latest models (BERT and GPT) rely on the transformer architecture (see Box B). BERT and GPT are referred to as language models, not word embeddings. They use the whole text as context, multiple paths to capture different meanings and neural networks with trillions of tunable parameters.

¹ Mikolov et al (2013) ² Pennington et al (2014), Peters et al (2018), Devlin et al (2018) and Brown et al (2019).

An implication is that more capable models tend to be larger models that need more data. Bigger models and larger data sets therefore go together and increase computational demands. The use of advanced techniques on vast troves of data would not have been possible without substantial increases in computing power – in particular, the computational resources employed by AI systems – which has been doubling every six months.⁶ The interplay between large amounts of data and computational resources implies that just a handful of companies provide cutting-edge LLMs, an issue revisited later in the chapter.

Some commentators have argued that AI has the potential to become the next **general-purpose technology**, profoundly impacting the economy and society. General-purpose technologies, like electricity or the internet, eventually achieve widespread usage, give rise to versatile applications and generate spillover effects that can improve other technologies. The adoption pattern of general-purpose technologies typically follows a J-curve: it is slow at first, but eventually accelerates. Over time, the pace of technology adoption has been speeding up. While it took electricity or the telephone decades to reach widespread adoption, smartphones accomplished the same in less than a decade. AI features two distinct characteristics that suggest an even steeper J-curve. First is its remarkable speed of adoption, reflecting ease of use and negligible cost for users. Second is its widespread use at an early stage by households as well as firms in all industries.

Of course, there is substantial uncertainty about the long-term capabilities of gen AI. Current LLMs can **fail elementary logical reasoning** tasks and struggle with counterfactual reasoning, as illustrated in recent BIS work.⁷ For example, when posed with a logical puzzle that demands reasoning about the knowledge of others and about counterfactuals, LLMs display a distinctive pattern of failure. They perform flawlessly when presented with the original wording of a puzzle, which they have likely seen during their training. They falter when the same problem is presented with small changes of innocuous details such as names and dates, suggesting a lack of true understanding of the underlying logic of statements. Ultimately, current LLMs do not know what they do not know. LLMs also suffer from the **hallucination problem**: they can present a factually incorrect answer as if it were correct, and even invent secondary sources to back up their fake claims. Unfortunately, hallucinations are a feature rather than a bug in these models. LLMs hallucinate because they are trained to predict the statistically plausible word based on some input. But they cannot distinguish what is linguistically probable from what is factually correct.

Do these problems merely reflect the limits posed by the size of the training data set and the number of model parameters? Or do they reflect more fundamental limits to knowledge that is acquired through language alone? Optimists acknowledge current limitations but emphasise the potential of LLMs to exceed human performance in certain domains. In particular, they argue that terms such as “reason”, “knowledge” and “learning” rightly apply to such models. Sceptics point out the limitations of LLMs in reasoning and planning. They argue that the main limitation of LLMs derives from their exclusive reliance on language as the medium of knowledge. As LLMs are confined to interacting with the world purely through language, they lack the tacit non-linguistic, shared understanding that can be acquired only through active engagement with the real world.⁸

Whether AI will eventually be able to perform tasks that require deep logical reasoning has implications for its long-run economic impact. Assessing which tasks will be impacted by AI depends on the specific cognitive abilities required in those tasks. The discussion above suggests that, at least in the near term, AI faces challenges in reaching human-like performance. While it may be able to perform tasks that require moderate cognitive abilities and even develop “emergent” capabilities, it is not yet able to perform tasks that require logical reasoning and judgment.

A primer on the transformer architecture

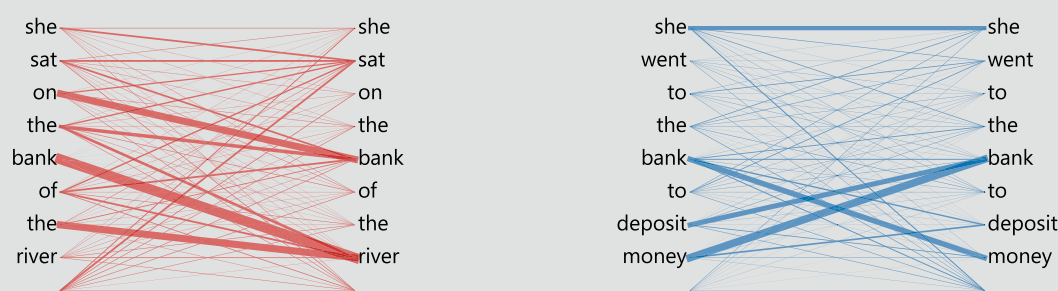
The transformer architecture¹ has been a breakthrough in natural language processing (NLP), laying the foundation for the development of advanced large language models (LLMs) such as BERT (Bidirectional Encoder Representations from Transformers)² and GPT (Generative Pre-trained Transformer).³ At the heart of the transformer architecture are two innovations: “multi-headed attention” and “positional encoding”.

Attention allows each word in a text to be understood in relation to every other word, enhancing the models’ capacity to take account of context and relationships within the text. Multi-headed attention allows several parallel neural networks to capture different meanings for the same word. For instance, it can discern the different meanings of “bank” in “She sat on the bank of the river” versus “She went to the bank to deposit money” by focusing on surrounding words like “river” versus “money” or “deposit” (Graph B1). The attention mechanism computes dot products of the input query vector (“bank”) with key vectors (each word in the text), resulting in a score matrix. This matrix is then used to obtain so-called attention weights, which measure the similarity between the query and key vectors. By combining the information into a value vector, the model relates each word to the most relevant parts of the query vector. This ability to interpret the meaning of words based on much wider contexts is a significant advance over previous NLP models.

Visualised examples of the attention mechanism in transformers¹

Graph B1

A. Understanding the meaning of “bank”: first example... B. ...and second example



¹ Examples of the attention mechanism within 10th head of specific layers in the BERT model. The connecting lines and their thickness represent the attention scores (ie the relevance) between words. This visualisation illustrates how the word “bank” varies in its attention to different words depending on the context.

Sources: Devlin et al (2018); Vig (2019); BIS.

Positional encoding enables transformers to process data concurrently rather than sequentially. This sets them apart from earlier neural network models such as Recurrent Neural Networks and Long Short-Term Memory. Sequential models are slow to train and memory-consuming, and they suffer from the so-called vanishing gradient problem (ie the signals that carry information about how to update the weights in a neural network eventually become too weak to effectively train deep layers). By embedding each word with positional information, transformers preserve the sequence of words, and the model can be parallelised during training. This capability allows training with more data and the building of bigger models, which leads them to accurately differentiate between sentences like “Inflation causes a rate hike” and “A rate hike causes inflation”, where word order determines meaning.

The integration of multi-headed attention and positional encoding has significantly enhanced the performance of language models. Transformer-based models exhibit remarkable proficiency in interpreting context and managing complex relationships within text. The result is superior accuracy and fluency in a variety of NLP tasks. The parallel computation capabilities of transformers, facilitated by advances in graphics processing unit technology, enable rapid processing of vast data sets and a complex linguistic structure.

¹ See Vaswani et al (2017). ² See Devlin et al (2018). ³ GPT is an underlying model of ChatGPT. See Radford et al (2018).

Financial system impact of AI

The financial sector is among those facing the greatest opportunities and risks from the rise of AI, due to its high share of cognitively demanding tasks and data-intensive nature.⁹ Table 1 illustrates the impact of AI in four key areas: payments, lending, insurance and asset management.

Across all four areas, AI can substantially **enhance efficiency and lower costs** in back-end processing, regulatory compliance, fraud detection and customer service. These activities give full play to the ability of AI models to identify patterns of interest in seemingly unstructured data. Indeed, “finding a needle in the haystack” is an activity that plays to the greatest strength of machine learning models. A striking example is the improvement of know-your-customer (KYC) processes through quicker data processing and the enhanced ability to detect fraud, allowing financial institutions to ensure better compliance with regulations while lowering costs.¹⁰ LLMs are also increasingly being deployed for customer service operations through AI chatbots and co-pilots.

In **payments**, the abundance of transaction-level data enables AI models to overcome long-standing pain points. A prime example comes from correspondent banking, which has become a high-risk, low-margin activity. Correspondent banks played a key role in the expansion of cross-border payment activity by enabling transaction settlement, cheque clearance and foreign exchange operations. Facing heightened customer verification and anti-money laundering (AML) requirements, banks have systematically retreated from the business (Graphs 2.A and 2.B). Such retreat fragments the global payment system by leaving some regions less connected (Graph 2.C), handicapping their connectivity with the rest of the financial system. The decline in correspondent banking is part of a general de-risking trend, with returns from processing transactions being small compared with the risks of penalties from breaching AML, KYC and countering the financing of terrorism (CFT) requirements.¹¹


A key use case of AI models is to improve KYC and AML processes by enhancing (i) the ability to understand the compliance and reputational risks that clients might carry, (ii) due diligence on the counterparties of a transaction and (iii) the analysis of

Opportunities, challenges and financial stability risks of AI in the financial sector Table 1


	Payments	Lending	Insurance	Asset management
General opportunities	Back-end processing, virtual assistants, co-pilots, fraud detection, regulatory compliance			
Sector-specific opportunities	Liquidity management, AML/KYC	Credit risk analysis, financial inclusion	Risk assessment, pricing, claims processing	Portfolio allocation, algorithmic trading, robo-advising, asset embeddings
General challenges	Lack of explainability, data silos, third-party dependencies, algorithmic collusion, hallucinations, cyber security risks			
Sector-specific challenges	Liquidity crises, sophisticated fraud and cyber attacks	Algorithmic discrimination, privacy concerns		Zero-sum arms race for private gains, herding, algorithmic coordination
Financial stability challenges	Herding, network interconnectedness and procyclicality, single point of failure, incorrect decisions based on short samples of non-representative data, spillovers from real sector			

Source: Adapted from Aldasoro, Gambacorta, Korinek, Shreeti and Stein (2024).


A. Banks have been retreating



B. The decline is global



C. Some regions are less connected



¹ See technical annex for details.

Sources: Garratt et al (2024); Rice et al (2020); CPMI (2023).

payment patterns and anomaly detection. By bringing down costs and reducing risks through greater speed and automation, AI holds the promise to reverse the decline in correspondent banking.

The ability of AI models to detect patterns in the data is helping financial institutions address many of these challenges. For example, financial institutions are using AI tools to enhance fraud detection and to identify security vulnerabilities. At the global level, surveys indicate that around 70% of all financial services firms are using AI to enhance cash flow predictions and improve liquidity management, fine-tune credit scores and improve fraud detection.¹²

In **credit assessment and lending**, banks have used machine learning for many years, but AI can bring further capabilities. For one, AI could greatly enhance credit scoring by making use of unstructured data. In deciding whether to grant a loan, lenders traditionally rely on standardised credit scores, at times combined with easily accessible variables such as loan-to-value or debt-to-income ratios. AI-based tools enable lenders to assess individuals' creditworthiness with alternative data. These can include consumers' bank account transactions or their rental, utility and telecommunications payments data. But they can also be of a non-financial nature, for example applicants' educational history or online shopping habits. The use of non-traditional data can significantly improve default prediction, especially among underserved groups for whom traditional credit scores provide an imprecise signal about default probability. By being better able to spot patterns in unstructured data and detect "invisible primes", ie borrowers that are of high quality even if their credit scores indicate low quality, AI can enhance financial inclusion.¹³

AI has numerous applications in **insurance**, particularly in risk assessment and pricing. For example, companies use AI to automatically analyse images and videos to assess property damage due to natural disasters or, in the context of compliance,

whether claims of damages correspond to actual damages. Underwriters, actuaries or claims adjusters further stand to benefit from AI summarising and synthesising data gathered during a claim's life cycle, such as call transcripts and notes, as well as legal and medical paperwork. More generally, AI is bound to play an increasingly important role in assessing different types of risks. For example, some insurance companies are experimenting with AI methods to assess climate risks by identifying and quantifying emissions based on aerial images of pollution. However, to the extent that AI is better at analysing or inferring individual-level characteristics in risk assessments, including those whose use is prohibited by regulation, existing inequalities could be exacerbated – an issue revisited in the discussion on the macroeconomic impact of AI.

In **asset management**, AI models are used to predict returns, evaluate risk-return trade-offs and optimise portfolio allocation. Just as LLMs assign different characteristics to each word they process, they can be used to elicit unobservable features of financial data (so-called asset embeddings). This allows market participants to extract information (such as firm quality or investor preferences) that is difficult to discern from existing data. In this way, AI models can provide a better understanding of the risk-return properties of portfolios. Models that use asset embeddings can outperform traditional models that rely only on observable characteristics of financial data. Separately, AI models are useful in algorithmic trading, owing to their ability to analyse large volumes of data quickly. As a result, investors benefit from quicker and more precise information as well as lower management fees.¹⁴

The widespread use of AI applications in the financial sector, however, brings new challenges. These pertain to cyber security and operational resilience as well as financial stability.

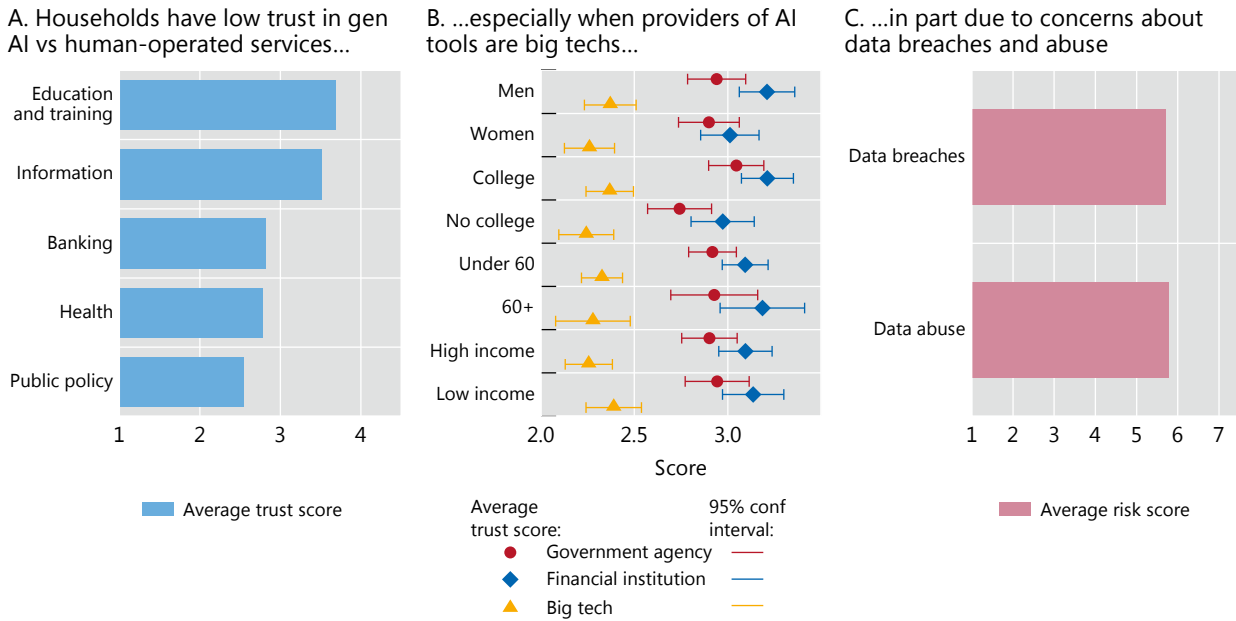
The reliance on AI heightens concerns about **cyber attacks**, which regularly feature among the top worries in the financial industry. Traditionally, phishing emails have been used to trick a user to run a malicious code (malware) to take over the user's device. Credential phishing is the practice of stealing a user's login and password combination by masquerading as a reputable or known entity in an email, instant message or another communication channel. Attackers then use the victim's credentials to carry out attacks on additional targets and gain further access.¹⁵ Gen AI could vastly expand hackers' ability to write credible phishing emails or to write malware and use it to steal valuable information or encrypt a company's files for ransom. Moreover, gen AI allows hackers to imitate the writing style or voice of individuals, or even create fake avatars, which could lead to a dramatic rise in phishing attacks. These developments expose financial institutions and their customers to a greater risk of fraud.

But AI also introduces altogether **new sources of cyber risk**. *Prompt injection attacks*, one of the most widely reported weaknesses in LLMs, refer to an attacker creating an input to make the model behave in an unintended way. For example, LLMs are usually instructed not to provide dangerous information, such as how to manufacture napalm. However, in the infamous *grandma jailbreak*, where the prompter asked ChatGPT to pretend to be their deceased grandmother telling a bedtime story about the steps to produce napalm, the chatbot did reveal this information. While this vulnerability has been fixed, others remain. *Data poisoning attacks* refer to malicious tampering with the data an AI model is trained on. For example, an attacker could adjust input data so that the AI model fails to detect phishing emails. *Model poisoning attacks* deliberately introduce malware, manipulating the training process of an AI system to compromise its integrity or functionality. This attack aims to alter the model behaviour to serve the attacker's purposes.¹⁶ As more applications use data created by LLMs themselves, such attacks could have increasingly severe consequences, leading to heightened operational risks among financial institutions.

Households' trust in generative AI (gen AI)¹

Score, 1 (lowest)–7 (highest)

Graph 3



¹ Based on a representative sample of US households from the Survey of Consumer Expectations. See technical annex for details.

Sources: Aldasoro, Armantier, Doerr, Gambacorta and Oliviero (2024a); Federal Reserve Bank of New York, *Survey of Consumer Expectations*.

Greater use of AI raises issues of **bias and discrimination**. Two examples stand out. The first relates to consumer protection and fair lending practices. As with traditional models, AI models can reflect biases and inaccuracies in the data they are trained on, posing risks of unjust decisions, excluding some groups from socially desirable insurance markets and perpetuating disparities in access to credit through algorithmic discrimination.¹⁷ Consumers care about these risks: recent evidence from a representative survey of US households suggests a lower level of trust in gen AI than in human-operated services, especially in high-stakes areas such as banking and public policy (Graph 3.A) and when AI tools are provided by big techs (Graph 3.B).¹⁸ The second example relates to the challenge of ensuring data privacy and confidentiality when dealing with growing volumes of data, another key concern for users (Graph 3.C). In the light of the high privacy standards that financial institutions need to adhere to, this heightens **legal risks**. The lack of explainability of AI models (ie their black box nature) as well as their tendency to hallucinate amplify these risks.

Another operational risk arises from relying on just a few providers of AI models, which increases **third-party dependency risks**. Market concentration arises from the centrality of data and the vast costs of developing and implementing data-hungry models. Heavy up-front investment is required to build data storage facilities, hire and train staff, gather and clean data and develop or refine algorithms. However, once the infrastructure is in place, the cost of adding each extra unit of data is negligible. This centrality leads to so-called data gravity: companies that already have an edge in collecting, storing and analysing data can provide better-trained AI tools, whose use creates ever more data over time. The consequence of data gravity is that only a few companies provide cutting-edge LLMs. Any failure among or cyber attack on these providers, or their models, poses risks to financial institutions relying on them.

The reliance of market participants on the same handful of algorithms could lead to **financial stability risks**. These could arise from AI's ubiquitous adoption throughout the financial system and its growing capability to make decisions independently and without human intervention ("automaticity") at a speed far beyond human capacity. The behaviour of financial institutions using the same algorithms could amplify procyclicality and market volatility by exacerbating herding, liquidity hoarding, runs and fire sales. Using similar algorithms trained on the same data can also lead to coordinated recommendations or outright collusive outcomes that run afoul of regulations against market manipulation, even if algorithms are not trained or instructed to collude.¹⁹ In addition, AI may hasten the development and introduction of new products, potentially leading to new and little understood risks.

Harnessing AI for policy objectives

Central banks stand at the intersection of the monetary and financial systems. As stewards of the economy through their monetary policy mandate, they play a pivotal role in maintaining economic stability, with a primary objective of ensuring price stability. Another essential role is to safeguard financial stability and the payment system. Many central banks also have a role in supervising and regulating commercial banks and other participants of the financial system.

Central banks are not simply passive observers in monitoring the impact of AI on the economy and the financial system. They can harness AI tools themselves in pursuit of their policy objectives and in addressing emerging challenges. In particular, the use of LLMs and AI can support central banks' key tasks of information collection and statistical compilation, macroeconomic and financial analysis to support monetary policy, supervision, oversight of payment systems and ensuring financial stability. As early adopters of machine learning methods, central banks are well positioned to reap the benefits of AI tools.²⁰

Data are the major resource that stand to become more valuable due to the advent of AI. A particularly rich source of data is the payment system. Such data present an enormous amount of information on economic transactions, which naturally lends itself to the powers of AI to detect patterns.²¹ Dealing with such data necessitates adequate privacy-preserving techniques and the appropriate data governance frameworks.


The BIS Innovation Hub's Project Aurora explores some of these issues. Using a synthetic data set emulating money laundering activities, it compares various machine learning models, taking into account payment relationships as input. The comparison occurs under three scenarios: transaction data that are siloed at the bank level, national-level pooling of data and cross-border pooling. The models undergo training with known simulated money laundering transactions and subsequently predict the likelihood of money laundering in unseen synthetic data.

The project offers two key insights. First, machine learning models outperform the traditional rule-based methods prevalent in most jurisdictions. Graph neural networks, in particular, demonstrate superior performance, effectively leveraging comprehensive payment relationships available in pooled data to more accurately identify suspect transaction networks. And second, machine learning models are particularly effective when data from different institutions in one or multiple jurisdictions are pooled, underscoring a premium on cross-border coordination in AML efforts (Graph 4).

The benefits of coordination are further illustrated by Project Agorá. This project gathers seven central banks and private sector participants to bring tokenised central bank money and tokenised deposits together on the same programmable platform.

Machine learning models' performance in different monitoring scenarios¹

Graph 4



¹ Transaction data visible on three different levels of analysis: the view of each financial institution (silo), the national view of a single country (national) and the cross-border view across countries (cross-border).


Source: BIS Innovation Hub (2023).

The tokenisation built into Agorá would allow the platform to harness three capabilities: (i) combining messaging and account updates as a single operation; (ii) executing payments atomically rather than as a series of sequential updates; and (iii) drawing on privacy-preserving platform resources for KYC/AML compliance. In traditional correspondent banking, information checks and account updates are made sequentially and independently, with significant duplication of effort (Graph 5.A). In contrast, in Agorá the contingent performance of actions enabled by tokenisation allows for the combination of assets, information, messaging and clearing into a single atomic operation, eliminating the risk of reversals (Graph 5.B). In turn, privacy-enhancing data-sharing techniques can significantly simplify compliance checks, while all existing rules and regulations are adhered to as part of the pre-screening process.²²


In the development of a new payment infrastructure like Agorá, great care must be taken to ensure potential gains are not lost due to fragmentation. This can be done via access policies to the infrastructure or via interoperability, as advocated in the idea of the Finternet. This refers to multiple interconnected financial ecosystems, much like the internet, designed to empower individuals and businesses by placing them at the centre of their financial lives. The Finternet leverages innovative technologies such as tokenisation and unified ledgers, underpinned by a robust economic and regulatory framework, to expand the range and quality of savings and financial services. Starting with assets that can be easily tokenised holds the greatest promise in the near term.²³

Central banks also see great benefits in using gen AI to **improve cyber security**. In a recent BIS survey of central bank cyber experts, a majority deem gen AI to offer more benefits than risks (Graph 6.A) and think it can outperform traditional methods in enhancing cyber security management.²⁴ Benefits are largely expected in areas such as the automation of routine tasks, which can reduce the costs of time-consuming activities traditionally performed by humans (Graph 6.B). But human expertise will remain important. In particular, data scientists and cyber security experts are expected to play an increasingly important role. Additional cyber-related benefits from AI

A. Silos lead to duplication of effort in AML screening



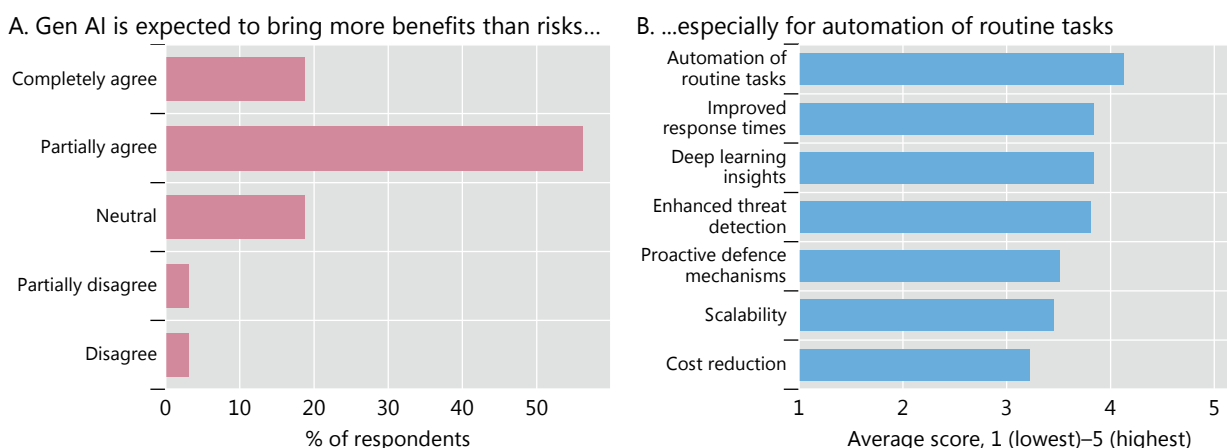
B. Collective effort and data-sharing



Source: Garratt et al (2024).

include the enhancement of threat detection, faster response times to cyber attacks and the learning of new trends, anomalies or correlations that might not be obvious to human analysts. In addition, by leveraging AI, central banks can now craft and deploy highly convincing phishing attacks as part of their cyber security training. Project Raven of the BIS Innovation Hub is one example of the use of AI to enhance cyber resilience (see Box C).

The challenge for central banks in using AI tools comes in two parts. The first is the availability of timely data, which is a necessary condition for any machine learning application. Assuming this issue is solved, the second challenge is to structure the data in a way that yields insights. This second challenge is where machine learning tools, and in particular LLMs, excel. They can transform unstructured data from a variety of sources into structured form in real time. Moreover, by converting time series data into



¹ Based on a survey of 32 members of the Global Cyber Resilience Group in January 2024. See technical annex for details.

Source: Aldasoro, Doerr, Gambacorta, Notra, Oliviero and Whyte (2024).

tokens resembling textual sequences, LLMs can be applied to a wide array of time series forecasting tasks. Just as LLMs are trained to guess the next word in a sentence using a vast database of textual information, LLM-based forecasting models use similar techniques to estimate the next numerical observation in a statistical series.

These capabilities are particularly promising for **nowcasting**. Nowcasting is a technique that uses real-time data to provide timely insights. This method can significantly improve the accuracy and timeliness of economic predictions, particularly during periods of heightened market volatility. However, it currently faces two important challenges, namely the limited usability of timely data and the necessity to pre-specify and train models for concrete tasks.²⁵ LLMs and gen AI hold promise to overcome both bottlenecks (see Box D). For example, an LLM fine-tuned with financial news can readily extract information from social media posts or non-financial firms’ and banks’ financial statements or transcripts of earning reports and create a sentiment index. The index can then be used to nowcast financial conditions, monitor the build-up of risks or predict the probability of recessions.²⁶ Moreover, by categorising texts into specific economic topics (eg consumer demand and credit conditions), the model can pinpoint the source of changes in sentiment (eg consumer sentiment or credit risk). Such data are particularly relevant early in the forecasting process when traditional hard data are scarce.

Beyond financial applications, AI-based nowcasting can also be useful to understand real-economy developments. For example, transaction-level data on household-to-firm or firm-to-firm payments, together with machine learning models, can improve nowcasting of consumption and investment. Another use case is measuring supply chain bottlenecks with NLP, eg based on text in the so-called Beige Book. After classifying sentences related to supply chains, a deep learning algorithm classifies the sentiment of each sentence and provides an index that offers a real-time view of supply chain bottlenecks. Such an index can be used to predict inflationary pressures. Many more examples exist, ranging from nowcasting world trade to climate risks.²⁷

Access to granular data can also enhance central banks’ ability to **track developments across different industries and regions**. For example, with the help of AI, data from job postings or online retailers can be used to track wage developments and employment dynamics across occupations, tasks and industries. Such a real-time and detailed view of labour market developments can help central

BIS Innovation Hub projects in artificial intelligence

The BIS Innovation Hub is exploring the use of artificial intelligence (AI) to support central banks and supervisors in their missions. So far, eight projects – Ellipse, Aurora, Gaia, Symbiosis, Raven, Neo, Spectrum and Insight – have employed AI methods. They cover a wide range of use cases from information collection and statistical compilation, payments oversight and supervision, and macroeconomic and financial analysis to monetary policy analysis (Table C1). These projects draw on both in-house expertise and that of external providers.

An overview of BIS Innovation Hub projects using AI

Table C1

	Ellipse	Aurora	Gaia	Symbiosis	Raven	Neo	Spectrum	Insight
Main use case	Match entities in news with those of supervisory interest	Enhance AML suspicious transaction monitoring across firms & borders	Extract climate-related data from ESG reports	Develop methods for Scope 3 emission disclosure	Process cyber security & resilience documents to generate answers to assessment questions	Create and forecast economic indicators using timely and granular data	Structure big data on micro prices for inflation nowcasting	Extract info & data on firm supply chain dependencies
BIS IH Centre	Singapore	Nordic	Eurosystem	Hong Kong	Nordic	Swiss	Eurosystem	Hong Kong
Status	Completed	Ongoing (Phase 2)	Completed	Ongoing	Ongoing	Ongoing	Ongoing	Ongoing
Key theme	Suptech/regtech		Green finance		Cyber security	Monetary policy tech		
NLP	✓	✗	✓	✓	✓	✗	✓	✗
LLM	✗	✗	✓	✓	✓	✓	✓	✗
SML	✗	✓	✗	✓	✗	✓	✗	✗
UML	✗	✓	✗	✓	✗	✓	✓	✗
Other	✗	✓	✗	✗	✓	✗	✗	✓ ¹

LLM = large language model; NLP = natural language processing; SML = supervised machine learning (ML); UML = unsupervised ML.

¹ As the project is in early stages, the types of AI technologies to be used are still to be determined.

Source: BIS Innovation Hub.

The experiments showcase the value of AI-enabled applications for central banks and the financial sector. For example, Project Gaia demonstrated the power of creating AI-enabled intelligent tools to automate data extraction from unstructured environmental, social and governance (ESG) reports for climate risk analysis. The combination of semantic search together with the iterative and systematic prompting of a large language model (LLM) enabled Gaia to navigate differences in disclosure frameworks, thereby enhancing the comparability of climate-related information. Project Aurora demonstrated that machine learning using graph neural networks and privacy-enhancing methods for sharing underlying network data can substantially improve detection of complex schemes such as “mule accounts” or “smurfing” activities in money laundering networks. Project Ellipse showed that machine learning methods are useful in risk assessment and analysis, alerting supervisors in real time about issues that might need further investigation.

Several technical challenges were also identified. These include LLMs’ long response times, randomness (non-repeatability) in their responses and hallucinations (Gaia), and computational costs and insufficient data quality and consistency across financial institutions (Aurora and Ellipse). Appropriate design choices and collaboration will help mitigate these challenges, as discussed further in the final section of this chapter.

Nowcasting with artificial intelligence

Most central banks support their economic analysis with so-called nowcasting models. The goal of nowcasting is to produce high-frequency assessments and forecasts, for example of gross domestic product (GDP) growth or inflation, that can be updated easily as soon as new data become available.

A crucial input to nowcasting models is timely data. As no single indicator suffices to accurately track economic activity in real time, nowcasting models often process large and complex data sets that contain dozens, if not hundreds, of indicators. Data series can include information ranging from industrial production to surveys on the sentiment of purchasing managers to credit card spending data. They can also include web data scraped from online retailers or social media. Recent advances in the econometrics of high-dimensional data, which enabled models to synthesise important parts of the complexity of the aggregate economy into just a few indicators, have substantially improved nowcasting. An example is the Federal Reserve Bank of New York's nowcasting model. It uses a wide range of data to extract the latent factors that drive movements in the data and produces a forecast of several economic series, in particular GDP.¹ However, the limited availability of timely data makes nowcasting challenging. For example, structured data on firms' hiring or consumer spending are only available with a lag, rendering the nowcasting exercise only as good as the data lags.

Nowcasting models could benefit from deep neural networks' embeddings and their ability to quickly convert unstructured data into a structured format. Thanks to data embeddings, large language models (LLMs) have the ability to process vast amounts of text- or image-based data, and they can create readily available data series from, for example, news reports, social media postings, web searches or aerial images, such as of car traffic in retailers' parking lots. Moreover, novel LLMs can process these data further to extract sentiment or tone in text or speech data. Such sentiment indices can substantially improve forecasts for key macroeconomic quantities, such as unemployment, GDP growth and inflation.² In this way, nowcasting models could draw on a much richer set of real-time data.

Another way LLMs can improve nowcasting models is through their "few-shot learning" abilities, ie the capacity to generate accurate responses with minimal or even no prior examples. Currently, most models used for time series forecasting are highly specialised, designed for specific settings, especially if they use natural language processing. A model that is used to nowcast GDP looks different from one used to nowcast the build-up of financial risks. Expert judgment is required to choose the right input variables and specify model parameters. This means that, as economic conditions and the focus of analysis change, it takes some time to adjust models accordingly. In contrast, few-shot learners display much greater versatility in what they can do. Recent work shows that LLMs trained to predict (forecast) the next word, by converting time series data into tokens resembling textual sequences (ie embedding), can be directly used for time series prediction.³ In other words, due to their advanced capabilities for few-shot learning, they can be readily applied to a wide array of time series forecasting tasks without additional training. This stands in contrast to existing forecasting models, for which optimisation often requires significant fine-tuning ex ante. Pre-trained foundation models have been shown to outperform state-of-the-art specialised forecasting models.

¹ Bok et al (2018). ² Sharpe et al (2023). ³ Jin et al (2023).

banks understand the extent of technology-induced job displacements, how quickly workers find new jobs and attendant wage dynamics. Similarly, satellite data on aerial pollution or nighttime lights can be used to predict short-term economic activity, while data on electricity consumption can shed light on industrial production in different regions and industries.²⁸ Central banks can thereby obtain a more nuanced picture of firms' capital expenditure and production, and how the supply of and demand for goods and services are changing.

Central banks can also use AI, together with human expertise, to better **understand factors that contribute to inflation**. Neural networks can handle more input variables compared with traditional econometric models, making it possible to work with detailed data sets rather than relying solely on aggregated data. They can further reflect intricate non-linear relationships, offering valuable insights during periods of rapidly changing inflation dynamics. If AI's impact varies by industry but materialises rapidly, such advantages are particularly beneficial for assessing inflationary dynamics.

Recent work in this area decomposes aggregate inflation into various sub-components.²⁹ In a first step, economic theory is used to pre-specify four factors shaping aggregate inflation: past inflation patterns, inflation expectations, the output gap and international prices. A neural network then uses aggregate series (eg the unemployment rate or total services inflation) and disaggregate series (eg two-digit industry output) to estimate the contribution of each of the four subcomponents to overall inflation, accounting for possible non-linearities.

The use of AI could play an important role in supporting **financial stability analysis**. The strongest suit of machine learning and AI methodologies is identifying patterns in a cross-section. As such, they can be particularly useful to identify and enhance the understanding of risks in a large sample of observations, helping identify the cross-section of risk across financial and non-financial firms. Again, availability of timely data is key. For example, during increasingly frequent periods of low liquidity and market dysfunction, AI could help prediction through better monitoring of anomalies across markets.³⁰

Finally, pairing AI-based insights with human judgment could help support **macroprudential regulation**. Systemic risks often result from the slow build-up of imbalances and vulnerabilities, materialising in infrequent but very costly stress events. The scarcity of data on such events and the uniqueness of financial crises limit the stand-alone use of data-intensive AI models in macroprudential regulation.³¹ However, together with human expertise and informed economic reasoning to see through the cycle, gen AI tools could yield large benefits to regulators and supervisors. When combined with rich data sets that provide sufficient scope to find patterns in the data, AI could help in building early warning indicators that alert supervisors to emerging pressure points known to be associated with system-wide risks.

In sum, with sufficient data, AI tools offer central banks an opportunity to get a much better understanding of economic developments. They enable central banks to draw on a richer set of structured and unstructured data, and complementarily, speed up data collection and analysis. In this way, the use of AI enables the analysis of economic activity in real time at a granular level. Such enhanced capabilities are all the more important in the light of AI's potential impact on employment, output and inflation, as discussed in the next section.

Macroeconomic impact of AI

AI is poised to increase productivity growth. For workers, recent evidence suggests that AI directly raises productivity in tasks that require cognitive skills (Graph 7.A). The use of generative AI-based tools has had a sizeable and rapid positive effect on the productivity of customer support agents and of college-educated professionals solving writing tasks. Software developers that used LLMs through the GitHub Copilot AI could code more than twice as many projects per week. A recent collaborative study by the BIS with Ant Group shows that productivity gains are immediate and largest among less experienced and junior staff (Box E).³²

Early studies also suggest positive effects of AI on firm performance. Patenting activity related to AI and the use of AI are associated with faster employment and output growth as well as higher revenue growth relative to comparable firms. Firms that adopt AI also experience higher growth in sales, employment and market valuations, which is primarily driven by increased product innovation. These effects have materialised over a horizon of one to two years. In a global sample, AI patent applications generate a positive effect on the labour productivity of small and medium-sized enterprises, especially in services industries.³³

Gen AI and labour productivity: a field experiment on coding

Generative artificial intelligence (gen AI) tools have the potential to enhance workers' productivity, but experimental evidence is scarce so far. To study the impact of gen AI on productivity, recent BIS work leverages data from Ant Group. In September 2023, Ant Group unveiled CodeFuse, a specific large language model (LLM) designed to assist software programmer teams in coding. Prior to its widespread release, this LLM was accessible for an initial six-week trial period only to a select group of programmers.


The evidence suggests that LLMs can boost productivity for coders. Comparing programmer groups with similar productivity levels and work experience but with or without access to the LLM shows a 55% increase in productivity (measured by the number of lines of code produced) on average for the group with access to the LLM. Roughly one third of this increase can be attributed directly to the code lines generated by the LLM, with the rest resulting from improved programmers' efficiency in coding elsewhere (likely reflecting additional time available for other programming tasks).

Dividing programmers by their experience levels reveals significant differences: productivity increased only among junior programmers (see Graphs E1.A and E1.B). Comparing the number of requests and acceptance rate of LLM suggestions by workers with different levels of experience sheds light on these differences, as senior programmers used the LLM less. The purple line in Graph E1.C indicates a negative correlation between the volume of requests following the LLM's introduction and the programmers' years of experience. At the same time, the acceptance rate (the frequency at which programmers used the LLM's suggestions) did not vary with the level of experience (orange line). These findings suggest that the lower impact of the LLM on senior programmers' productivity stems from their lower engagement with the LLM rather than a lack of usefulness.


Effect of generative AI on productivity for different levels of work experience

Graph E1

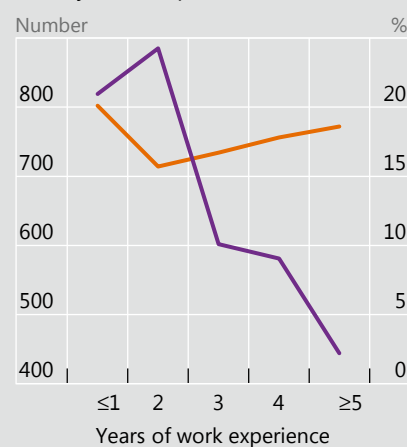
A. Junior programmers become much more productive...¹



B. ...whereas gains for senior programmers are more modest¹



C. Generative AI use and acceptance rate by work experience²



— / — Average effect

— 95% confidence interval

— Number of requests per user (lhs)
— Acceptance rate (rhs)

¹ Based on a difference-in-differences analysis to evaluate the effects on labour productivity between two groups of programmers: those with access to CodeFuse (treatment group) and those without (control group). The comparison spans six weeks before the introduction of CodeFuse and eight weeks afterwards (with time 0 marking the two-week introduction period). The y-axis approximates the growth rate in the number of lines of code produced (in logarithm). Junior programmers (panel A) are defined as those with up to one year of experience. Senior programmers (panel B) have more than one year of experience. ² The number of requests per user is determined by the average number of times a programmer, categorised by years of work experience, has requested assistance from the LLM in the eight weeks following its introduction. The acceptance rate represents the proportion of these requests for which a programmer has accepted the suggestions offered by the LLM application with less than 50% human modification.


Source: Gambacorta, Qiu, Rees and Shian (2024).

AI and productivity


In per cent

Graph 7

A. Generative AI makes workers more productive...¹



B. ...but higher-income or better-educated workers expect greater benefits²



¹ See technical annex for details. ² Based on 893 respondents from the Survey of Consumer Expectations conducted by the Federal Reserve Bank of New York.

Sources: Aldasoro, Armantier, Doerr, Gambacorta and Oliviero (2024a); Brynjolfsson et al (2023); Nielsen (2023); Noy and Zhang (2023); Peng et al (2024); Federal Reserve Bank of New York, *Survey of Consumer Expectations*.

The **macroeconomic impact of AI** on productivity growth could be sizeable. Beyond directly enhancing productivity growth by raising workers' and firms' efficiency, AI can spur innovation and thereby future productivity growth indirectly. Most innovation is generated in occupations that require high cognitive abilities. Improving the efficiency of cognitive work therefore holds great potential to generate further innovation. The estimates provided by the literature for AI's impact on annual labour productivity growth (ie output per employee) are thus substantive, although their range varies.³⁴ Through faster productivity growth, AI will expand the economy's productive capacity and thus raise aggregate supply.

Higher productivity growth will also affect aggregate demand through changes in firms' **investment**. While gen AI is a relatively new technology, firms are already investing heavily in the necessary IT infrastructure and integrating AI models into their operations – on top of what they already spend on IT in general. In 2023 alone, spending on AI exceeded \$150 billion worldwide, and a survey of US companies' technology officers across all sectors suggests almost 50% rank AI as their top budget item over the next years.³⁵

An additional boost to investment could come from **improved prediction**. AI adoption will lead to more accurate predictions at a lower cost, which reduces uncertainty and enables better decision-making.³⁶ Of course, AI could also introduce new sources of uncertainty that counteract some of its positive impact on firm investment, eg by changing market and price dynamics.

Another substantial part of aggregate demand is household **consumption**. AI could spur consumption by reducing search frictions and improving matching, making markets more competitive. For example, the use of AI agents could improve consumers' ability to search for products and services they want or need and help firms in advertising and targeting services and products to consumers.³⁷

AI's impact on household consumption will also depend on how it affects labour markets, notably labour demand and wages. The overall impact depends on the

relative strength of three forces (Graph 8): by how much AI raises productivity, how many new tasks it creates and how many workers it displaces by making existing tasks obsolete.

If AI is a true general-purpose technology that raises total factor productivity in all industries to a similar extent, the **demand for labour** is set to increase across the board (Graph 8, blue boxes). Like previous general-purpose technologies, AI could also **create altogether new tasks**, further increasing the demand for labour and spurring wage growth (green boxes). If so, AI would increase aggregate demand.

However, the **effects of AI might differ across tasks and occupations**. AI might benefit only some workers, eg those whose tasks require logical reasoning. Think of nurses who, with the assistance of AI, can more accurately interpret x-ray pictures. At the same time, gen AI could make other tasks obsolete, for example summarising documents, processing claims or answering standardised emails, which lend themselves to automation by LLMs. If so, increased AI adoption would lead to displacement of some workers (Graph 8, red boxes). This could lead to declines in employment and lower wage growth, with distributional consequences. Indeed, results from a recent survey of US households by economists in the BIS Monetary and Economic Department in collaboration with the Federal Reserve Bank of New York indicate that men, better-educated individuals or those with higher incomes think that they will benefit more from the use of gen AI than women and those with lower educational attainment or incomes (Graph 7.B).³⁸

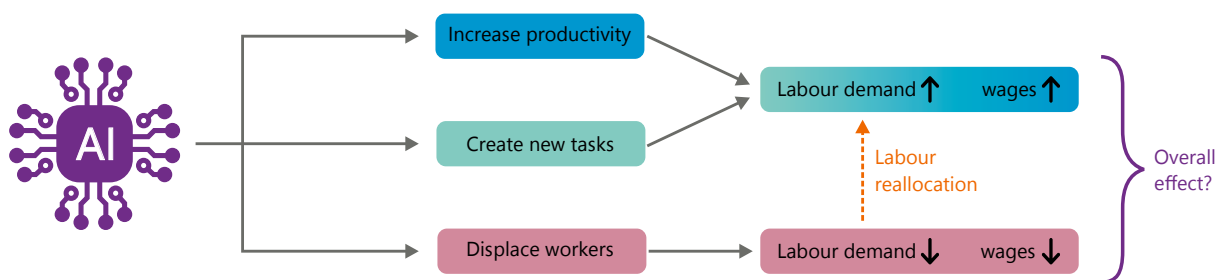
These considerations suggest that AI could have implications for **economic inequality**. Displacement might eliminate jobs faster than the economy can create new ones, potentially exacerbating income inequality. A differential impact of benefits across job categories would strengthen this effect. The “digital divide” could widen, with individuals lacking access to technology or with low digital literacy being further marginalised. The elderly are particularly at risk of exclusion.³⁹

Through the effects on productivity, investment and consumption the deployment of AI has implications for output and inflation. A BIS study illustrates the key mechanisms at work.⁴⁰ As the source of a permanent increase in productivity, AI will raise aggregate supply. An increase in consumption and investment raises aggregate demand. Through higher aggregate demand and supply, **output** increases (Graph 9.A). In the short term, if households and firms fully anticipate that they will be richer in the future, they will increase consumption at the expense of investment, slowing down output growth.

The response of **inflation** will also depend on households’ and businesses’ anticipation of future gains from AI. If the average household does not fully anticipate gains, it will increase today’s consumption only modestly. AI will act as a disinflationary

The impact of AI on labour demand and wages

Graph 8




Source: Aldasoro, Doerr, Gambacorta, Gelos and Rees (2024).

The impact of AI on output and inflation¹


Changes relative to the initial steady state, in per cent

Graph 9

A. AI will raise output in the short and long run...



B. ...but could be inflationary or disinflationary in the short run



— No anticipation — Full anticipation

¹ The vertical axis measures the change relative to the initial steady state value of output (panel A) and inflation (panel B). No anticipation refers to the case where households and firms do not anticipate the future effects of AI on productivity growth. Full anticipation refers to the case where households and firms perfectly anticipate the future effects of AI. For more details, see Aldasoro, Doerr, Gambacorta and Rees (2024).

Source: Adapted from Aldasoro, Doerr, Gambacorta and Rees (2024).

force in the short run (blue line in Graph 9.B), as the impact on aggregate supply dominates. In contrast, if households anticipate future gains, they will consume more, making AI's initial impact inflationary (red line in Graph 9.B). Since past general-purpose technologies have had an initial disinflationary impact, the former scenario appears more likely. But in either scenario, as economic capacity expands and wages rise, the demand for capital and labour will steadily increase. If these demand effects dominate the initial positive shock to output capacity over time, higher inflation would eventually materialise. How quickly demand forces increase output and prices will depend not only on households' expectations but also on the **mismatch in skills** required in obsolete and newly created tasks. The greater the skill mismatch (other things being equal), the lower employment growth will be, as it takes displaced workers longer to find new work. It might also be the case that some segments of the population will remain permanently unemployable without retraining. This, in turn, implies lower consumption and aggregate demand, and a longer disinflationary impact of AI.

Another aspect that warrants further investigation is the effect of AI adoption on **price formation**. Large retail companies that predominately sell online use AI extensively in their price-setting processes. Algorithmic pricing by these retailers has been shown to increase both the uniformity of prices across locations and the frequency of price changes.⁴¹ For example, when gas prices or exchange rates move, these companies quickly adjust the prices in their online stores. As the use of AI becomes more widespread, also among smaller companies, these effects could become stronger. Increased uniformity and flexibility in pricing can mean greater and quicker pass-through of aggregate shocks to local prices, and hence inflation, than in the past. This can ultimately change inflation dynamics. An important aspect to consider is how these effects could differ depending on the degree of competition in the AI model and data market, which could influence the variety of models used.

Finally, the impact of AI on **fiscal sustainability** remains an open question. All things equal, an AI-induced boost to productivity and growth would lead to a reduced debt burden. However, to the extent that faster growth is associated with higher interest rates, combined with the potential need for fiscal programmes to manage AI-induced labour relocation or sustained spells of higher unemployment rates, the impact of AI on the fiscal outlook might be modest. More generally, the AI growth dividend is unlikely to fully offset the spending needs that may arise from the green transition or population ageing over the next decades.

Toward an action plan for central banks

The use of AI models opens up new opportunities for central banks in pursuit of their policy objectives. A consistent theme running through the chapter has been the availability of data as a critical precondition for successful applications of machine learning and AI. Data governance frameworks will be part and parcel of any successful application of AI. Central banks' policy challenges thus encompass both **models** and **data**.

An important trade-off arises between **using "off-the-shelf" models versus developing in-house fine-tuned ones**. Using external models may be more cost-effective, at least in the short run, and leverages the comparative advantage of private sector companies. Yet reliance on external models comes with reduced transparency and exposes central banks to concerns about dependence on a few external providers. Beyond the general risks that market concentration poses to innovation and economic dynamism, the high concentration of resources could create significant operational risks for central banks, potentially affecting their ability to fulfil their mandates.

Another important aspect relates to central banks' role as **users, compilers and disseminators of data**. Central banks use data as a crucial ingredient in their decision-making and communication with the public. And they have always been extensive compilers of data, either collecting them on their own or drawing on other official agencies and commercial sources. Finally, central banks are also providers of data, to inform other parts of government as well as the general public. This role helps them fulfil their obligations as key stakeholders in national statistical systems.

The rise of machine learning and AI, together with advances in computing and storage capacity, have cast these aspects in an urgent new light. For one, central banks now need to make sense of and use increasingly large and diverse sets of structured and unstructured data. And these data often reside in the hands of the private sector. While LLMs can help process such data, hallucinations or prompt injection attacks can lead to biased or inaccurate analyses. In addition, commercial data vendors have become increasingly important, and central banks make extensive use of them. But in recent years, the cost of commercial data has increased markedly, and vendors have imposed tighter use conditions.

The decision on whether to use external or internal models and data has far-reaching implications for central banks' investments and human capital. A key challenge is **setting up the necessary IT infrastructure**, which is greater if central banks pursue the road of developing internal models and collecting or producing their own data. Providing adequate computing power and software, as well as training existing or hiring new staff, involves high up-front costs. The same holds for creating a data lake, ie pooling different curated data sets. Yet a reliable and safe IT infrastructure is a prerequisite not only for big data analyses but also to prevent cyber attacks.

Hiring new or retaining existing staff with the right mix of economic understanding and programming skills can be challenging. As AI applications

increase the sophistication of the financial system over time, the premium on having the right mix of skills will only grow. Survey-based evidence suggests this is a top concern for central banks (Graph 10). There is high demand for data scientists and other AI-related roles, but public institutions often cannot match private sector salaries for top AI talent. The need for staff with the right skills also arises from the fact that the use of AI models to aid financial stability monitoring faces limitations, as discussed above. Indeed, AI is not a substitute for human judgment. It requires supervision by experts with a solid understanding of macroeconomic and financial processes.

How can central banks address these challenges and mitigate trade-offs? The answer lies, in large part, in cooperation paired with sound data governance practices.

Collaboration can yield significant benefits and relax constraints on human capital and IT. For one, the **pooling of resources and knowledge** can lower demands among central banks and could ease the resource constraints on collecting, storing and analysing big data as well as developing algorithms and training models. For example, central banks could address rising costs of commercial data, especially for smaller institutions, by sharing more granular data themselves or by acquiring data from vendors through joint procurement. Cooperation could also facilitate training staff through workshops in the use of AI or the sharing of experiences in conferences. This would particularly benefit central banks with fewer staff and resources and with limited economies of scale. Cooperation, for example by re-using trained models, could also mitigate the environmental costs associated with training algorithms and storing large amounts of data, which consume enormous amounts of energy.

Central bank collaboration and the sharing of experiences could also help identify areas in which AI adds the most value and how to **leverage synergies**.


Challenges in recruitment and retention

As a percentage of respondent central banks

Graph 10

A. Has recruitment or retention become more difficult in the past five years?¹

B. For which areas or skills has recruitment/retention become more difficult?²



¹ Based on a survey of 52 members of the Central Bank Governance Network conducted in May 2024. ² Shares based on the subset of central banks experiencing more difficulties with recruitment and/or retention.

Source: CBGN (2024).

Common data standards could facilitate access to publicly available data and facilitate the automated collection of relevant data from various official sources, thereby enhancing the training and performance of machine learning models. Additionally, dedicated repositories could be set up to share the open source code of data tools, either with the broader public or, at least initially, only with other central banks. An example is a platform such as BIS Open Tech, which supports international cooperation and coordination in sharing statistical and financial software. More generally, central banks could consider sharing domain-adapted or fine-tuned models in the central banking community, which could significantly lower the hurdles for adoption.⁴² Joint work on AI models is possible without sharing data, so they can be applied even where there are concerns about confidentiality.

An example of how collaboration supports data collection and dissemination is the jurisdiction-level statistics on international banking, debt securities and over-the-counter derivatives by the BIS. These data sets have a long history – the international banking statistics started in the 1970s. They are a critical element for monitoring developments and risks in the global financial system. They are compiled from submissions by participating authorities under clear governance rules and using well established statistical processes. At a more granular level, arrangements for the sharing of confidential bank-level data include the quantitative impact study data collected by the Basel Committee on Banking Supervision and the data on large global banks collected by the International Data Hub. Other avenues to explore include sharing synthetic or anonymised data that protect confidential information.

The rising importance of data and emergence of new sources and tools call for **sound data governance practices**. Central banks must establish robust governance frameworks that include guidelines for selecting, implementing and monitoring both data and algorithms. These frameworks should comprise adequate quality control and cover data management and auditing practices. The importance of metadata, in particular, increases as the range and variety of data expand. Sometimes referred to as “the data about the data”, metadata include the definitions, source, frequency, units and other information that define a given data set. This metadata is crucial when privacy-preserving methods are used to draw lessons from several data sets overseen by different central banks. Machine readability is greatly enhanced when metadata are standardised so that the machines know what they are looking for. For example, the “Findable, Accessible, Interoperable and Reusable” (FAIR) principles provide guidance in organising data and metadata to ease the burden of sharing data and algorithms.⁴³

More generally, metadata frameworks are crucial for a better understanding of the comparability and limits of data series. Central banks can also cooperate in this domain. For example, the Statistical Data and Metadata Exchange (SDMX) standard provides a common language and structure for metadata. Such standards are crucial to foster data-sharing, lower the reporting burden and facilitate interoperability. Similarly, the Generic Statistical Business Process Model lays out business processes for official statistics with a unified framework and consistent terminology. Sound data governance practices would also facilitate the sharing of confidential data.

In sum, there is an urgent need for central banks to collaborate in fostering the development of a **community of practice** to share knowledge, data, best practices and AI tools. In the light of rapid technological change, the exchange of information on policy issues arising from the role of central banks as data producers, users and disseminators is crucial. Collaboration lowers costs, and such a community would foster the development of common standards. Central banks have a history of successful collaboration to overcome new challenges. The emergence of AI has hastened the need for cooperation in the field of data and data governance.

Endnotes

- ¹ “Structured data” refers to organised, quantitative information that is stored in relational databases and is easily searchable, such as categorical and numeric information. In contrast, unstructured data are not organised based on pre-defined models and can include information such as audio, video, emails, presentations, satellite images, etc.
- ² For an analysis of how gen AI impacts household and firm behaviour, see Aldasoro, Armantier, Doerr, Gambacorta and Oliviero (2024a) and McKinsey & Company (2023).
- ³ For technical and non-technical introductions to AI, see Russell and Nordvig (2021) and Mitchell (2019), respectively. For a definition of machine learning, see Murphy (2012).
- ⁴ See Park et al (2024).
- ⁵ See Belkin et al (2019).
- ⁶ See Aldasoro, Gambacorta, Korinek, Shreeti and Stein (2024). An important (and open) question is whether increasing (or “scaling”) the number of parameters and the amount of input data in training AI models will continue to deliver proportional gains in capabilities – the so-called scaling hypothesis; see Korinek and Vipra (2024).
- ⁷ See Perez-Cruz and Shin (2024).
- ⁸ See Bender and Koller (2020), Browning and LeCun (2022), Bubeck et al (2023) and Wei et al (2022).
- ⁹ Common measures of exposure, as computed for example in Felten et al (2021) and Aldasoro, Doerr, Gambacorta and Rees (2024), show that the financial sector ranks highest in exposure to AI.
- ¹⁰ See BCBS (2024).
- ¹¹ On the decline of correspondent banking, see Rice et al (2020). For an analysis of the negative real effects suffered by jurisdictions severed from the network of correspondent banking relationships, see Borchert et al (2023).
- ¹² See BCBS (2024).
- ¹³ See Doerr et al (2023) for a discussion on alternative data, Berg et al (2022) for evidence on default prediction, Cornelli, Frost, Gambacorta, Rau, Wardrop and Ziegler (2023) and Di Maggio et al (2023) for evidence on the use of alternative data by fintechs and big techs to detect “invisible primes”, and Gambacorta, Huang, Qiu and Wang (2024) for evidence on financial inclusion.
- ¹⁴ On asset embeddings, see Zhu et al (2023) and Gabaix et al (2023). On fees, see OECD (2021).
- ¹⁵ See Doerr, Gambacorta, Leach, Legros and Whyte (2022).

- ¹⁶ See Hitaj et al (2022).
- ¹⁷ For example, there is evidence from machine learning-based credit scoring models that, in the US mortgage market, Black and Hispanic borrowers were less likely to benefit from lower interest rates than borrowers from other communities; see Fuster et al (2019).
- ¹⁸ Distrust in big techs to safely handle data, relative to traditional financial institutions and the government, has been shown to also exist in other countries (Chen et al (2023)). A related risk is that alternative data are correlated with certain consumer characteristics that lenders are, for good reason, not allowed to use in their credit assessment process (eg gender or minority status). Moreover, gen AI could exacerbate and perpetuate biases by creating biased data itself (either because of biased training data or hallucination), which are then used by other models.
- ¹⁹ For a general policy discussion of financial stability risks from AI and machine learning, see Hernández de Cos (2024), Gensler and Bailey (2020), and OECD (2023). Calvano et al (2020) present an academic treatment of AI, algorithmic pricing and collusion. Assad et al (2024) find evidence of collusion by pricing algorithms in retail gasoline markets. More specifically for financial markets, Georges and Pereira (2021) find that even if traders pay a lot of attention to model selection, the risk of destabilising speculation cannot be entirely eliminated.
- ²⁰ See Doerr et al (2021) and Araujo et al (2024) for more information on the use of big data and AI in central banks.
- ²¹ See Desai et al (2024) for a recent application of machine learning to anomaly detection in payment systems.
- ²² See Garratt et al (2024).
- ²³ See Carstens and Nilekani (2024) for details on the Finternet. See Aldasoro et al (2023) for a primer on tokenisation.
- ²⁴ See Aldasoro, Doerr, Gambacorta, Notra, Oliviero and Whyte (2024), who investigate the link between gen AI and cyber risk in central banks by drawing on the results of an ad hoc survey of members of the Global Cyber Resilience Group in January 2024.
- ²⁵ See a recent report by Bernanke (2024), which argues that central banks need to improve their forecasting abilities by gathering more timely data and integrating advanced data analytics.
- ²⁶ See Du et al (2024).
- ²⁷ See Barlas et al (2021) for a discussion on using transaction-level data for nowcasting consumption and investment. The Beige Book aggregates narratives that are collected from business contacts to summarise the economic conditions of each of the 12 Federal Reserve districts. See Soto (2023).
- ²⁸ See Bricongne et al (2021), Dasgupta (2022) and Lehman and Möhrle (2022).
- ²⁹ See Buckman et al (2023).

- ³⁰ See Aquilina et al (2024).
- ³¹ Synthetic data are unlikely to help. In most cases, their generation relies on known data-generating processes, which do not apply to financial crises. And continued reliance on these data by AI models diminishes the information coming from the tails of the distribution (ie rare but highly consequential events), which are of particular concern for macroprudential policy. See Shumailov et al (2023).
- ³² Brynjolfsson et al (2023) document that access to a gen AI-based conversational assistant improved customer support agents' productivity by 14%. Noy and Zhang (2023) find that support by the chatbot ChatGPT raised productivity in solving writing tasks for college-educated professionals from a variety of fields, reducing the time required by 40% and raising output quality by 18%. For evidence on productivity improvements in coding, see Gambacorta, Qiu, Rees and Shian (2024) and Peng et al (2024).
- ³³ On employment and output growth, see Yang (2022) and Czarnitzki et al (2023) for Taiwan and Germany, respectively. For revenue growth, see Alderucci et al (2020). Babina et al (2024) present evidence for product innovation, whereas Damoli et al (2021) do so for labour productivity.
- ³⁴ On innovation, see Brynjolfsson et al (2018). Estimates range from 0.5 to 1.5 percentage points over the next decade; see eg Baily et al (2023) and Goldman Sachs (2023). Acemoglu (2024) provides lower yet still positive estimates.
- ³⁵ See Statista and Anwah and Rosenbaum (2023).
- ³⁶ See Agrawal et al (2019, 2022) as well as Ahir et al (2022).
- ³⁷ "AI agents" refers to systems that build on advanced LLMs and are endowed with planning capabilities, long-term memory and, typically, access to external tools (eg the ability to execute code, use the internet or perform market trades).
- ³⁸ See Aldasoro, Armantier, Doerr, Gambacorta and Oliviero (2024a, b). See also Pizzinelli et al (2023) on the effects of AI on labour markets and inequality.
- ³⁹ Cornelli, Frost and Mishra (2023) find that investments in AI are associated with greater inequality and a shift from mid-skill jobs to high-skill and managerial positions. On the digital divide, see Doerr, Frost, Gambacorta and Qiu (2022). For more details on the impact of AI on income inequality, see Cazzaniga et al (2024).
- ⁴⁰ See Aldasoro, Doerr, Gambacorta and Rees (2024).
- ⁴¹ See Cavallo (2019).
- ⁴² See Gambacorta, Kwon, Park, Patelli and Zhu (2024).
- ⁴³ See Wilkinson et al (2016).

Technical annex

Graph 1.A: The adoption of ChatGPT is proxied by the ratio of the maximum number of website visits worldwide for the period November 2022–April 2023 and the worldwide population with internet connectivity. For more details on computer see US Census Bureau; for electric power, internet and social media see Comin and Hobijn (2004) and Our World in Data; for smartphones, see Statista.

Graph 1.B: Based on an April 2023 global survey with 1,684 participants.

Graph 1.C: Data for capital invested in AI companies for 2024 are annualised based on data up to mid-May. Data on the percentage of AI job postings for AU, CA, GB, NZ and US are available for the period 2014–23; for AT, BE, CH, DE, ES, FR, IT, NL and SE, data are available for the period 2018–23.

Graph 2.A: Three-month moving averages.

Graphs 2.B and 2.C: Correspondent banks that are active in several corridors are counted several times. Averages across countries in each region. Markers in panel C represent subregions within each region. Grouping of countries by region according to the United Nations Statistics Division; for further details see unstats.un.org/unsd/methodology/m49/.

Graph 3.A: Average scores in answers to the following question: “In the following areas, would you trust artificial intelligence (AI) tools less or more than traditional human-operated services? For each item, please indicate your level of trust on a scale from 1 (much less trust than in a human) to 7 (much more trust).”

Graph 3.B: Average scores and 95% confidence intervals in answers to the following question: “How much do you trust the following entities to safely store your personal data when they use artificial intelligence tools? For each of them, please indicate your level of trust on a scale from 1 (no trust at all in the ability to safely store personal data) to 7 (complete trust).”

Graph 3.C: Average scores (with scores ranging from 1 (lowest) to 7 (highest)) in answers to the following questions: (1) “Do you think that sharing your personal information with artificial intelligence tools will decrease or increase the risk of data breaches (that is, your data becoming publicly available without your consent)?”; (2) “Are you concerned that sharing your personal information with artificial intelligence tools could lead to the abuse of your data for unintended purposes (such as for targeted ads)?”.

Graph 6.A: The bars show the share of respondents to the question, “Do you agree that the use of AI can provide more benefits than risks to your organisation?”.

Graph 6.B: The bars show the average score that respondents gave to each option when asked to “Rate the level of significance of the following benefits of AI in cyber security”; the score scale of each option is from 1 (lowest) to 5 (highest).

Graph 7.A: The bars correspond to estimates of the increase in productivity of users that rely on generative AI tools relative to a control group that did not.

References

Acemoglu, D (2024): "The simple macroeconomics of AI", *Economic Policy*, forthcoming.

Agrawal, A, J Gans and A Goldfarb (2019): "Exploring the impact of artificial intelligence: prediction versus judgment", *Information Economics and Policy*, vol 47, pp 1–6.

——— (2022): *Prediction machines, updated and expanded: the simple economics of artificial intelligence*, Harvard Business Review Press, 15 November.

Ahir, H, N Bloom and D Furceri (2022): "The world uncertainty index", *NBER Working Papers*, no 29763, February.

Aldasoro, I, O Armantier, S Doerr, L Gambacorta and T Oliviero (2024a): "Survey evidence on gen AI and households: job prospects amid trust concerns", *BIS Bulletin*, no 86, April.

——— (2024b): "The gen AI gender gap", *BIS Working Papers*, forthcoming.

Aldasoro, I, S Doerr, L Gambacorta, R Garratt and P Koo Wilkens (2023): "The tokenisation continuum", *BIS Bulletin*, no 72, April.

Aldasoro, I, S Doerr, L Gambacorta, G Gelos and D Rees (2024): "Artificial intelligence, labour markets and inflation", *mimeo*.

Aldasoro, I, S Doerr, L Gambacorta, S Notra, T Oliviero and D Whyte (2024): "Generative artificial intelligence and cybersecurity in central banking", *BIS Papers*, no 145, May.

Aldasoro, I, S Doerr, L Gambacorta and D Rees (2024): "The impact of artificial intelligence on output and inflation", *BIS Working Papers*, no 1179, April.

Aldasoro, I, L Gambacorta, A Korinek, V Shreeti and M Stein (2024): "Intelligent financial system: how AI is transforming finance", *BIS Working Papers*, no 1194, June.

Alderucci, D, L Branstetter, E Hovy, A Runge and N Zolas (2020): "Quantifying the impact of AI on productivity and labor demand: evidence from US census microdata", paper presented at the Allied Social Science Associations 2020 Annual Meeting.

Allcot, D (2023): "How to invest in artificial intelligence in 2023", *GOBankingRates*, 21 June.

Anwah, O and E Rosenbaum (2023): "AI is now the biggest spend for nearly 50% of top tech executives across the economy: CNBC survey", *CNBC*, 23 June.

Aquilina, M, D Araujo, G Gelos, T Park and F Perez-Cruz (2024): "Harnessing artificial intelligence for monitoring financial markets", *BIS Working Papers*, forthcoming.

Araujo, D, S Doerr, L Gambacorta and B Tissot (2024): "Artificial intelligence in central banking", *BIS Bulletin*, no 84, January.

Assad, S, R Clark, D Ershov and L Xu (2024): "Algorithmic pricing and competition: empirical evidence from the German retail gasoline market", *Journal of Political Economy*, vol 132, no 3.

Babina, T, A Fedyk, A He and J Hodson (2024): "Artificial intelligence, firm growth and product innovation", *Journal of Financial Economics*, vol 151, 103745.

Baily, M, E Brynjolfsson and A Korinek (2023): "Machines of mind: the case for an AI-powered productivity boom", Brookings, 10 May.

Barlas, A, S G Mert, B O Isa, A Ortiz, T Rodrigo, B Soybilgen and E Yazgan (2021): "Big data information and nowcasting: consumption and investment from bank transactions in Turkey", [arXiv:2107.03299v1](https://arxiv.org/abs/2107.03299v1).

Basel Committee on Banking Supervision (BCBS) (2024): *Digitalisation of finance*, May.

Belkin, M, D Hsu, S Ma and S Mandal (2019): "Reconciling modern machine learning practice and the classical bias-variance trade-off", *Proceedings of the National Academy of Sciences*, vol 116, no 32, pp 15849–54.

Bender, E and A Koller (2020): "Climbing towards NLU: on meaning, form and understanding in the age of data", in D Jurafsky, J Chai, N Schluter and J Tetreault (eds), *Proceedings of the 58th annual meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, July, pp 5185–98.

Berg, T, A Fuster and M Puri (2022): "Fintech lending", *Annual Review of Financial Economics*, vol 14, pp 182–207.

Bernanke, B (2024): *Forecasting for monetary policy making and communication at the Bank of England: a review*, April.

BIS Innovation Hub (2023): *Project Aurora: the power of data, technology and collaboration to combat money laundering across institutions and borders*, May.

Bok, B, D Caratelli, D Giannone, A Sbordone and A Tambalotti (2018): "Macroeconomic nowcasting and forecasting with big data", *Annual Review of Economics*, vol 10, pp 615–43.

Borchert, L, R De Haas, K Kirschenmann and A Schultz (2023): "Broken relationships: de-risking by correspondent banks and international trade", *EBRD Working Papers*, no 285, December.

Bricongne J-C, B Meunier and T Pical (2021): "Can satellite data on air pollution predict industrial production?", *Bank of France Working Papers*, no 847.

Brown, T et al (2019): "Language models are few-shot learners", in H Larochelle, M Ranzato, R Hadsell, M Balcan and H Lin (eds), *Advances in neural information processing systems 32: 33rd conference on neural information processing systems (NeurIPS 2019)*, Neural Information Processing Systems Foundation.

Browning, J and Y LeCun (2022): "AI and the limits of language", *Noema*, 23 August.

Brynjolfsson, E, D Li and L Raymond (2023): "Generative AI at work", *NBER Working Papers*, no 31161.

Brynjolfsson, E, D Rock and C Syverson (2018): "Artificial intelligence and the modern productivity paradox: a clash of expectations and statistics", in A Agrawal, J Gans and A Goldfarb (eds), *The economics of artificial intelligence: an agenda*, University of Chicago Press, pp 23–57.

Bubeck, S, V Chandrasekaran, R Eldan, J Gehrke, E Horvitz, E Kamar, P Lee, Y T Lee, Y Li, S Lundberg, H Nori, H Palangi, M Ribeiro and Y Zhang (2023): "Sparks of artificial general intelligence: early experiments with GPT-4", arxiv.org/abs/2303.12712.

Buckman, M, G Potjagailo and P Schnattinger (2023): "Dissecting UK service inflation via a neural network Phillips curve", *Bank Underground*, 10 July.

Calvano, E, G Calzolari, V Denicolò and S Pastorello (2020): "Artificial intelligence, algorithmic pricing and collusion", *American Economic Review*, vol 110, no 10, pp 3267–97.

Carstens, A and N Nilekani (2024): "Finternet: the financial system for the future", *BIS Working Papers*, no 1178.

Cavallo, A (2019): "More Amazon effects: online competition and pricing behaviors", *Jackson Hole Economic Symposium Conference Proceedings*.

Cazzaniga, M, F Jaumotte, L Li, G Melina, A Panton, C Pizzinelli, E Rockall and M Tavares (2024): "Gen-AI: artificial intelligence and the future of work", *International Monetary Fund Staff Discussion Notes*, no 1, January.

Central Bank Governance Network (CBGN) (2024): Survey on talent management and new ways of working at central banks.

Chen, S, S Doerr, J Frost, L Gambacorta and H S Shin (2023): "The fintech gender gap", *Journal of Financial Intermediation*, vol 54.

Comin, D and B Hobijn (2004): "Cross-country technology adoption: making the theories face the facts", *Journal of Monetary Economics*, vol 51, no 1, pp 39–83.

Committee on Payments and Market Infrastructures (CPMI) (2023): *CPMI quantitative review of correspondent banking data*, May.

Cornelli, G, J Frost, L Gambacorta, R Rau, R Wardrop, and T Ziegler (2023): "Fintech and big tech credit: drivers of the growth of digital lending", *Journal of Banking and Finance*, vol 148, March.

Cornelli, G, J Frost and S Mishra (2023): "Artificial intelligence, services globalisation and income inequality", *BIS Working Papers*, no 1135, October.

Czarnitzki, D, G P Fernández, and C Rammer (2023): "Artificial intelligence and firm-level productivity", *Journal of Economic Behavior & Organization*, vol 211, pp 188–205.

Damioli, G, V Van Roy and D Vertesy (2021): "The impact of artificial intelligence on labor productivity", *Eurasian Business Review*, vol 11, pp 1–25.

Dasgupta, N (2022): "Using satellite images of nighttime lights to predict the economic impact of Covid-19 in India", *Advances in Space Research*, vol 70, no 4, pp 863–79.

Desai, A, A Kosse and J Sharples (2024): "Finding a needle in a haystack: a machine learning framework for anomaly detection in payment systems", *BIS Working Papers*, no 1188, May.

Devlin, J, M W Chang, K Lee and K Toutanova (2018): "BERT: pre-training of deep bidirectional transformers for language understanding", *NAACL-HLT*, vol 1 (2019), pp 4171–86, [arXiv:1810.04805](https://arxiv.org/abs/1810.04805).

Di Maggio, M, D Ratnadiwakara and D Carmichael (2023): "Invisible primes: fintech lending with alternative data", *NBER Working Papers*, no 29840.

Doerr, S, J Frost, L Gambacorta and H Qiu (2022): "Population ageing and the digital divide", *SUERF Policy Briefs*, no 270.

Doerr, S, L Gambacorta, L Guiso and M Sanchez del-Villar (2023): "Privacy regulation and fintech lending", *BIS Working Papers*, no 1103, June.

Doerr, S, L Gambacorta, T Leach, B Legros and D Whyte (2022): "Cyber risk in central banking", *BIS Working Papers*, no 1039, September.

Doerr, S, L Gambacorta and J-M Serena (2021): "Big data and machine learning in central banking", *BIS Working Papers*, no 930, March.

Du S, K Guo, F Haberkorn, A Kessler, I Kitschelt, S J Lee, A Monken, D Saez, K Shipman and S Thakur (2024): "Do anecdotes matter? Exploring the Beige Book through textual analysis from 1970 to 2023", *IFC Bulletin*, forthcoming.

Felten, E, M Raj and R Seamans (2021): "Occupational, industry and geographic exposure to artificial intelligence: a novel data set and its potential uses", *Strategic Management Journal*, vol 42, no 12, pp 2195–217.

Fuster, A, M Plosser, P Schnabl and J Vickery (2019): "The role of technology in mortgage lending", *The Review of Financial Studies*, vol 32, no 5, pp 1854–99.

Gabaix, X, R Koijen, R Richmond and M Yogo (2023): "Asset embeddings", mimeo.

Gambacorta, L, Y Huang, H Qiu and J Wang (2024): "How do machine learning and non-traditional data affect credit scoring? New evidence from a Chinese fintech firm", *Journal of Financial Stability*, forthcoming.

Gambacorta, L, B Kwon, T Park, P Patelli and S Zhu (2024): "CB-LMs: language models for central banking", *BIS Working Papers*, forthcoming.

Gambacorta, L, H Qiu, D Rees and S Shian (2024): "Generative AI and labour productivity: a field experiment on code programming", mimeo.

Garratt, R, P Koo Wilkens and H S Shin (2024): "Next generation correspondent banking", *BIS Bulletin*, no 87, May.

- Gensler, G and L Bailey (2020): "Deep learning and financial stability", mimeo.
- Georges, C and J Pereira (2021): "Market stability with machine learning agents", *Journal of Economic Dynamics and Control*, vol 122.
- Goldman Sachs (2023): "Generative AI: hype or truly transformative?", *Global Macro Research*, no 120.
- Grand, G, I Blank, F Pereira and E Fedorenko (2022): "Semantic projection recovers rich human knowledge of multiple object features from word embeddings", *Nature Human Behaviour*, vol 6, pp 975–87.
- Hernández de Cos, P (2024): "Managing AI in banking: are we ready to cooperate?", keynote speech at the Institute of International Finance Global Outlook Forum, Washington DC, 17 April.
- Hitaj, D, G Pagnotta, B Hitaj, L Mancini and F Perez-Cruz (2022): "MaleficNet: hiding malware into deep neural networks using spread-spectrum channel coding", in V Atluri, R Di Pietro, C Jensen and W Meng (eds), *Computer security – ESORICS 2022*, Springer Cham, 23 September.
- Jin, M, S Wang, L Ma, Z Chu, J Zhang, X Shi, P-Y Chen, Y Liang, Y-F Li, S Pan and Q Wen (2023): "Time-LLM: time series forecasting by reprogramming large language models", [arXiv:2310.01728](https://arxiv.org/abs/2310.01728).
- Korinek, A and J Vipra (2024): "Concentrating intelligence: scaling laws and market structure in generative AI", prepared for the 79th Economic Policy panel meeting, Brussels, April.
- Lehman, R and S Möhrle (2022): "Forecasting regional industrial production with high-frequency electricity consumption data", *CESifo Working Papers*, no 9917, August.
- Maslej, N, L Fattorini, R Perrault, V Parli, A Reuel, E Brynjolfsson, J Etchemendy, K Ligett, T Lyons, J Manyika, J Niebles, Y Shoham, R Wald and J Clark (2024): *Artificial intelligence index report 2024*, AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford, CA, April.
- McKinsey & Company (2023): "The state of AI in 2023: generative AI's breakout year", August.
- Mikolov, T, K Chen, G Corrado and J Dean (2013): "Efficient estimation of word representations in vector space", International Conference on Learning Representations.
- Mitchell, M (2019): *Artificial intelligence: a guide for thinking humans*, Picador Paper Publishing, 15 October.
- Murphy, K (2012): *Machine learning: a probabilistic perspective*, MIT Press, 24 August.
- Nielsen, J (2023): "AI improves employee productivity by 66%", Nielsen Norman Group.

Noy, S and W Zhang (2023): "Experimental evidence on the productivity effects of generative artificial intelligence", *Science*, vol 381, no 6654, pp 187–92.

Organisation for Economic Co-operation and Development (OECD) (2021): "Artificial intelligence, machine learning and big data in finance: opportunities, challenges and implications for policy makers".

——— (2023): "Generative artificial intelligence in finance", *OECD Artificial Intelligence Papers*, no 9.

Park T, H S Shin and H Williams (2024): "Mapping the space of economic ideas with LLMs", *BIS Working Papers*, forthcoming.

Peng, S, W Swiatek, A Gao, P Cullivan and H Chang (2024): "AI revolution on chat bot: evidence from a randomized controlled experiment", [arXiv:2401.10956](https://arxiv.org/abs/2401.10956).

Pennington J, R Socher and C Manning (2014): "GloVe: global vectors for word representation", in A Moschitti, B Pang and W Daelemans (eds), *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, Association for Computational Linguistics, pp 1532–43.

Perez-Cruz, F and H S Shin (2024): "Testing the cognitive limits of large language models", *BIS Bulletin*, no 83, January.

Peters, M, M Neumann, M Iyyer, M Gardner, C Clark, K Lee and L Zettlemoyer (2018): "Deep contextualized word representations", in M Walker, H Ji and A Stent (eds), *Proceedings of the 2018 conference of the North American chapter of the Association for Computational Linguistics: human language technologies*, vol 1, pp 2227–37.

Pizzinelli C, A Panton, M Tavares, M Cazzaniga and L Li (2023): "Labor market exposure to AI: cross-country differences and distributional implications", *International Monetary Fund*, October.

Radford, A, J Wu, R Child, D Luan, D Amodei and I Sutskever (2018): "Language models are unsupervised multitask learners", mimeo.

Rice, T, G von Peter and C Boar (2020): "On the global retreat of correspondent banks", *BIS Quarterly Review*, March, pp 37–52.

Russell, S and P Norvig (2021): *Artificial intelligence: a modern approach*, Pearson, 4th edition.

Sharpe, S, N Sinha and C Hollrah (2023): "The power of narrative sentiment in economic forecasts", *International Journal of Forecasting*, vol 39, no 3, pp 1097–121.

Shumailov, I, Z Shumaylov, Y Zhao, Y Gal, N Papernot and R Anderson (2023): "The curse of recursion: training on generated data makes models forget", [arXiv:2305.17493v3](https://arxiv.org/abs/2305.17493v3).

Soto P (2023): "Measurement and effects of supply chain bottlenecks using natural language processing", *FEDS Notes*, 6 February.

Vaswani, A, N Shazeer, N Parmar, J Uszkoreit, L Jones, A Gomez, Ł Kaiser and I Polosukhin (2017): "Attention is all you need", in I Guyon, U Von Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan and R Garnett (eds), *Advances in neural information processing systems 30: 31st annual conference on neural information processing systems (NIPS 2017)*, Neural Information Processing Systems Foundation.

Vig, J (2019): "A multiscale visualization of attention in the transformer model", in M Costa-Jussà and E Alfonseca (eds), *Proceedings of the 57th annual meeting of the Association for Computational Linguistics: system demonstrations*, Association for Computational Linguistics, pp 37–42.

Wei, J, Y Tay, R Bommasani, C Raffel, B Zoph, S Borgeaud, D Yogatama, M Bosma, D Zhou, D Metzler, E Chi, T Hashimoto, O Vinyals, P Liang, J Dean and W Fedus (2022): "Emergent abilities of large language models", *Transactions on Machine Learning Research*, August.

Wilkinson, M et al (2016): "The FAIR guiding principles for scientific data management and stewardship", *Scientific Data*, no 3.

Yang, C H (2022): "How artificial intelligence technology affects productivity and employment: firm-level evidence from Taiwan", *Research Policy*, vol 51, no 6.

Zhu L, H Wu and M Wells (2023): "News-based sparse machine learning models for adaptive asset pricing", *Data Science in Science*, vol 2, no 1.