Artificial Intelligence and the economy: implications for central banks

Speech by Hyun Song Shin
Economic Adviser and Head of Research, Bank for International Settlements

on the occasion of the Bank's Annual General Meeting
in Basel on 30 June 2024

The rise of artificial intelligence

Artificial intelligence (AI) has taken the world by storm and set off a gold rush across the economy, with an unprecedented pace of adoption and investment in the technology. This year's special chapter discusses the impact of AI on the financial sector and the real economy and lays out the implications for central banks.

The technology behind AI can be traced back to the early days of computing itself. But it was the advent of deep learning in the 2010s, based on the combination of massive amounts of data and computing power, that set the stage for today's AI applications. Today's machine learning models excel at imposing mathematical structure on unstructured data, to identify patterns of interest in vast amounts of data. (Graph 1).

Turning unstructured into structured data — Graph 1

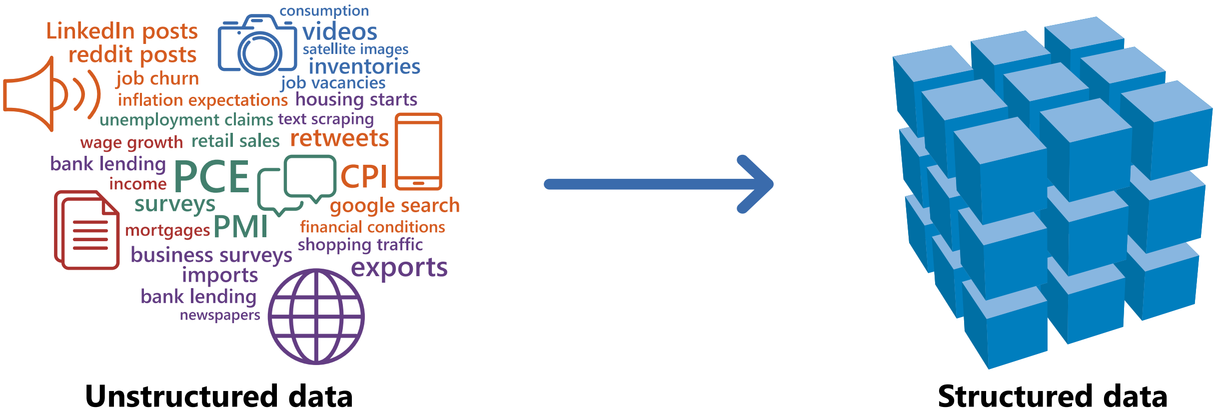

Source: BIS.

One of the leading machine learning methods is the "embedding" of words in a vector space so that words become arrays of numbers (Graph 2). The vectors preserve meaning, in that similar words are closer together in the vector space. For example, the word "turtle" is close to "hawk" and "alligator", but far from words for sports or for the weather.

The latest models take this concept one step further. Thanks to so-called transformer architecture, they take account of the surrounding context in the embedding of a word rather than having just one embedding for each word. Think of the word "bond". It could refer to a fixed income security, a connection or link, or 007 James Bond. Depending on the context, the vector embedding for the word "bond" will be different. It could lie geometrically close to words such as "treasury" and "yield", or to words such as "family" and "cultural", or perhaps to words such as "spy" and "martini".

Embedding turns words into arrays of numbers; similar words are closer together — Graph 2

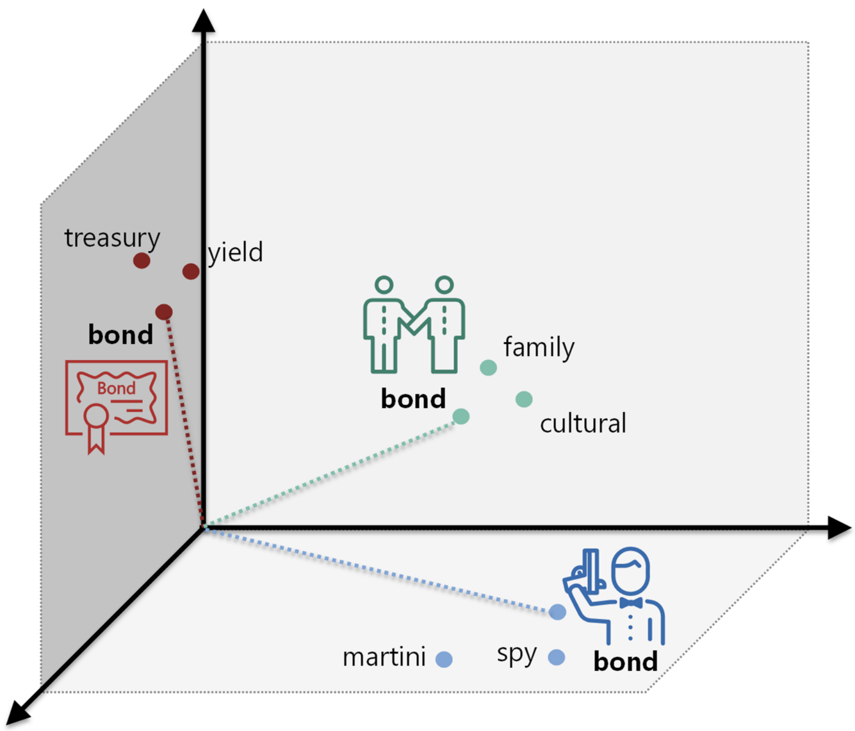

Source: BIS.

This ability to take account of context sets transformer-based large language models apart from previous models. Previous expert systems were tailor-made for specific applications. They needed skilled operators to develop and refine them. For example, AlphaGo, the Go-playing AI system that made headlines in 2016 by beating the world champion Lee Sedol is highly specialised. It is great at playing Go but would struggle to answer the question, "What would John Maynard Keynes have said?"

In contrast, the latest AI applications are versatile so-called "zero-shot learners" that can tackle previously unseen tasks. At worst, they are "few-shot learners" that need only minimal additional training to become conversant in a new, unfamiliar domain. This means they are suitable for use cases beyond textual analysis.

The source of this versatility lies in the combination of vast reservoirs of data and the massive computing power of the latest generation of hardware (Graph 3). The latest large language models have been trained on the totality of the text and non-text data on the internet. In this way, AI has moved from narrow systems that solve specific tasks to more general systems that deal with a wide range of tasks, and all in ordinary language rather than in specialised code.

Computing power and data enable AI to create new insights — Graph 3

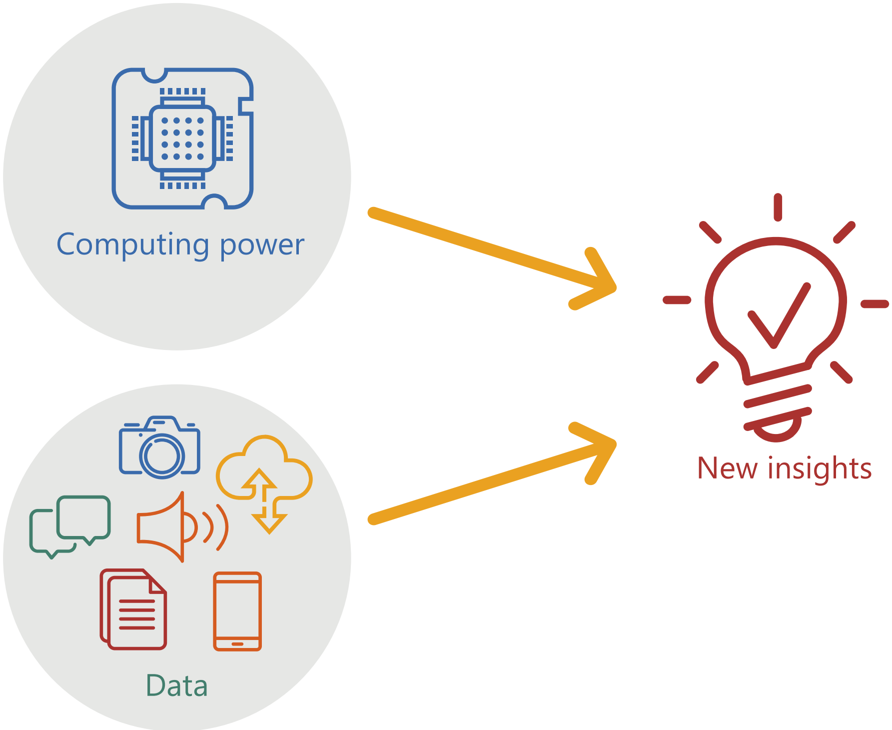

Source: BIS.

Financial sector applications and central bank use cases

The financial sector is a particularly promising area for the application of AI. Already, machine learning has made substantial inroads in the business processes of private financial institutions. Examples include credit assessment and lending, assessing damages in insurance and various applications in asset management. Important use cases also arise in fraud detection and compliance tasks such as customer verification.

Is the reason for AI's prevalence in the financial sector because AI has some kind of magic ability to see things? The answer is no. Rather, the "secret ingredient" is data, or more precisely, a lot of data. The ability to impose mathematical structure on unstructured data makes AI ideally suited to identify patterns that are otherwise obscured.

This ability “to find the needle in the haystack”, even in previously unseen haystacks, could offer breakthroughs in nowcasting economic activity and in the monitoring of financial systems for the buildup of risks. AI thus stands poised to impact central banks as users of the technology.

A key application of large language models is nowcasting real activity or inflation. Traditionally, nowcasting has been hindered by the limited availability of timely data and the need to develop and train models for narrowly defined tasks.

Large language models could help overcome the narrow scope of previous nowcasting models (Graph 4). As versatile zero-shot learners, they can provide forecasts or nowcasts without fine-tuning, and so find needles in previously unseen haystacks. Just as large language models are trained to guess the next word in a sentence using a vast database of textual information, macroeconomic forecasting models can use the same techniques to forecast the next numerical observation. For the AI, any input is just an array of numbers, and the pattern recognition abilities that apply to words can be equally applied to statistical series.

Zero-shot models could enhance the capabilities of nowcasting models — Graph 4

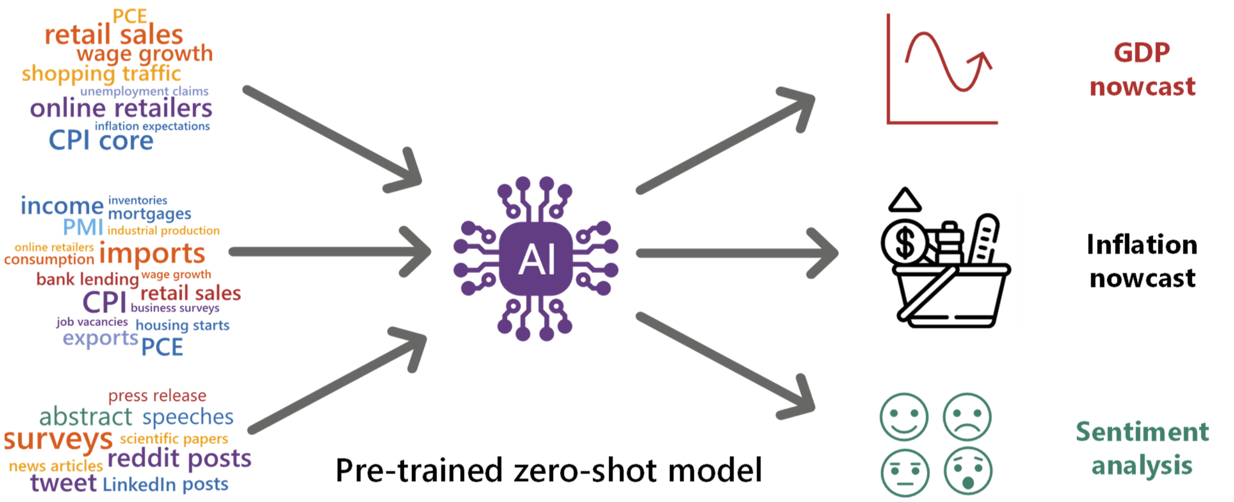

Source: BIS.

Combining time series data with other forms of unstructured data could further enhance the capabilities of nowcasting models. For example, adding non-standard data, such as satellite images, text from social media and so on, could provide additional context to the numerical time series data. The AI model could then be further refined and trained for the nowcasting exercise.

Although these are powerful tools, central banks should not succumb to “magical thinking” – that somehow the tools alone will bring miraculous outcomes. Timely and plentiful data are key to the success of nowcasting applications. AI excels at finding “needles in the haystack”, but there needs to be a haystack with the needles to be found.

But perhaps even more than in nowcasting, it is in the payment system where AI holds the greatest potential. Money laundering networks exploit the complexity of interconnections across firms both within and across borders to obscure the nature of financial transactions.

AI tools can improve the detection of money laundering networks, as illustrated by Project Aurora from the BIS Innovation Hub. Aurora uses simulated instances of money laundering activities that are sprinkled into the payment data. Aurora compares the performance of various machine learning tools with that of the prevailing rule-based approach to assess how well the instances of money laundering are caught by the various approaches.

The comparison occurs under three scenarios: first, transaction data that are siloed at the bank level. Second, national-level pooling of data. And third, cross-border data cooperation using privacy-preserving methods that do not divulge the underlying data to the authorities in other jurisdictions. The results show that machine learning models outperform the traditional rule-based methods prevalent in most jurisdictions (Graph 5). The pooling of data at the national level gives another boost to performance. Most strikingly, machine learning methods really excel when data from different jurisdictions are shared in a privacy-preserving way. Data cooperation improves detection dramatically over the current rule-based method.

AI tools can improve the detection of money laundering networks[1] Graph 5

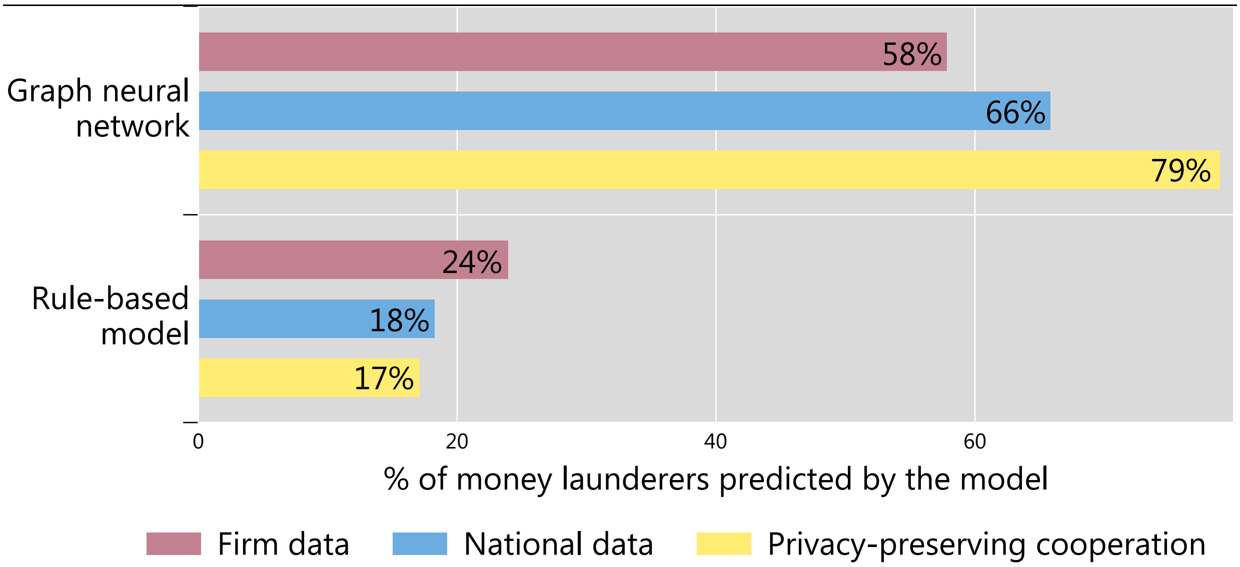

[1] Transaction data visible on three different levels of analysis: the view of each financial institution (firm data), the national view of a single country (national data) and the cross-border view across countries (privacy-preserving cooperation).

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

There are of course risks arising from AI, two of which deserve particular attention. First is financial stability. Reliance on the same handful of algorithms could amplify procyclicality and market volatility by exacerbating herding, liquidity hoarding, runs and fire sales. But AI could also be harnessed for more effective financial stability monitoring. It could help in building early warning indicators that alert supervisors to emerging pressure points known to be associated with system-wide risks.

The second risk is a greater prevalence of cyber attacks. As well as more sophisticated versions of familiar tricks such as phishing, there could be entirely new sources of cyber risk that exploit weaknesses in large language models to make the model behave in unintended ways, or to reveal sensitive information. But here again, just as AI increases cyber risks, it can be harnessed by cyber defenders in their threat analysis and monitoring of computer networks. In a recent BIS survey of central bank cyber experts, most central banks deem AI to be effective or moderately effective at combatting cyber attacks (Graph 6). They believe that AI systems can outperform traditional methods in enhancing cyber security management, especially in areas such as the automation of routine tasks or threat detection.

Just as AI increases cyber risks, it can also be harnessed by cyber defenders[1]

As a percentage of respondents — Graph 6

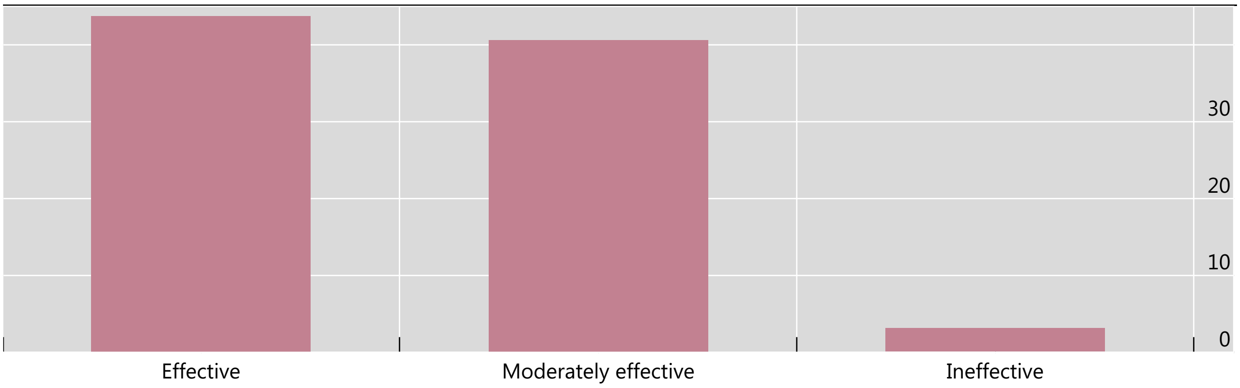

[1] Based on a survey of 32 members of the Global Cyber Resilience Group in January 2024. The bars report the share of respondents that selected each answer to the question, "How effective do you believe AI is in identifying and responding to cyber threats compared to traditional methods?".

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

AI and the macroeconomy

When it comes to the labour market and the macroeconomy, the impact of AI will depend on different channels: how many workers AI displaces, by how much it raises productivity and how many new tasks it creates (Graph 7). The relative strength of these channels will determine aggregate employment and wage dynamics, and it also has implications for inequality.

The impact of AI on employment and wages

Graph 7

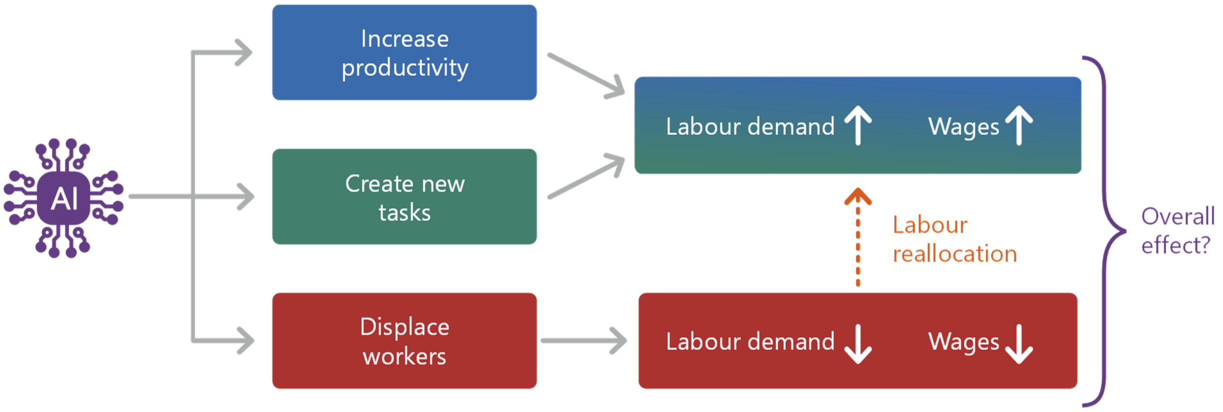

Source: BIS.

What is clear is that AI will expand both aggregate demand and supply – and thereby lead to an increase in output (Graph 8). But the effects on inflationary pressures in the near term depend on the relative impact on aggregate demand versus supply. If households anticipate higher incomes tomorrow, they will spend more today. Inflationary pressures depend on whether this additional spending outstrips supply. Conversely, when there is less anticipation, AI will be less inflationary in the short run.

AI will expand aggregate demand and supply, with uncertain effects on inflation[1]

Changes relative to the initial steady state, in per cent

Graph 8

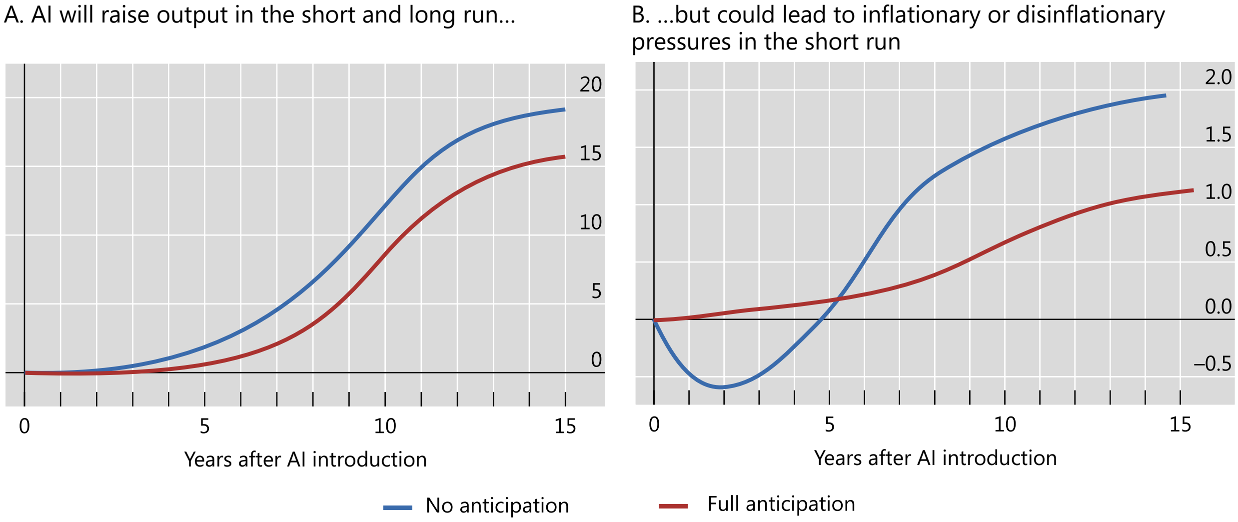

[1] The vertical axis measures the change relative to the initial steady state value of output (panel A) and inflation (panel B). No anticipation refers to the case where households and firms do not anticipate the future effects of AI on productivity growth. Full anticipation refers to the case where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

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

Towards a community of practice among central banks

Central banks will need to be attuned to these dynamics. But more generally, there is a need to rethink central banks' role in collecting and using data, and how central banks should respond to the challenges.

Traditionally, most data were collected and hosted within statistical agencies, including the central bank, with clearly defined access rights. And public institutions have traditionally acted as data providers to private sector firms and the general public (Graph 9).

But our intuition for "data" appeals to existing structured data sets organised around traditional statistical classifications. The age of AI will rely increasingly on unstructured data drawn from all walks of life, with more and more data collected by autonomous AI agents. Central banks are already using AI and unstructured data to fulfil their mandates. But much of the unstructured data reside in the hands of the private sector, which increasingly acts as data provider.

Unstructured data often reside in the hands of the private sector | Graph 9

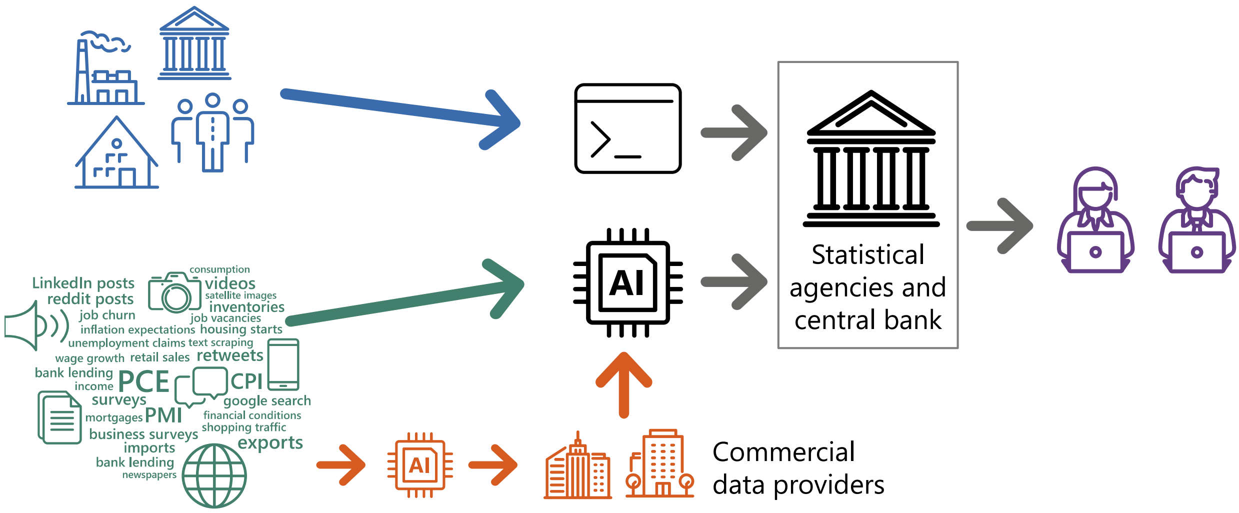

Source: BIS.

So, one important question is how much central banks would rely on in-house data and how much they would source externally. Another key challenge is setting up the necessary IT infrastructure. This is very expensive but crucial in the age of AI. Staffing also arises as a key priority. The challenge of having the right mix of skills will only grow.

How can central banks address these challenges? For one, the rising importance of data and emergence of new data sources call for even greater attention to sound data governance practices. Another issue is the importance of metadata, or the "data about the data". In the future, data will be assembled increasingly by autonomous machine learning applications that need to be guided on what to look for. Metadata frameworks are hence crucial for data retrieval, as well as better comparability.

In all this, cooperation among central banks holds the key to meeting the challenges in the age of AI. The pooling of resources and knowledge can mitigate resource constraints and lower the barriers for central banks in using AI tools. In particular, central banks would benefit from sharing domain-adapted or fine-tuned models. Comparing notes on policy issues and on the role of central banks as data producers, users and disseminators would yield great benefits.

There is an urgent need for central banks to come together to foster a "community of practice". In this community, central banks can share knowledge, data, best practices and AI tools.

Conclusion

Let me conclude.

Central banks have a history of successful collaboration in overcoming new challenges. In the era of AI, they are well-placed to expand cooperation in the fields of data and data governance, as well as in the technology itself.

We should not underestimate the efforts needed to harness the full potential of AI. But the fruits of cooperation in a community of practice will be considerable and the BIS stands ready to play its part.