Artificial intelligence: a supply-side perspective

Artificial intelligence (AI), in a broad sense, is the ability of machines and computational systems to replicate human intelligence in the perception, synthesis and inference of information, thereby performing tasks which previously could not be carried out or traditionally required human cognitive abilities, such as language comprehension, pattern recognition and decision-making.

While the development of this technology dates back to the 1950s-60s, advances in large language models (LLMs) over the past decade, along with improvements in processing power and mass data collection, have driven more recent developments in so-called generative AI, capable of producing text, code and audiovisual material based on patterns learned from large datasets.

The AI value chain: complex and heterogeneous from country to country

The development of AI relies on a complex value chain composed of several interdependent links.[1] At its core lies access to critical minerals used in the manufacture of semiconductors, the «physical brain» of AI. These components are integrated into a broader computational infrastructure comprising data centres, communication networks, power grids and cloud computing services, together forming the «body» that enables large-scale data processing. On top of this infrastructure lies access to large volumes of data for training models. The development of so-called foundational models represents in itself the next link in the chain and requires sophisticated algorithms and neural networks for deep learning. Finally, the value of generative AI is realised in the development of specific applications built on foundational models, such as virtual assistants or content generation systems, and in its integration into digital products and services, which are the main point of contact with the end user.

Structure of the AI value chain by country

(% of total valuation) *(% of GDP)*

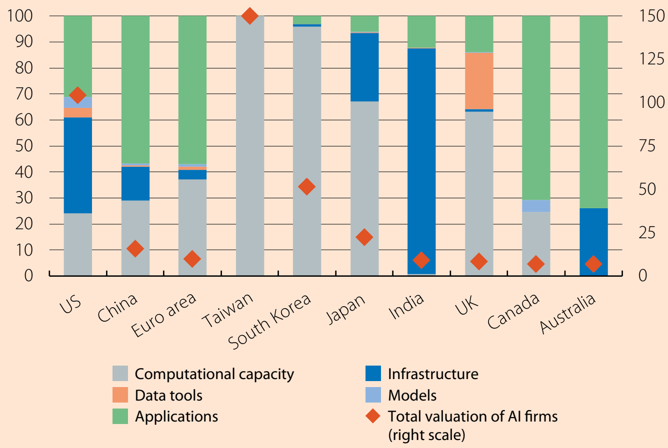

Note: *In the case of Taiwan, the total valuation of AI firms (as a % of GDP) is 207.2%.*
Source: *CaixaBank Research, based on data from the BIS (K. Rishabh and V. Shreeti (2026), «The geography of AI firms», BIS Working Papers 1343).*

Currently, AI value chains vary widely across the major economies (see first chart). On one hand, in the US and several Asian nations the sector plays a significant role in the economy, measured in terms of AI firm valuations as a percentage of GDP. Among these, Taiwan and South Korea are highly specialised in computational capacity, while the US has a more diversified value chain. On the other hand, in China and several advanced economies, AI plays a somewhat smaller role in the economy and they also have different specialisations. Japan and the United Kingdom show greater specialisation in computational capacity (and data tools, in the case of the latter), while in China and the euro area there is a greater emphasis on applications.

AI deployment, from innovation to adaptation

The deployment of AI can be divided into four key phases: the innovation phase, the development of new infrastructure, the diffusion and widespread adoption of the new technology, and the adaptation of business models and markets to the new technology.

In this context, the global economy is still immersed in the first two phases of AI deployment. There is ample evidence of an investment boom related to innovation and infrastructure construction, which is particularly evident in countries such as the US and some Asian nations.[2] In this context, AI capabilities are improving at an exponential rate. This progress is being supported by

1. See O. García Retuerta and D. García Retuerta (2026), «La cadena de valor de la inteligencia artificial: estrategias de autonomía para España», IEEE Opinion Document 03/2026, Spanish Ministry of Defence, Spanish Institute for Strategic Studies; and McKinsey & Company (2023), «Exploring opportunities in the generative AI value chain», QuantumBlack, AI by McKinsey.
2. See the article «The AI buzz in financial markets», in this same Dossier.

hyperscaling, driven by rapid advances in the amount of data used to train models, the number of parameters, and computational capacity. At the same time, the surge in supply and demand is creating infrastructure bottlenecks.[3]

AI supply indicators in the US, the EU and China

	China	US	EU-27
Innovation, infrastructure and procurement			
Market share for logic chip production, by manufacturing stage[1]			
Design	9.0	61.0	0.0
Fabrication	12.0	27.0	2.0
Assembly, testing and packaging	14.0	28.0	0.0
Granted AI patents (per 100,000 people)[2]	7.0	4.7	2.6
Scholarly publications on AI (per million people)[1]	72.7	145.7	139.4
Highest score achieved by an AI model (MMLU)[3]	90.6	92.5	84.0
Number of data centres[2]	449.0	5,427.0	1,461.0
Number of notable data models (2021-2025)[2]	108.0	331.0	44.0
Number of cumulative GitHub stars (millions)[2]	9.0	30.0	13.0
Innovation Index (Global AI Vibrancy Tool)[2]	5.1	20.7	3.0
Adoption, diffusion and adaptation			
AI adoption rate among the population[2]	17.0	28.3	31.2
Industrial robots in operation (per 10,000 employees)[4]	166.0	307.0	266.0
Number of industrial robots installed in the last year (thousands)[4]	276.3	37.6	50.3
Electricity consumption of data centres (% of total electricity demand)[5]	1.1	4.4	2.3
Digital infrastructure (AI Preparedness Index)[6]	0.19	0.19	0.17
Innovation and economic integration (AI Preparedness Index)[6]	0.15	0.18	0.16
Human capital and labour market policies (AI Preparedness Index)[6]	0.15	0.18	0.16
Graduates from STEM programmes (% of total, tertiary education)[7]	41.0	20.0	25.0
Digital-intensive sectors' share in total employment (% of total employment)[8]	28.0	47.8	47.4
ICT goods and services (% of international trade)[8]	22.2	9.8	7.5
Digital Services Trade Restrictiveness Index[8]	0.29	0.06	0.11

Note: *The latest available year is used for each series, unless otherwise stated. In cases where data for the EU aggregate is unavailable, the average is calculated using the available countries. Values in red indicate more negative performance, yellow medium, and green more positive. Sources: [1] Our World in Data, [2] Stanford University Institute for Human-Centered AI, [3] MMLU-Pro Benchmark Leaderboard, [4] International Federation of Robotics (IFR), [5] International Energy Agency (IEA), [6] International Monetary Fund (IMF), [7] World Bank and Center for Security and Emerging Technology - Georgetown University, [8] OECD.*
Source: *CaixaBank Research, based on various sources.*

As has been the case with other technologies in the past, some economies will not play a decisive role in the innovation phase, but will benefit from adoption, diffusion and adaptation to the technology. If we focus on the comparison between the US, the EU and China, we can see important nuances in different phases of the deployment process. In the innovation phase, the US economy is taking a clear lead, as is particularly evident in indicators related to output (such as the performance of models at the technological cutting edge, academic publications and open-source development) and infrastructure (such as the number of data centres and chip design). Yet China has shown a remarkable ability to catch up with the technological frontier in recent years. Its most advanced models show a very similar level of performance to their American counterparts, while the dynamism observed in the granting of patents and the development of models points to a boom in innovation. The EU, meanwhile, is not so well positioned in terms of innovation, according to most indicators. In particular, its low market share in chip production makes it highly dependent in this sphere, while it is lagging behind China and the US in the development of AI models. Finally, China is leading in the supply of materials, due to its access to critical minerals and its processing capacity for chip and semiconductor manufacturing.[4]

In indicators related to adoption, diffusion and adaptation, the picture is somewhat more homogeneous. Adoption in the three major economies, in terms of the proportion of the population currently using AI, stands at around 30% in the EU and the US, while for China it is just under 20%.[5] Similarities are also found in their degree of readiness for adoption, diffusion and adaptation,

3. For example, METR, a metric which measures AI performance based on the maximum task length (time horizon) a model can handle, shows that in recent months it can now satisfactorily perform tasks that would require several hours, whereas a year ago models could only handle tasks with a duration of minutes. See also «The AI Index 2026 Annual Report», by Stanford University's Institute for Human-Centered AI. The main infrastructure bottlenecks are found in the chip market, but also in data centre capacity and in the energy market.
4. See «China's alchemy: how it transforms critical minerals into global power» in the MR01/2026.
5. These are significant figures that indicate a substantially higher adoption rate than that of previous technologies. Business adoption figures, meanwhile, show greater dispersion by function, sector and degree of implementation. See «The AI Index 2026 Annual Report», by Stanford University's Institute for Human-Centered AI.

albeit with a slight advantage for the US. On the other hand, in recent years China's manufacturing sector has undergone a very rapid modernisation – and «robotisation» – process, anchored in an aggressive industrial policy and significant investment in infrastructure and human capital. This places it in a strong position to benefit from the difussion of and adaptation to AI, especially as a global supplier of advanced technologies. Finally, the European and US economies are more intensive in digital services, positioning them as potential leaders in the adaptation phase, the speed and magnitude of which will be key to determining the macroeconomic effects of AI.[6]

The global economy, in the early miles of the AI marathon

The «AI race» is still in its early stages. While the US has taken the lead in the innovation phase, the pack is closing in, led by China, and it is unlikely that the race will be decided between just two participants. Given its transformative potential, the success of the deployment of AI and its macroeconomic impact will depend on the business sector's ability to adapt and manage the frictions associated with this new technology. However, AI will also require an active role from states, both in its regulation and in its adoption, diffusion, adaptation and coordination at a global level, promoting the necessary improvements in terms of institutions, infrastructure and human capital.[7] The task ahead is complex and will require new tools of public policy and economic diplomacy. Moreover, the AI supply model that is ultimately adopted – whether in fragmented blocks centred around the US or China, or more integrated globally – will have implications that extend beyond the economy. The AI marathon is only just begun, and everyone is taking part.

Isabela Lara White and Luís Pinheiro de Matos

6. For further details, see the article «Productivity and employment in the face of generative AI: what do we know?», in this same Dossier.
7. For further details, see the article «Differentiated strategies for governing AI: towards cooperation or conflict?», in this same Dossier.

Differentiated strategies for governing AI: towards cooperation or conflict?

Generative artificial intelligence (AI) is a critical area of economic and strategic competition among the major powers, and its development depends both on the dynamism of the private sector and on state action. Together, they shape the scope and effects of a technology whose complex ecosystem spans innovation and its monetisation, positioning in the value chain, its diffusion and adoption, and the management of its externalities. This article reviews – from a geoeconomic perspective – the strategies adopted by the US, China and the EU in key areas such as regulation, the role of the state in the industrial model, public support instruments, and cross-cutting policies such as professional training and sustainability. We conclude with a reflection on the future interaction of these governance models and the potential areas of friction and cooperation that may arise.

AI governance: from development to adoption

Analytical framework for international comparison

AI development
Technological frontier and strategic capabilities

Innovation & technological frontier
- Foundation models
- Research & development
- Computational capacity & semiconductors

Industrial model
- Position in the global value chain
- Vertical integration
- Access to critical minerals

Role of the state
- Planning & coordination
- Public-private collaboration
- National priorities

Legal framework
- *Ex ante* and *ex post* regulation
- Horizontal or sectoral scope
- Security & content control

AI adoption and diffusion
Productive deployment and impacts

Adoption mechanisms
- Market & platforms
- Public procurement
- Public facilitation & support

Regulation of AI use
- Risk-based approach
- Sectoral supervision
- Testing environments

Training & labour market
- Applied higher education
- Technical & vocational training
- Job retraining

Infrastructure, energy & sustainability
- Data centres
- Energy costs
- Environmental restrictions

Source: *CaixaBank Research.*

The US seeks to shape the technological frontier

The potential of AI lies in the complexity, speed and reliability with which it performs tasks. Its development relies on a combination of cutting-edge knowledge for the design of language models, computers equipped with high-capacity processing chips, and a robust architecture – both physical (data centres) and digital (cloud infrastructure) – for information storage and model training.

In this field, the US has consolidated its position at the global AI frontier thanks to its human capital, technological capabilities and a favourable business environment.[1] It boasts an innovative ecosystem based on elite universities and a concentration of STEM and international research talent. It also benefits from public support as an early incubator, led by both civilian agencies (NSF) and defence agencies (DARPA), and from a business cluster with large tech firms that are embedded into the economy's industrial base and have both financial muscle and an appetite for risk. Added to this are favourable tax and regulatory frameworks, with minimal intervention during the development phase, still lacking a comprehensive federal law[2] and with a predominance of *ex post* intervention. The Trump administration's action plan has placed even greater focus on the technological frontier with a marked geostrategic emphasis,[3] setting an explicit goal for US semiconductors, models and applications to be hegemonic on a global scale and become the new «gold standard».[4]

In contrast, state planning, coordination and guidance are the foundation of the Chinese model. While it is private companies that have capitalised on the exponential improvement of technological capabilities in the last decade, AI research and development

1. According to estimates based on data from Epoch AI, the US accounts for two-thirds of the world's AI-related computational capacity, followed by China with around 20%, while the EU accounts for just 5%.
2. The only general AI law in effect in the US is the one passed by the state of Colorado in 2024.
3. White House (2025), «America's AI Action Plan».
4. It thus shifts the previous focus on the coordination of the innovative ecosystem and industrial resilience as outlined in the National Artificial Intelligence Initiative Act (2020) and the CHIPS and Science Act (2022).

are aligned with national priorities. Whereas the US seeks to define the technological frontier, China prioritises key links in the global industrial value chain,[5] scale, technological self-sufficiency and security. Subsidies, tax incentives and public financing mechanisms, both at the central and provincial levels, are all contributing to this. This approach is complemented by the preventive control of socially sensitive content, including registration and assessment requirements *ex ante* for recommendation systems in digital applications.[6] Recent regulation reinforces the limits on public dissemination of information while maintaining greater relative freedom in the research, development and training of models for productive or strategic uses.[7]

The EU, for its part, is seeking to establish a common governance framework to overcome the prevalence of national frameworks in AI development. The main strength of the European innovative ecosystem is its scientific and research base, with universities and centres of excellence. However, it suffers from insufficient supranational coordination and limited prioritisation of its framework programmes, such as Horizon Europe. The financial system is less geared towards risk-taking and, together with the fragmentation of the internal market, hinders the transfer and monetisation of knowledge, as well as the scaling-up of the technology.[8] To protect its citizens, the EU's regulatory framework prioritises *ex ante* regulation of the uses of AI based on risk,[9] which can shift its development away from the cutting edge of innovation. Added to this is a high external dependency on advanced semiconductors and foundation models, which the EU is attempting to mitigate through open strategic autonomy and diversification of economic partners.[10]

China prioritises adoption and diffusion with productive uses

Beyond technological development, the economic and social impact of AI largely depends on how its adoption and diffusion are governed. In these areas, the approach adopted by the main players also varies widely.

In the US, it is private entrepreneurial initiative and competition that is taking the lead, with the major tech platforms and software providers acting as natural channels for scaling up towards businesses and consumers. State action focuses on removing barriers, providing critical infrastructure, and using public procurement – especially in defence and security – as a driving mechanism for adoption. Regulation is mostly *ex post*, guided by voluntary standards of cross-sector application defined by a federal scientific agency (NIST), along with sectoral oversight in sensitive areas, such as the protection of healthcare patients and financial services clients. Based on this logic of minimal intervention, the state acts as a facilitator and largely leaves the management of cross-cutting areas to the market, although the new national regulatory framework includes recommendations for professional retraining and to limit the impact of the expansion of data centres on electricity costs.[11]

In the Chinese model, as in the development phase, the public sector plays a key role. The state acts as the coordinator of the ecosystem, *ex ante* regulator, financier and demander, channelling substantial public investment through large state-owned enterprises and into strategic sectors such as advanced industry, logistics, energy and security. The planned schedule includes sectoral and territorial penetration objectives at different time horizons, with a roadmap that is due to culminate in a fully «smart» economy and society by 2035.[12] To this end, vertical programmes for the transformation of the industrial value chain have been defined,[13] with controlled competitive environments that facilitate the assessment of scalability without transferring risks to the wider system, such as regulatory sandboxes and pilot zones. This approach is accompanied by the integration of AI into higher education and technical and professional training programmes. Energy and infrastructure planning forms part of the deployment strategy, while sustainability is subordinated to national economic security priorities.

Unlike in the US, where the diffusion of AI relies on large private platforms, and in China, where the state acts as a centralised demander, in the EU the adoption and diffusion of AI is primarily structured through an approach based on regulation and public support. The fragmentation of the internal market and *ex ante* regulatory obligations for high-risk uses limit the pace and scale of

5. See the Focus «China's alchemy: how it transforms critical minerals into global power» in the MR01/2026.
6. Cyberspace Administration of China, CAC (2021), «Algorithm Recommendation Provisions». CAC (2023), «Interim Measures for the Management of Generative AI Services», CAC (2023), «Deep Synthesis Provisions» and CAC (2025), «AI-generated Content Labeling Rules».
7. CAC (2023), «Interim Measures for the Management of Generative AI Services», CAC (2023), «Deep Synthesis Provisions» and CAC (2025), «AI-generated Content Labeling Rules».
8. M. Draghi (2024), «The Future of European Competitiveness».
9. EU (2024), Artificial Intelligence Act.
10. The «AI Continent» Action Plan, presented by the Commission in 2025, extends the strategic public intervention approach applied to semiconductors in the European Chips Act (2023), complemented by the objectives of the Critical Raw Materials Act (2024) to ensure a secure and sustainable supply of critical raw materials.
11. White House (2026), «Artificial Intelligence: national policy framework».
12. These objectives are defined by the work programme of the AI Plus initiative launched in 2024 by the State Council, similar to the Internet Plus initiative of 2015.
13. e.g. the AI+ Manufacturing initiative launched in 2025 under the umbrella of AI Plus.

adoption.[14] Public-sector action combines regulation with EU-wide instruments – such as the Apply AI Strategy – and practical support – such as hubs and testing environments – aimed at facilitating sectoral implementation and reducing legal uncertainty.[15] This approach tends to increase adoption costs and slow down dissemination, especially among SMEs, where fixed costs and skills deficits weigh more heavily. Added to this are structural limiting factors, such as high energy costs and environmental commitments associated with the deployment of compute-intensive infrastructures.[16]

The EU seeks its place in AI geopolitics

The rivalry between the US and China in the AI era is unfolding amid significant strategic uncertainty.[17] It is unclear whether the advantage of being at the technological frontier will generate persistent revenue streams that are difficult to replicate or whether competition will shift towards dissemination, deployment, and the ability to scale up applications in key sectors. In both scenarios, the power associated with AI will largely depend on the control of key assets – advanced chips, computing capacity, energy, talent, and industrial integration – so betting on a single trajectory could prove costly if the evolution of the technology diverges from the initial assumptions.

This framework tends to place middle powers in a position of technological dependence.[18] The concentration of talent, investment, and computing capacity in the US and China limits the scope of influence over the direction of technological change and amplifies the economic and social adjustment costs associated with AI. For the EU, the risk of falling behind reinforces the debate surrounding the balance between regulation, competitiveness and scale. In particular, the Draghi report's diagnosis on internal market frictions and the difficulty of scaling up innovation aligns with the recent shift towards approaches based on simplification and regulatory proportionality. The goal of this shift is to prevent legal certainty from ultimately penalising adoption and scale-up, especially among SMEs.[19]

Nevertheless, AI governance is not necessarily reduced to bloc logic. Even in a context of strategic rivalry, recent multilateral initiatives show areas for coordinating principles and practices. Thus, the focus on security and regulation at the summits in London (2023) and Seoul (2024) has expanded to encompass a more comprehensive agenda of innovation, digital skills, labour impact, and sustainability in Paris (2025), and the emphasis on capacity gaps between advanced and emerging economies in New Delhi (2026). In this vein, the framework promoted at the United Nations suggests a more inclusive and distributed global architecture based on common principles and mechanisms that are complementary to national and regional strategies.[20] For the EU, the challenge will be translating this cooperative agenda into real adoption and scale-up capabilities.

David Martínez Turégano

14. M. Draghi, *op. cit.*
15. The AI Act (2024) establishes support mechanisms for AI deployment to facilitate regulatory compliance in high-risk uses, while the Apply AI Strategy (2025) integrates them into an action plan aimed at accelerating adoption, especially among SMEs and government entities.
16. IEA (2025), «Energy and AI».
17. Foreign Affairs (2026), «Geopolitics in the Age of Artificial Intelligence: Strategy and Power in an Uncertain AI Future».
18. Foreign Affairs (2026), «The AI Divide: How U.S.-Chinese Competition Could Leave Most Countries Behind».
19. The European Commission's proposal set out in the Digital Omnibus package of November 2025 – currently being negotiated among co-legislators – introduces a more pragmatic tone in the regulatory approach, with adjustments aimed at reducing the burden and facilitating technological adoption without altering protection objectives.
20. United Nations (2024), «Governing AI for Humanity».

AI adoption in Spanish firms is advancing rapidly but remains limited and uneven

The adoption of artificial intelligence (AI) in Spanish firms has accelerated in recent years, but the process has been uneven and remains incomplete. This article analyses the extent of AI penetration considering four key dimensions: company size, sectoral differences, specific uses within organisations and the main barriers hindering its deployment, as well as a comparison with the rest of Europe. Understanding how and where AI is being incorporated is particularly important from a business and macroeconomic perspective, as its adoption influences efficiency and productivity gains and can widen gaps between firms, sectors and workers in a productive fabric like Spain's, which is dominated by SMEs and microenterprises.

Business adoption: size matters

Between 2021 and 2025, the adoption of AI technologies in Spain's economy has more than doubled in firms with over 10 employees,[1] going from 8% to 21%, suggesting that AI is no longer an experimental technology. Nevertheless, in 2025, nearly 8 out of 10 firms were still not using it, suggesting that widespread adoption has not yet occurred.[2]

Companies employing AI technologies

(% of all companies)

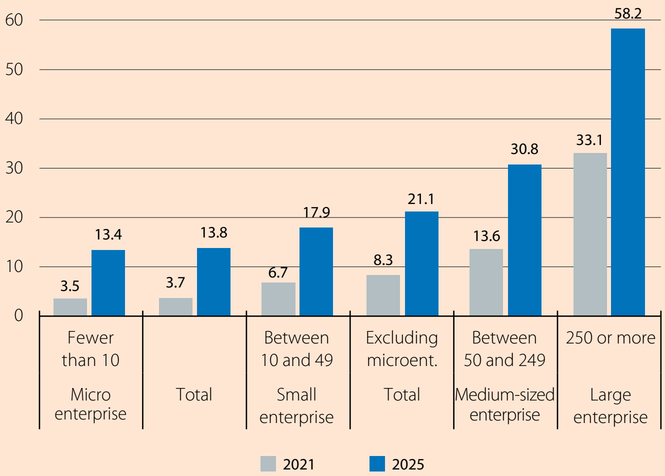

Source: *CaixaBank Research, based on data from the National Statistics Institute.*

The first key conclusion is that company size is a decisive factor. AI is present in nearly 3 out of every 5 large firms, but only in 18% of those with fewer than 50 employees. This gap reflects barriers that go beyond technology, related to financial resources, the availability of data, qualified personnel and organisational capacity.

These differences have not diminished in the period 2021-2025. Although adoption is growing across all company sizes, progress has been far more rapid in large firms (+25 pps) than in small ones (+11 pps). Medium-sized companies (between 50 and 250 employees) reflect a turning point: their level of adoption (31%) is significantly higher than that of small businesses (18%), suggesting they meet a threshold of sufficient resources to experiment with AI.

If microenterprises were included in the analysis, the aggregate adoption rate would fall considerably: instead of 21%, it would be 14%, as these entities have a very low adoption rate (13%) and account for 95% of Spain's productive fabric.[3]

Highly mixed picture of AI adoption by sector

In 2025, the sectors with the highest adoption of AI were information and communications and the ICT sector,[4] with percentages of around 60% – significantly above the levels observed in 2021 (26-27%). In these sectors, AI has become quite prevalent, in line with their greater intensity in intangible capital, the availability of data and their proximity to technological supply.

In a second tier are knowledge-intensive services, such as professional and scientific activities (38.5%) and real estate activities (35%). Thirdly, a wide range of sectors show moderate adoption levels, at around 20%-26%, spanning both services and industry: basic supplies, hospitality and the automotive and electronics sectors.

1. Excluding companies with fewer than 10 employees is also the most commonly used metric for international comparison purposes.
2. For this analysis of the adoption of AI tools by Spanish firms, we use data from the National Statistics Institute's Survey on the Use of ICT and E-commerce in Enterprises. Given that the survey separates companies with fewer and more than 10 employees, the total figure is obtained by weighting both groups according to the business structure (National Statistics Institute's Central Business Register): according to our calculations, AI adoption rises from 4% in 2021 to 14% in 2025.
3. It is also interesting to look at adoption not in terms of companies, but employees. Weighted by employment (and including all companies, even microenterprises), it is estimated that in 2025 around 31% of workers in Spain were employed in companies using AI technologies. This is because adoption is concentrated in medium and large enterprises, which account for a much larger proportion of total employment. This result is obtained by combining AI adoption rates by company size with data on job distribution by size, according to the European Commission's 2025 SME Country Fact Sheet (Eurostat/JRC).
4. The ICT sector includes activities such as the manufacture of electronic components and equipment, software publishing, telecommunications, IT and data processing services, and the repair of ICT equipment (sections 261–264, 268, 465, 582, 61, 6201–6209, 631 and 951 of the National Classification of Economic Activities (CNAE)).

At the lower end we find construction, metallurgy, and transport and logistics, where adoption remains low. Overall, although some sectors are approaching widespread AI usage, the economy remains in an intermediate phase, with significant margin for future dissemination concentrated in the sectors and firms that are lagging the furthest behind.

What are our companies using AI for?

The next issue is to identify the specific functions for which AI is being used in each sector. To this end, we distinguish between two major types of applications. On one hand, we find AI used in processes, aimed at improving organisational and commercial efficiency, such as internal management, administrative tasks or sales support. On the other hand, there is AI linked to the productive core, that is, uses applied directly in the production of goods or in the provision of the main service.[5]

Companies employing AI technology by sector

(% of all companies with 10 or more employees)

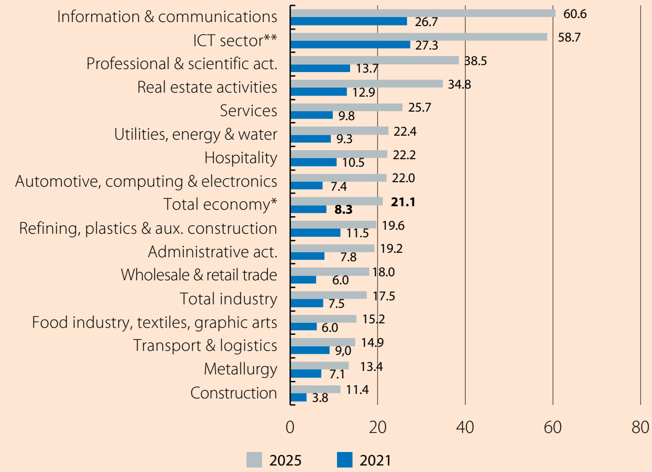

Notes: ** The financial sector is excluded. ** Manufacturing of electronic components and equipment, software publishing, telecommunications, IT and data processing services, and repair of ICT equipment (sections 261-264, 268, 465, 582, 61, 6201-6209, 631 and 951 of the National Classification of Economic Activities (CNAE)).*

Source: *CaixaBank Research, based on data from the National Statistics Institute.*

In this context, adoption is considered superficial when AI is primarily used in processes rather than in the production of goods or services; and it is considered deep when AI is employed to a greater extent in production rather than in management tasks. There is another group of sectors where AI adoption is currently in a transition phase, moving from its use in processes towards greater use linked to production. Most of the adoption in sectors such as hospitality, commerce and administrative activities is found in processes. In these cases, AI is initially introduced in cross-functional management and commercial support roles, where implementation costs and organisational risks are lower, before extending to the productive core.

In an intermediate position we find sectors such as food and textiles, professional and scientific activities, and construction, with a transitional adoption where both types of uses coexist. AI is beginning to permeate operational decisions, but it does not yet play a dominant role in core business activities.

*Level of AI adoption according to the use of these technologies by sector**

(% of all companies with 10 or more employees)

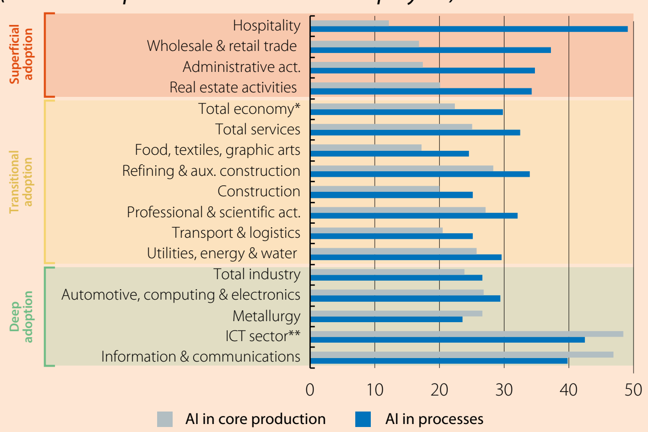

Notes: ** The financial sector is excluded. ** Manufacturing of electronic components and equipment, software publishing, telecommunications, IT and data processing services, and repair of ICT equipment (sections 26-264, 268, 465, 582, 61, 6201—6209, 631 and 951 of the National Classification of Economic Activities (CNAE)).*

Source: *CaixaBank Research, based on data from the National Statistics Institute.*

Finally, sectors such as information and communications and the ICT sector show a high level of adoption both in processes and in the productive core. In these fields, AI is not merely an efficiency tool but an integral part of the product, service and digital infrastructure.

Overall, the pattern observed suggests that the recent progress in adoption within the Spanish economy has been primarily driven by horizontal uses, which are quick to implement, low-cost, and more oriented towards administrative tasks than production processes. This is consistent with what Daron Acemoglu has documented for the US, where improving quality and the reliability of processes is the main motivation for adopting AI.[6] The big unknown for the impact on aggregate productivity is when the use of AI will become widespread in operational and core business uses, which require greater integration, investment, and the redesign of processes.

5. We consider process AI to include uses in business administration and management, marketing and sales, accounting and finance, ICT security, and support for information analysis. Conversely, production AI is considered that directly linked to the production of goods, the direct provision of services, logistics and operations, as well as advanced R&D, such as automation, simulation and process optimisation.
6. See D. Acemoglu *et al.* (2022). «Automation and the workforce: a firm-level view from the 2019 annual business survey». NBER Working Paper, 30659, National Bureau of Economic Research.

The main barrier to greater adoption is the lack of skills

Among the reasons hindering AI adoption, we can distinguish three groups: (i) economic-technological barriers (primarily high costs), (ii) skills, and (iii) data and governance – the quality and availability of data, privacy, and legal clarity. In sectors with below-average adoption, the main and most commonly cited barrier is a skills issue. Below, we set out the barriers related to data and governance, before covering economic-technological barriers.

What are the main barriers to AI adoption?

(% of all companies with 10 or more employees)

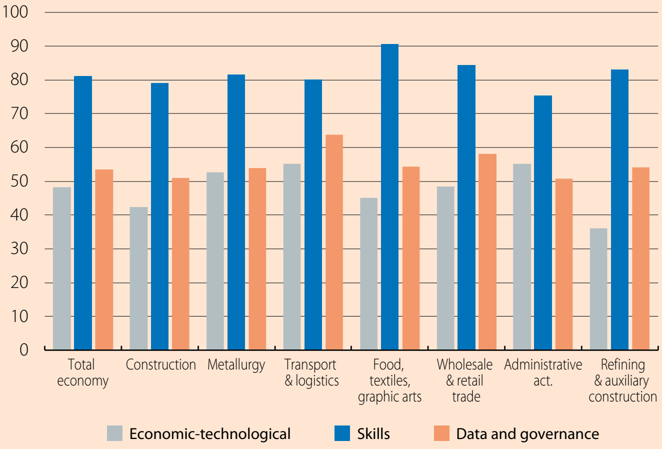

Note: *The chart considers sectors with AI adoption below the economy's average in Q1 2025.*
Source: *CaixaBank Research, based on data from the National Statistics Institute.*

International comparison: we are catching up, but remain behind the European leaders

Data from the European Commission reveal that AI adoption rates in 2025 among Spanish firms with over 10 employees are close to the euro area average (albeit still 5 points below). In large firms, the adoption rate is practically identical between Spain and the euro area; the difference is greater (–7 pps) in medium-sized firms.

Compared to the main European economies, the percentage of firms using AI in Spain exceeds that of Portugal, Italy and France, but is below that of Germany and the Netherlands. The recent acceleration in Spain is particularly noteworthy: between 2021 and 2024, adoption increased by just 3 pps (half that of the euro area), whereas between 2024 and 2025 the increase reached 9 pps – outpacing the euro area average and that of Germany, France or Italy.

Euro area: companies using AI in 2025

(%)

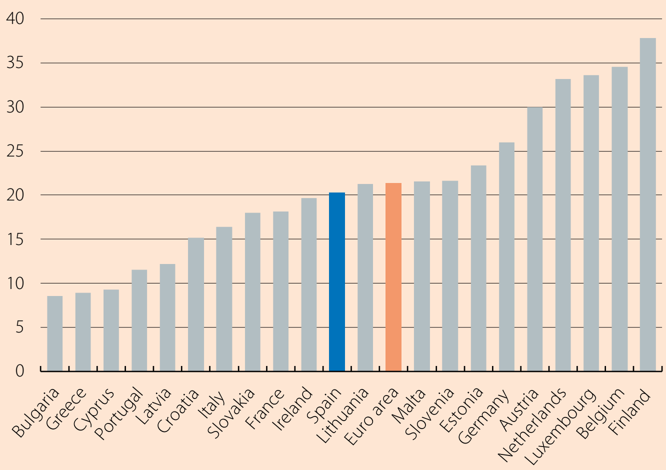

Source: *CaixaBank Research, based on data from Eurostat.*

In summary, there is progress in AI adoption among Spanish firms, but it remains limited and highly uneven. Firstly, company size is crucial: large firms are adopting AI to a much greater extent than SMEs and microenterprises – which are predominant in the country's productive fabric – potentially widening productivity gaps. Secondly, AI is primarily introduced in cross-functional management and commercial support roles, while its penetration into core production tasks is progressing more slowly, hence its still limited overall impact. Finally, the main barriers are related to a lack of skills, highlighting the importance of strengthening human capital through educational and training policies. The challenge, therefore, is no longer to demonstrate the usefulness of AI, but to facilitate its deep and widespread adoption where the greatest obstacles are currently concentrated.

Pedro Álvarez Ondina and Javier Garcia-Arenas.

Productivity and employment in the face of generative AI: what do we know?

Generative artificial intelligence (AI) has traits of a general-purpose technology: applications in many sectors, rapid improvement of the technology itself and a catalyst for complementary innovations. This has already happened with technologies such as electricity and the internet. Even so, having high potential does not necessarily mean an immediate or uniform macro impact. The final magnitude of AI's impact will depend on the speed of its adoption and the ability of firms to reorganise processes. This article examines how AI could affect productivity growth and what it means for the labour market.

Striking productivity increases at the micro level

Since the emergence of ChatGPT in 2022, research on the impact of AI on worker productivity has surged. A review by the OECD indicates that, on average, using AI tools can boost individual productivity by around 30%, and some studies find improvements exceeding 50% in specific tasks.[1,2] Many of these studies, conducted in controlled environments where one group of workers is given access to the tool and another is not, find vast productivity improvements in tasks where the technology has a direct application, such as programming or writing.

These results should not be read as an automatic estimate of the impact on the entire economy. Firstly, they focus on specific tasks, and secondly, they often exclude implementation costs (training, process adaptation, organisational changes, legal or technical frictions). In short, they show what AI can do under favourable conditions, but not necessarily what it will do immediately on an aggregate scale. Even so, they represent a floor. As the technology advances, further improvements can be expected, and the evidence available to date suggests a rapid rate of improvement. Many of the available studies, for example, were conducted before the arrival of autonomous AI agents capable of executing complete tasks without human intervention; if this type of solution becomes widespread, then the productivity gains could increase substantially. There is also a pattern that is repeated in many jobs: among workers performing the same task, AI tends to be of greater help to those who started with a lower level of productivity. In this regard, it acts as a «leveller».

Estimated productivity gains due to AI
(pps)

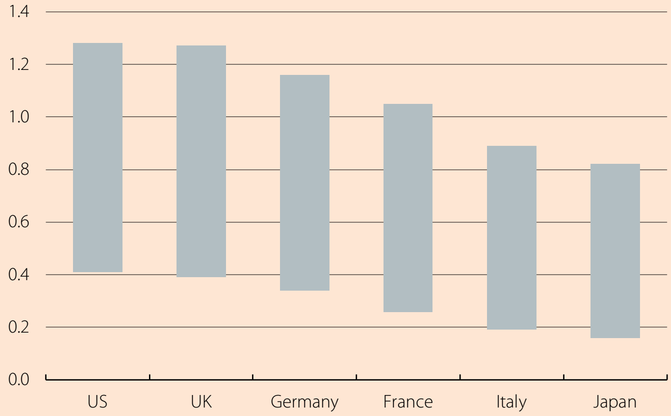

Notes: *Average annual increase over the next 10 years. The range captures differences in the speed of the technology's adoption and differences in the sectoral structure of each country.*
Source: *CaixaBank Research, based on data from the OECD.*

The leap from micro to macro is not automatic

Small-scale advances do not always translate into macro figures. If AI were to have a big impact on a few occupations, then the overall impact could be limited. The Nobel laureate in Economics, Daron Acemoglu, proposes a simple framework for considering this leap.[3] AI boosts productivity in two ways: it automates tasks (replacing human labour) or it complements workers (enabling them to do more and better). Both increase productivity, but with different implications for employment, wages and inequality.

Under certain assumptions, the author shows that the impact of AI on aggregate productivity can be approximated based on two ingredients: (i) the proportion of tasks or occupations actually affected by the new technology and (ii) the average productivity gain in those tasks.[4] Unfortunately, there is significant uncertainty regarding the magnitude of both of these ingredients.

1. «Macroeconomic productivity gains from Artificial Intelligence in G7 economies», OECD Artificial Intelligence Papers, June 2025, nº 41.
2. The productivity metric varies according to the study. In some cases, it refers to time savings, while in others it refers to increases in production within the same time interval. In general, they can be interpreted as savings in labour costs.
3. D. Acemoglu (2025). «The simple macroeconomics of AI». Economic Policy 40, nº 121, pages 13-58.
4. Economic literature distinguishes the concept of a task from that of an occupation. An occupation is a set of tasks, and the automation of a task does not necessarily mean that the occupation is automated. For the sake of simplicity, in this article we will use the words task and occupation as synonyms.

For instance, Acemoglu assumes that 20% of tasks are susceptible to being automated and that, among these, it will only be economically viable to automate 23% of them within the next 10 years. Other authors find higher figures, with 60% of tasks being susceptible to automation and 80% viability among these cases.[5]

Aggregate estimates vary significantly depending on assumptions about the proportion of tasks affected and the average productivity gains. At one end of the spectrum, Acemoglu suggests modest productivity gains of around 0.1 pp per year. With more favourable assumptions, the figures are higher. For example, the OECD estimates that, over the next 10 years, annual productivity growth will increase by between 0.4 and 1.3 pps in the US and by between 0.2 and 0.8 pps in other advanced economies.[6] These are broad ranges, depending on different assumptions about the speed of the technology's adoption and the sectoral structure of each economy, but in no case are they negligible figures.

These exercises do not exhaust all impact channels. AI can facilitate new occupations and business models, and could accelerate scientific innovation. The OECD, for example, notes signs of a virtuous circle of innovation: there is an increase in generative AI patents cited in developments in other fields and, in turn, an increase in generative AI patents that cite innovations from other fields that cited generative AI patents.[7] In other words, AI facilitates innovation in other fields and these accelerate innovation in AI itself.

The adverse effects also need to be included. The economy does not always function as the sum of isolated tasks. A simple example is the so-called Baumol effect: if productivity increases significantly in some sectors but little in others, wages tend to move similarly across sectors. Otherwise, workers would end up moving to where the pay is better. In order to retain them, less productive sectors have to raise wages, even if they do not produce more. The rise in wages in these sectors translates into higher prices and, therefore, the weight of these sectors on final expenditure increases and dilutes the impact of the productivity gains in the more advanced sectors. OECD simulations suggest that this effect could subtract around one-sixth of the potential increase in productivity growth associated with AI.[8]

Furthermore, AI can have harmful uses – disinformation, manipulation, cyberattacks or addictive advertising – that generate negative externalities. If these costs are not reflected in standard metrics, macro gains may overestimate the social benefits.

The labour market: a great unknown

The net effect of AI on employment is ambiguous. On the one hand, automation reduces the demand for labour in the affected tasks. On the other hand, new technologies also create new jobs – the reinstatement effect. This is an important channel. In the four decades following the Second World War, the emergence of new occupations completely offset job destruction due to automation.[9] The big question is whether AI will replicate that pattern and at what pace. There is also a third channel: by boosting productivity, AI could result in lower costs, lower prices, and better products, which could stimulate demand and, therefore, the demand for labour too.

Wage inequality does not follow a single direction either. Unlike other technological waves, such as robotics, which disproportionately affected certain groups, exposure to AI seems to be relatively widespread across occupations of different skill levels, potentially limiting the increase in wage inequality. The IMF notes, however, that higher-income workers are, on the one hand, those at greater risk of having their jobs replaced by AI, but at the same time, those with more potential to benefit from its complementarity.[10]

The institution simulates three scenarios and finds that the effect of AI on wage inequality depends on whom it helps and whom it harms more: if task substitution dominates, inequality could decrease (because higher-paid jobs would be more affected). If complementarity prevails, inequality would tend to increase (because workers with higher qualifications would benefit more). And if AI increases aggregate productivity, wages can rise for everyone, but more so for those who benefit from greater complementarity with AI, once again widening the gap.

5. For a review of the estimates made, see P. Aghion and S. Bunel (2024).«AI and Growth: Where do we Stand?», Policy Note.
6. See footnote 1.
7. «Is Generative AI a General-Purpose Technology? Implications for Productivity and Policy», OECD Artificial Intelligence Papers, June 2025, nº 40.
8. The impact is greater the more unequal the productivity gains are between sectors and the greater the difficulty households have in redirecting their spending towards more productive sectors.
9. D. Acemoglu and P. Restrepo (2019) «Automation and new tasks: How technology displaces and reinstates labor», Journal of Economic Perspectives 33, nº 2, pages 3-30.
10. M. Giovanni, A. Panton, C. Pizzinelli, E. Rockall and M.M. Tavares (2024). «Gen-ai: Artificial intelligence and the future of work». IMF, 979, pages 1-37.

Competition will be a key element

The distribution of the benefits will also depend on the competitive environment. AI can reduce barriers to entry in some markets. Cheaper tools for programming, translating, designing or analysing data can enable small businesses to do things that previously required greater scale. In competitive markets, part of the profits would translate into lower prices and more widely distributed benefits. If, on the other hand, companies capture most of the income – through patents or market power – then the distribution may be unequal.

This tension is particularly relevant in the AI market itself. Economies of scale (larger size, greater efficiency), economies of scope (a single model can be adapted for multiple uses at a relatively low cost) and bottlenecks in access to data to train models, as well as the cost of computing and human capital, naturally drive this market towards greater concentration. It is not inevitable, but it is a plausible risk. Oversight by authorities will therefore be important: not to hinder innovation, but to prevent a technology with the potential to enhance well-being from being sequestered by excessively closed market structures.

In summary, AI will be transformative. Its potential to increase productivity is real, but its deployment will be gradual. Initially, time-saving in specific tasks will prevail. The most significant changes will come later, when companies redesign entire processes and when AI helps accelerate the generation of knowledge and new ideas.

The most reasonable scenario is thus one of increasing benefits in the medium term. This increase will be more intense and rapid in the US than in Europe, given the faster pace of technological adoption and the prominence of the tech sector in the US compared to Europe.[11] In this context, it seems plausible to expect productivity improvements of up to 1 pp annually in the US over a 5 to 10-year horizon, and about half of that in Europe. This would not be an instant revolution, but it would represent a step change for growth.

Oriol Carreras Baquer

11. For further details, see the articles «Artificial intelligence: a supply-side perspective» and «Differentiated strategies for governing AI: towards cooperation or conflict?», in this same Dossier.

The AI buzz in financial markets

Artificial intelligence (AI) has accounted for much of the recent economic growth[1] and stock market performance in the US. Since the emergence of ChatGPT three years ago, the so-called Magnificent Seven[2] account for 60% of the cumulative increase in the market capitalisation of the S&P 500 and now represent around 35% of the index. The rise of AI has led to hopes of a new industrial revolution and, at the same time, fears of another bubble. This ambivalence extends to stock market valuations: they rest on expectations of vast revenue growth, but at the same time, there are doubts about their sustainability, either because the expectations themselves may disappoint or due to the eye-watering spending and investment plans being drawn up by firms in the sector.[3]

Magnificent Seven: stock market capitalisation and valuation

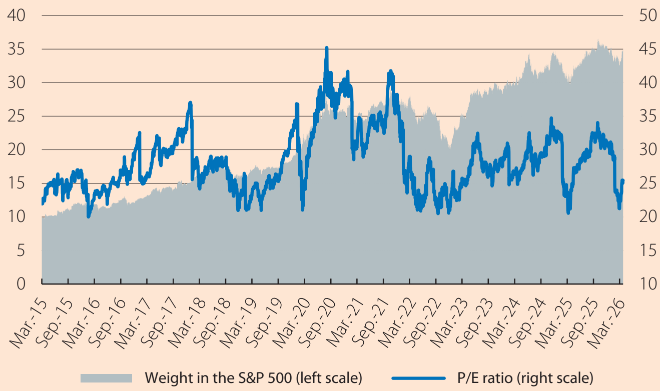

Notes: *The Magnificent Seven are Alphabet, Amazon, Apple, Meta, Microsoft, NVIDIA and Tesla. The P/E ratio is that between the observed market capitalisation and the expected earnings over the next 12 months (Forward Price-to-Earnings ratio). A higher P/E ratio can suggest both higher expectations of long-term earnings and the risk of over-valuation of the stock.*
Source: *CaixaBank Research, based on data from Bloomberg.*

The role of market structure

Although the Magnificent Seven are considered global leaders in AI today,[4] one of the questions for determining whether they will be able to monetise their vast investment plans in time is the shape that the AI market will take and which companies will emerge as winners and losers when the technology matures.

The AI value chain provides insights into the potential evolution of the market. This chain has five layers.[5] Firstly, computing power, with the design of microprocessors and memory chips that manage intensive calculations, where NVIDIA currently stands out in design and TSMC in production. Secondly, infrastructure, with data centres and cloud services and a notable presence of Amazon (Amazon Web Services) and Microsoft (Azure). Thirdly, datasets (images, text, audio) for training AI. The fourth link is large AI models, such as OpenAI's GPT or Anthropic's Claude, which can be adapted to a wide range of tasks. The final layer is applications that adapt these large models for specific use cases, such as Copilot, ChatGPT or Claude Code.

US: P/E ratio of leading AI companies

(Share price divided by expected earnings per share over the next 12 months)

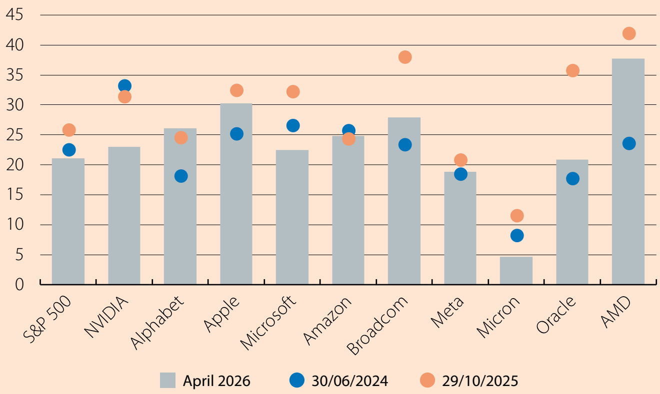

Note: *Companies ranked by market capitalisation from left (highest) to right (lowest). These companies dominate the Bloomberg Global Artificial Intelligence C-Series Index (the index does not include Tesla among its constituent stocks).*
Source: *CaixaBank Research, based on data from Bloomberg.*

Generally speaking, the current technology involves significant investment needs and high fixed costs, which can create barriers to entry in the AI value chain. This applies especially to the first two layers, where network effects are also observed that reinforce the competitiveness of incumbent firms against potential new entrants. The last three layers (data, models, and applications) are, *a priori*, more open to competition (training with public data, open-source code for models and applications), but they also exhibit dynamics that may favour market concentration. For example, if public data as a training source become exhausted (something various experts see as imminent), it will be necessary to resort to private data. Here, established companies such as Meta, Google or Microsoft could reinforce their competitive advantage thanks to the high volume of users of their applications (social networks, such as Instagram or LinkedIn, navigation [Google Maps] or office software [Microsoft 365]).

The complementary connections between the different layers in the chain also favour the dominance of firms that encompass multiple layers in the AI value chain – an integration already exhibited by the established big tech firms. For example, Google also produces its own hardware (TPU chips), builds models (Gemini) and integrates its products with one another.

1. See the article «Productivity and employment in the face of generative AI: what do we know?», in this same Dossier.
2. Alphabet (Google), Amazon, Apple, Meta, Microsoft, NVIDIA and Tesla.
3. This ambivalence is reflected in the P/E ratios (price per share divided by earnings per share, a standard valuation metric) shown in the first two charts: tech firms have above-average P/E ratios, but they have experienced corrections in recent months.
4. J. Frost, K. Rishabh and V. Shreeti (2026). «Global giants in the AI supply chain», Bank for International Settlements.
5. L. Gambacorta and V. Shreeti (2026). «The AI supply chain», Review of Network Economics.

A step change in investment needs

AI not only requires cutting-edge research but also vast investment in infrastructure, particularly related to the computational needs for storing data and training and using models. This investment includes data centres, computer servers, cooling systems, energy facilities, etc. Among the Magnificent Seven, this investment ambition has translated into capex growth of 50% and 60% in 2024-2025, accelerating to 70% in 2026, according to estimates and forecasts by Bloomberg's analyst consensus.

Magnificent Seven: investment in capital goods and revenues

(USD billions) *(USD billions)*

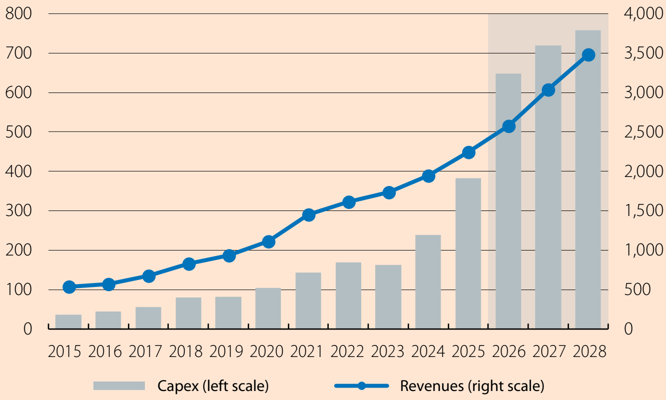

Notes: *The Magnificent Seven are Alphabet, Amazon, Apple, Meta, Microsoft, NVIDIA and Tesla. Forecasts from 2026 onwards.*
Source: *CaixaBank Research, based on data and forecasts from Bloomberg.*

The sharp growth in investment has led to a shift in financing strategies. In recent years, tech firms have taken advantage of their low debt ratios and highly profitable operations to fund their investments with the cash flow they themselves generated. But their spending plans have grown so much that they have become more reliant on external financing (corporate bonds, loans and private credit and venture capital[6]).[7]

A common structure for obtaining external financing links data centres, private equity, and cross-investments among large AI firms.[8] Typically, this formula involves forming a consortium of actors to create a new entity, which will own data centres. That consortium includes, as a minority shareholder, the AI company itself which will pay the rent and operate the data centres. To obtain financing, the entity issues debt, often channelled through private credit[9] or institutional investors, and the servicing of that debt is backed by the income generated from the rental of the data centres. According to the Bank for International Settlements itself,[10] this structure can create circularity and opacity around AI firms' actual level of indebtedness. Furthermore, it tends to create links between the established major AI firms (when they converge in consortia), just like in other cross-investment operations among the leading companies.[11]

US: AI investments through venture capital

(USD billions)

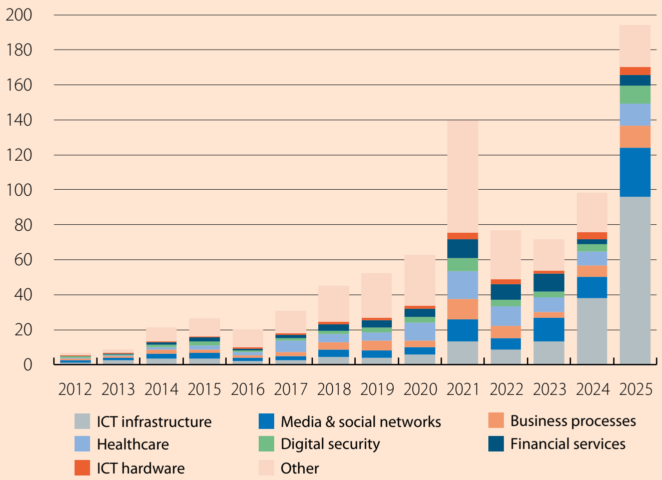

Source: *CaixaBank Research, based on data from the OECD.*

Overall, the current AI value chain and the financing strategies and strategic deals between established tech firms exhibit traits that are conducive to market concentration and dominance by incumbents. Besides helping to explain current market valuations, such concentration could pose a source of instability.[12] Nonetheless, how the AI market will actually develop remains uncertain and could result in very different configurations. Regulation, the ease of building new models, and supply dependencies (such as specialised chips) will be key to determining its final structure.

6. Venture capital is a form of investment that involves providing capital to new or growing companies with a perceived high long-term growth potential.
7. I. Aldasoro, S. Doerr and D. Rees (2026). «Financing the AI boom: from cash flows to debt», Bank for International Settlements.
8. Eren *et al.* (2026). «Financing the AI infrastructure boom: on- and off-balance sheet borrowing», Bank for International Settlements.
9. i.e. non-bank credit granted by specialist investment funds, negotiated directly between lender and borrower.
10. Eren *et al.* (2026), *op. cit.*
11. Bloomberg (2026). «A Guide to the Circular Deals Underpinning the AI Boom», describes various such circular arrangements. e.g. in 2025 NVIDIA agreed to invest 100 billion dollars in OpenAI, while OpenAI committed to operating its data centres intensively with NVIDIA chips. OpenAI and AMD also formed a strategic alliance whereby OpenAI could end up becoming a major shareholder of AMD and, at the same time, committed to purchasing AMD chips worth tens of billions of dollars.
12. For example, by exposing a large portion of the economy to the difficulties of a handful of agents or to bottlenecks, or by increasing the correlation between agents (e.g. correlated market movements that amplify moments of stress). S. Breeden (2024), *Engaging with the machine: AI and financial stability*, speech at the HKMA-BIS Joint Conference on Opportunities and Challenges of Emerging Technologies in the Financial Ecosystem.