



BANCA D'ITALIA  
EUROSISTEMA

## Data Science in Central Banking

Closing remarks by Paolo Angelini,  
Deputy Governor of the Bank of Italy

4<sup>th</sup> Irving Fisher Committee and Bank of Italy Workshop  
Rome, 20 February 2025

Good afternoon ladies and gentlemen,

I am happy to be here with you today as we mark the conclusion of this workshop on data science in central banking. As highlighted in the Irving Fisher Committee [Annual Report for 2024](#), launching this periodic workshop series jointly with the Bank of Italy, back in 2019, was a far-sighted strategic decision.<sup>1</sup>

The focus of this fourth edition of the workshop, Generative Artificial intelligence (AI) and its applications in central banking, aptly underlines the ongoing transformation in the field.<sup>2</sup>

There is broad consensus that Generative AI has the potential to significantly boost productivity worldwide. According to some estimates, it could yield annual productivity gains in the \$2.6 - 4.4 trillion range globally.<sup>3</sup> Others predict an increase in output of the order of 15-20 percent, largely phased in over 15 years after adoption.<sup>4</sup> Survey evidence collected from global corporations suggests that for over 40 percent of respondents the return on investment in the most advanced Generative AI initiatives falls within the range of 11-30 percent, and meets expectations (Fig. 1).

---

<sup>1</sup> The primary goal of these workshops is to showcase projects, share expertise among central banks, and reduce reliance on external service providers. The first workshop, in January 2019, focused on "[Computing Platforms for Big Data and Machine Learning](#)". The second focused on "Data Science in Central Banking" and was organized in two parts, dealing respectively with "[Machine learning applications](#)" (October 2021), and with "[Applications and tools](#)" (February 2022). The third workshop was also on "Data Science in Central Banking", with a focus on "[Data sharing and data access](#)" (October 2023).

<sup>2</sup> An excellent primer on AI and central banks is in Bank for International Settlements, "[Annual report](#)", June 2024, pp. 91-127.

<sup>3</sup> McKinsey, "[The economic potential of generative AI](#)", June 2023.

<sup>4</sup> P. Aghion, S. Bunel, and X. Jaravel, "[What AI Means for Growth and Jobs](#)", Project Syndicate, January 14, 2025. I. Aldasoro, L. Gambacorta, A. Korinek, V. Shreeti and M. Stein, "[Intelligent financial system: how AI is transforming finance](#)", BIS Working Papers, no 1194, June 2024.

The uncertainty surrounding any estimate remains significant, largely due to the rapidly evolving nature of the phenomenon itself. Consider for instance the launch of the ChatGPT-like AI model “R1” by DeepSeek, a Chinese startup, at the end of January. The event cast doubt on the widespread belief that AI requires massive amounts of hardware and energy, raising questions about the leadership of the American tech industry and causing strong stock markets fluctuations.

The adoption rate of AI by firms is also uncertain, but growing rapidly. In early 2024, 65 percent of respondents in a survey of international corporations reported that their organizations regularly use Generative AI, nearly double the percentage observed in 2023.<sup>5</sup> Adoption rates vary across industries and appear to be higher among larger firms. Moreover, the speed of adoption within firms seems uneven: functions such as IT, operations, marketing, customer service, and cybersecurity appear to be at the forefront.<sup>6</sup> Additionally, perceptions of Generative AI differ markedly across corporate hierarchies, with top executives tending to have a more optimistic view of their organization’s speed of adoption and implementation.<sup>7</sup>

Uncertainty has led some commentators to avoid making precise predictions about AI’s future impact, instead considering two broad scenarios. In one scenario, the gradual adoption of Generative AI leads to incremental economic progress. In the other, more transformative scenario, AI fundamentally reshapes the world as we know it.<sup>8</sup> While the latter scenario may seem like science-fiction, some ongoing developments, such as the emergence of the so-called Agentic AI, suggest that it could in some form become a reality.

Regardless of the reliability of available estimates, there is no doubt that Generative AI is rapidly expanding across the business world and reaching consumers as well. Looking ahead, further growth seems inevitable, driven by increasing volume and diversity of data being generated and stored, the very fuel of AI, which Generative AI itself contributes to. The amount of data created worldwide each year is projected to continue its astonishing expansion, from 149 zettabytes (i.e.,  $10^{21}$  bytes) in 2024 to more than 394 zettabytes by 2028 (fig. 2).

As this event draws to a close, I would like to reflect on some of the key topics that emerged from our discussion and extrapolate a little to broader considerations.

Over the past three days, we have explored several use cases of AI and data science that are gradually permeating data analysis, real-time economic monitoring, and decision-making in central banking. One of the key takeaways from our discussions is that AI can enhance the speed, accuracy, and depth of forecasting and nowcasting (e.g., for inflation and output, or energy demand). We have also seen applications in regulatory compliance, financial

---

<sup>5</sup> McKinsey, “[The state of AI in early 2024: Gen AI adoption spikes and starts to generate value](#)”.

<sup>6</sup> Deloitte, “[Now decides next: Generating a new future](#) [Deloitte’s State of Generative AI in the Enterprise](#)”, Quarter four report, January 2025.

<sup>7</sup> Deloitte, op. cit.

<sup>8</sup> M. S. Barr, “[Artificial Intelligence: Hypothetical Scenarios for the Future](#)”, speech at the Council on Foreign Relations, New York, New York February 18, 2025.

supervision, and legal analysis. Large Language Models (LLMs), when combined with Retrieval-Augmented Generation (RAG) techniques<sup>9</sup> can extract insights from thousands of pages of regulatory filings, or legal documents, reducing the time devoted to document review and processing. Promising applications have also been showcased also in the field of financial stability monitoring. These findings confirm that AI applications are set to become ubiquitous in modern central banking, and beyond.

Several presentations also highlighted the risks associated with the introduction of AI. First, there is the issue of explainability in AI models. Many machine learning (ML) algorithms function as black boxes, making it difficult for users to understand the rationale behind specific results or recommendations. This is particularly concerning for applications in regulation and supervision, where transparency and accountability are crucial. Second, biases embedded in historical financial data may be reinforced by AI models, raising ethical concerns. Third, Generative AI is prone to hallucinations and misinformation. We have seen how LLMs can fabricate imaginary regulatory clauses, or misinterpret complex financial concepts; more generally, AI can generate incorrect answers that may be difficult to detect.<sup>10</sup> Fourth, the widespread application of Generative AI to investment decisions raises financial stability concerns, partly due to its potential to amplify herding behavior.<sup>11</sup> Fifth, while AI has useful applications in cybersecurity, fraud detection and anti-money laundering, it can also serve be a powerful tool for criminals and fraudsters. I am sure that many other issues could be added to this list.<sup>12</sup>

While these problems are significant, they have been clearly identified, and solutions are actively being developed. For instance, hallucination problems have been reduced in more recent releases of LLMs and can be further mitigated by grounding AI-generated responses in verified datasets through RAG models. Human oversight remains essential, but viable avenues for improvement are emerging.

But alongside these problems we also face other potential issues that are more difficult to identify and address. I am referring to the effects that the systematic use of these tools may have on the cognitive processes of individual users. For instance, as the technology becomes more reliable with each new release, there may be a growing tendency to reduce human oversight and accept the answers generated by AI engines as “the truth”. While this could lead to efficiency gains, it could also result in errors and the spread of incorrect or distorted information and beliefs, not only in financial markets, but across all fields. Furthermore, as responses and outcomes that closely resemble the truth

---

<sup>9</sup> RAG models enhance AI capabilities by adding information retrieval from specialized documents to the general purpose LLM text generation capabilities.

<sup>10</sup> Among the respondents to the 2023 DigitalOcean [survey on startups and SMEs trends](#), 34 percent mentioned security concerns, while 29 percent pointed to ethical and legal issues, including AI hallucinations and privacy risks.

<sup>11</sup> Financial Stability Board, [“The Financial Stability Implications of Artificial Intelligence”](#), 2024; G7 Finance Ministers and Central Bank Governors under the 2024 Italian Presidency, [“A High-Level Panel of Experts’ Report to the G7”](#).

<sup>12</sup> For instance, F. Natoli and A. G. Gazzani, [“The Macroeconomic Effects of AI-based Innovation”](#), confirm that AI gradually raises total factor productivity, but also tends to increase income and wealth inequality.

become instantly available at virtually no cost, users may have less incentive to memorize information. These shifts are occurring in a world where the role of machines, already significant, is expected to expand rapidly (fig. 3).

Generative AI is just the latest addition to a long list of innovations that may significantly impact how the human brain functions (consider for example the pervasive use of mobile devices, the internet, or social media). A parallel phenomenon that comes to mind is remote working, another major trend reshaping how we work and interact. Both AI and remote work bring clear and tangible benefits, but they may also have potential side effects that are harder to identify and assess, as they may only materialize in the longer term. We should all be mindful of these risks.<sup>13</sup>

That said, I am not suggesting that we should delay AI adoption until uncertainty is reduced, or that its use should be resisted. However, low visibility calls for prudent driving. To ensure that AI is used in a robust, effective, fair, and accountable manner, adopters should focus on three key areas, which were echoed in various workshop presentations.

First, developing a sound AI governance framework should be a priority. At the macro level, as outlined in the High-Level Panel of Experts' Report to the G7, this requires "policy preparedness", based on alternative AI development scenarios.<sup>14</sup> At the firm level, Generative AI significantly expands the range of users, moving from a relatively small group of specialists to virtually the entire workforce. Unlike previous technologies, it has the potential to radically transform operational models and introduce new risks. This calls for a high level of responsibility and strategic attention.

The governance argument is particularly strong in the specific field of data management. Data is the fuel of AI, and no engine can function properly if fed with contaminated or low-quality fuel. Until now, proper data management has enabled efficient and effective search, retrieval and usage. Looking ahead, it will be crucial for extracting value from AI applications. Proprietary data are likely to be fundamental for the successful adoption of Small Language Models. These models hold great promise, as they can be easily tailored to a firms' specific needs while also mitigating the high costs and large energy consumption associated with LLMs. Examples relevant to central banks were provided during the workshop.

Second, the workforce should prepare for AI adoption. In the case of central banks, experts in economics and law have traditionally dominated the workforce. However, the digital revolution has gradually increased the pressure to reorient the talent mix

---

<sup>13</sup> See e.g. I. Dergaa, ["From tools to threats: a reflection on the impact of artificial-intelligence chatbots on cognitive health"](#), *Frontiers in Psychology*, April 2024; S. Ahmad, ["Impact of artificial intelligence on human loss in decision making, laziness and safety in education"](#), *Nature - Humanities and Social Sciences Communications*, June 2023; T. Ray, ["The paradox of innovation and trust in Artificial Intelligence"](#), *Observer Reserve Foundation*, February 2024. While these side effects, if indeed they exist, are quite general, we should ask ourselves if there is anything specific to finance, or central banking, which might need attention.

<sup>14</sup> G7 Finance Ministers and Central Bank Governors under the 2024 Italian Presidency, ["A High-Level Panel of Experts' Report to the G7"](#). This approach emphasizes readiness and flexibility, ensuring that policies can be adjusted and refined as new evidence emerges and technology matures.

toward experts in IT, engineering, and mathematics. AI, with its transformative impact on work processes, will further intensify this pressure, necessitating adjustments in hiring policies. But in such a fast-changing environment, simply hiring specialists cannot be the main solution. Training, reskilling and upskilling the workforce will be essential. Proper attention to these aspects will not only help organizations fully leverage new technology, but also address employees' resistance to change, often a major obstacle to successful adoption.

Third, cross-institutional cooperation among authorities (not just central banks and bank supervisors) should be encouraged, given the increasingly blurred regulatory boundaries in data-intensive technologies such as AI.<sup>15</sup> At the start of a new journey, it makes sense to learn from each other's mistakes and share best practices rather than rely on a trial-and-error approach in isolation. In today's challenging global environment, multilateralism is under strain, and international cooperation is facing strong headwinds. Given its long tradition of cooperation in international fora such as the BIS and the FSB, the central banking community is well-positioned to foster cooperation in this field.<sup>16</sup> Workshops like this one embody that spirit.

A clear area for targeted cooperation is cross-border data sharing. Many of you are familiar with the difficulties of enabling data access, both nationally and internationally, due to the complex intersections between personal data protection, national security, intellectual property, and sectoral regulation.<sup>17</sup> In a sense, AI will raise the opportunity cost of maintaining exclusive control over information, calling for smart ways to reconcile public policy objectives with data-drive innovation.<sup>18</sup>

Concluding on this final note, I encourage all of you to stay in touch, actively collaborate, and continue the conversations initiated during this workshop. These discussions can serve as the foundation for future partnerships and joint initiatives.

I would like to extend my sincere gratitude to the speakers and organizers whose dedication and hard work have made this workshop possible. I wish you all continued success in your endeavors.

---

<sup>15</sup> P. Hernández de Cos, "[Managing AI in banking: are we ready to cooperate?](#)" Keynote speech at the Institute of International Finance Global Outlook Forum Washington DC, 17 April 2024.


<sup>16</sup> O. Borgogno and A. Perrazzelli, "[From principles to practice: The case for coordinated international LLMs supervision](#)", Cambridge Forum on AI: Law and Governance. Vol. 1, e13, 2025.

<sup>17</sup> O. Borgogno and M.S. Zangrandi, "Data governance: a tale of three subjects" (2022) *Journal of Law, Market & Innovation*, 1(2), 50-75.

<sup>18</sup> The EU has been working on various data sharing measures that have the potential to enhance AI model training and development. European Commission, "[Working Document on Common European Data Spaces](#)", 24 January 2024 SWD(2024) 21 final, designs a plan to develop Common European Data Spaces, particularly in finance, energy, and health. Similarly, regulatory sandboxes enshrined in the EU AI Act could provide a valuable opportunity to facilitate the responsible use of personal data collected for other purposes. See [Regulation \(EU\) 2024/1689](#) of 13 June 2024 laying down harmonized rules on artificial intelligence (Artificial Intelligence Act), Art. 59, Rec. 140.

Figure 1

**Distribution of the estimated return on investment (ROI)  
for the most advanced scaled Generative AI initiatives  
in a sample of international corporates (1)**




Source: Deloitte, "Now decides next: Generating a new future. Deloitte's State of Generative AI in the Enterprise", Quarter four report, January 2025.

The panels portray the frequency distribution of ROI across surveyed firms. For instance, in the left-hand side panel, from the bottom, 5 percent of firms report that they did not attempt to measure the return on their investment, 9 percent estimate a ROI to be less than 5 percent, and so on.

Figure 2

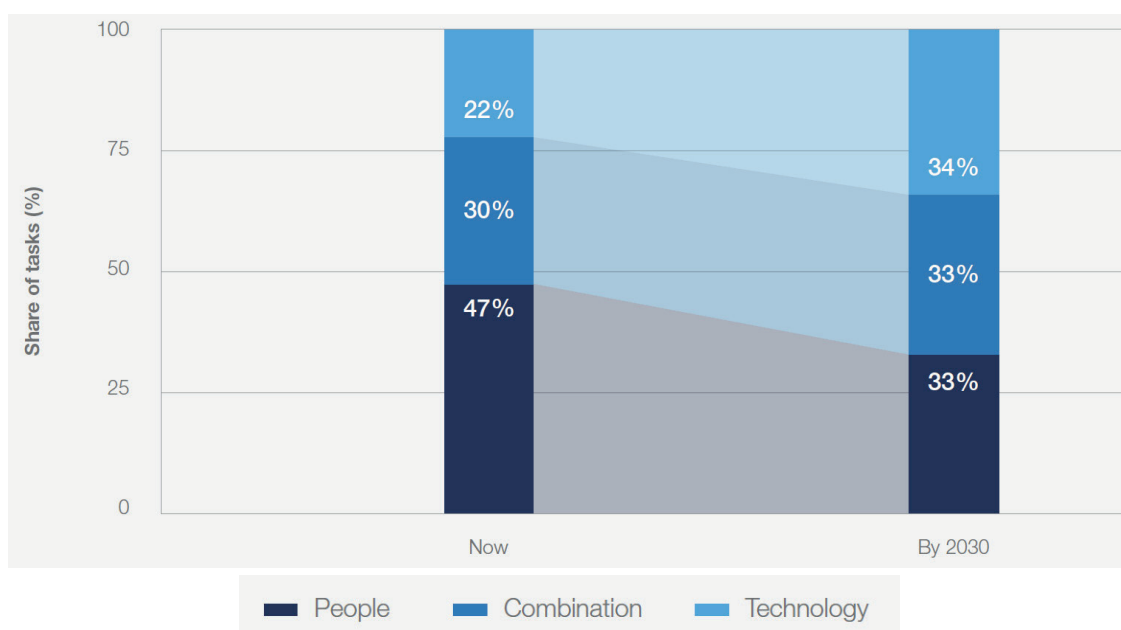
**Volume of data/information created, captured, copied,  
and consumed worldwide each year from 2010 to 2023,  
with forecasts from 2024 to 2028**



Source: IDC; Statista estimates.  
© Statista 2024.

Figure 3

Share of total work tasks expected to be delivered predominantly by human workers, by technology (machines and algorithms), or by a combination of both



Source: World Economic Forum, "Future of Jobs Report", January 2025.