



BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

Artificial intelligence and the US economy:
an accounting perspective on investment and production

by Luisa Carpinelli, Filippo Natoli and Marco Taboga

March 2026

Number

1006



BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

Artificial intelligence and the US economy:
an accounting perspective on investment and production

by Luisa Carpinelli, Filippo Natoli and Marco Taboga

Number 1006 – March 2026

The series Occasional Papers presents studies and documents on issues pertaining to the institutional tasks of Banca d'Italia and the Eurosystem. The Occasional Papers appear alongside the Working Papers series which are specifically aimed at providing original contributions to economic research.

The Occasional Papers include studies conducted within Banca d'Italia, sometimes in cooperation with the Eurosystem or other institutions. The views expressed in the studies are those of the authors and do not involve the responsibility of the institutions to which they belong.

The series is available online at www.bancaditalia.it.

ARTIFICIAL INTELLIGENCE AND THE US ECONOMY: AN ACCOUNTING PERSPECTIVE ON INVESTMENT AND PRODUCTION

by Luisa Carpinelli*, Filippo Natoli** and Marco Taboga***

Abstract

Artificial intelligence (AI) has moved to the centre of policy, market and academic debates, but its macroeconomic footprint is still only partly understood. This paper provides an overview on how the reverberations of the AI wave are captured in US national accounts, combining a simple macro-accounting framework with a stylized description of the AI production process. We highlight the crucial role played by data centres, which constitute the backbone of the AI ecosystem and attracted formidable investment in 2025 to meet the worldwide demand for AI services. In the first three quarters of the year, the boom in IT and AI-related capital expenditure clearly boosted aggregate US demand; the contribution to GDP was more limited, as a result of the high import incidence in AI hardware. Furthermore, simple calculations suggest that, at current utilization rates and pricing, the production of services originating in new AI data centres could boost GDP in the near future on a scale comparable to that recorded so far by investment spending. Short reinvestment cycles and uncertainty about future AI demand, while yet to act as a drag on growth, can nevertheless fuel macroeconomic risks over the medium term.

JEL Classification: E2, O3, O4.

Keywords: artificial intelligence, capital expenditures, data centres, national accounts.

DOI: 10.32057/0.QEF.2026.1006

* Bank of Italy, Economics, Statistics and Research Department; luisa.carpinelli@bancaditalia.it

** Bank of Italy, Economics, Statistics and Research Department; filippo.natoli@bancaditalia.it

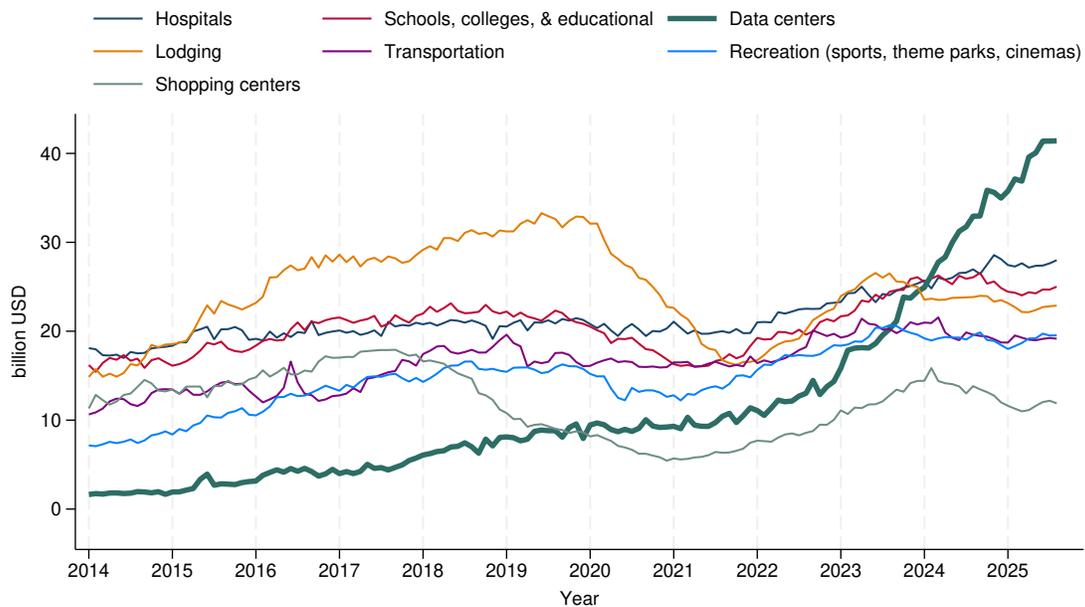
*** Bank of Italy, Economics, Statistics and Research Department; marco.taboga@bancaditalia.it

1 Introduction

The extraordinary rise of Artificial Intelligence (AI) has moved to the center of policy, market, and academic debates. From a financial stability perspective, a recurring question is whether the sharp appreciation in the share prices of technology companies is justified by a credible and sustainable expansion in their underlying profitability. At the same time, the key macroeconomic issue is whether – and how much - AI will leave a durable imprint on economic growth, productivity, and employment. A definitive empirical assessment of the long-run productivity effects of generative AI remains premature. These questions are not new. They echo earlier debates on the productivity payoff of information technology in the 1990s (Brynjolfsson and Hitt, 2000) and on the macroeconomic role of intangible capital (Corrado et al., 2005).

What is distinctive in the current episode is the sheer scale of investment, connected to the massive surge in data center construction (Figure 1). These numbers have attracted a large attention, leading some commentators to identify AI investment as a dominant driver of the 2025 US GDP growth, which was solid above expectations.

Figure 1: PRIVATE CONSTRUCTION SPENDING BY SELECTED INDUSTRIES.



Source. U.S. Census Bureau, Value of Construction Put in Place. Last observation: August 2025.

This paper provides a tentative assessment of how the current AI wave is already contributing to US GDP. Adopting a deliberately mechanical accounting approach and a short-term horizon, we abstract from second-round effects that may eventually trigger deeper structural changes. Also, we admittedly abstract from the financing structure of the underlying investment, as well as from its cost and long-term financial sustainability – dimensions that have attracted increasing media attention. We combine a macro-accounting perspective with a detailed description of the AI production ecosystem to answer a simple question: how, and through which channels, is the current AI wave showing up in measured US GDP? We focus on the parts of the AI value chain where investments is already large enough to matter for the short-term growth dynamics -- notably investment in data-center infrastructure and the provision of cloud and AI services. We also attempt to distinguish between domestic value added and AI-related spending directed at foreign products.

Although conjunctural, this assessment is of first-order importance: quantifying AI's current macroeconomic contribution informs policy-making in many aspects. For monetary policy, it can improve the understanding of inflationary pressures and labor market dynamics – for instance, by helping to reconcile the apparent dichotomy between resilient overall growth and a deteriorating labor market. It can provide a clear mapping of how AI investment and production enter GDP is essential to evaluate the strength and composition of demand, to gauge the domestic versus foreign share of value added in AI supply chains. It can feed into fiscal and industrial policies aimed at supporting high-value activities. Finally, it also matters for financial stability analysis.

Any quantification is nevertheless fraught with methodological and measurement challenges. First, isolating AI-related activity in national accounts is difficult because the whole AI production is embedded in broader aggregates – such as IT equipment, semiconductors, software, R&D, and IT services. In practice, statistics frequently allow to capture only a broad “technology” sector rather than AI specifically. Second, firm-level figures on AI revenues and expenditures are also rarely observable: key players in the AI production ecosystem are large, multi-activity firms whose financial statements combine costs, revenues, and profits across many lines of business. As a result, isolating the contribution of AI to costs, revenues and profits is extremely challenging. Third, the broader “second-order” effects triggered by

the current AI transformation – such as productivity gains and employment reallocation associated – are also difficult to assess, absent reliable and standardized information on the link between AI adoption and employment. Our accounting framework will deliberately abstract from these effects, and mechanically evaluate how AI initially manifests itself across GDP components.

Three findings stand out. First, the boom in IT and AI-related capital expenditure has provided exceptional support to aggregate demand, but its net contribution to US GDP growth is more limited than headline investment figures suggest, because a large share of the associated hardware is imported. This reflects the fact that activities like chip design, data center operation, cloud services, and model development mainly occur in the US, but only part of the value created along the AI hardware supply chain accrues to the US, as the fabrication and assembly of servers and accelerators are concentrated abroad.

Second, data centers are the backbone of the AI ecosystem. They aggregate the computational capacity required for AI model training, inference, and cloud-based services, converting AI infrastructure investments into a service. From a macroeconomic standpoint, the construction of data center facilities generates significant demand for construction services, high-performance hardware and ancillary inputs. At the same time, the computational services they provide enter in the national accounts as final consumption (e.g. households, firms or government purchasing a cloud/AI services) and constitute intermediate inputs for other industries (e.g. a bank using cloud AI to run operations). Hence, data center activity is a pivotal channel through which AI investment feeds aggregate demand and contributes to GDP growth, making it a central focus of any macroeconomic analysis of the AI wave.

Third, according to our analysis, as the new generation of AI data centers becomes operational and starts selling computational and AI services to final users, the resulting revenue stream will add to GDP at a scale comparable to the capital expenditure needed to expand the underlying AI production capacity. This follows from the high utilisation rates of existing data center infrastructure, the current pricing of GPU-based services, and the short pay-back periods estimated in the paper.

The rest of the paper is organized as follows. Section 2 provides a roadmap to the AI industry, describing the main actors in the ecosystem – hardware vendors, providers of

cloud infrastructure and AI labs – and highlighting the central role of data centers and global supply chains. Section 3 turns to AI investment and quantifies how much of the strength in US demand and output in the first three quarters of 2025 can be attributed to technology-related capital expenditures once imports and the sectoral distribution of value added are taken into account. Section 4 describes the production of AI services using the example of an AI data center to illustrate the channels through which computational services enter national accounts. Sections 5 and 6 discuss future developments and draw conclusions.

2 A roadmap to the AI industry

In this section we provide a stylized description of how the AI industry is structured, which helps to understand the channels through which AI activity affects the economy.

The AI ecosystem is organized around three often overlapping groups of actors: 1) hardware vendors, who design, build and sell the computer equipment used in AI data centers; 2) cloud infrastructure providers, who own the data centers and rent out the computing capacity; 3) providers of AI software, who demand computing capacity to run AI models and sell the derived services to their customers.

We describe these groups paying particular attention to where they are headquartered and where different stages of production take place. This geographical dimension is crucial to understand how much of the value created is domestic output and how much accrues to foreign trading partners.

2.1 Hardware vendors

Modern AI workloads are computationally intensive and need to be run on specialized hardware (so-called AI chips), such as Graphical Processing Units (GPUs) and Tensor Processing Units (TPUs). By leveraging thousands of computational cores working in parallel, these chips accelerate the mathematical operations, such as matrix multiplication, that recur most often in AI-model training and inference. The acceleration factor is several orders of magnitude with respect to classical CPU (Central Processing Units) chips.

The production of AI chips – the most costly item in AI data centers – and their packaging into boards and server racks happens along a complex global supply network:

- Chip design: the design of AI chips is mostly performed by large US public companies such as NVIDIA (the leading player in this space), AMD and Broadcom (which co-designs proprietary chips with large data-center operators like Apple and Google).
- Chip manufacturing: once chips are designed, they are manufactured in semiconductor fabrication plants (so-called fabs). These are mostly operated by Asian firms such as TSMC (Taiwan) and Samsung (Korea), but some US plants have begun production on the back of public financial incentives provided by the Chips Act. Memory modules, expensive components tightly integrated in AI chips, are provided both by US manufactures (Micron) and Asian ones.
- Server assembly: packaged GPU modules go to server makers (e.g., SuperMicro, Hewlett Packard, Dell), who integrate them into full servers and racks comprising thousands of other less costly components (e.g., networking modules, mechanical, electrical and cooling parts).¹ Given the high number of components involved, this segment of the value chain is much more difficult to trace, although it is known to be dislocated not only in Asia and the US, but also in other countries that are geographically close to data-center locations (e.g., Mexico).²

This already complex and geographically spread supply chain is complemented by numerous other US and foreign companies both upstream (providers of chip manufacturing equipment; e.g., ASML, Applied Materials) and downstream (integrators and providers of data-center equipment; e.g. Schneider Electric).

From an economic perspective, these segments of the hardware industry function as oligopolies. They are dominated by a small number of firms that enjoy high profit mar-

¹The manufacturers either sell under their own brand (*Original Equipment Manufacturers*, OEM), or build as custom for someone else's brand (*Original Design Manufacturers*, ODM). Big-name server OEMs are Dell Technologies (PowerEdge), HPE (ProLiant, Synergy), Lenovo (ThinkSystem); major server ODMs are Quanta Computer (and QCT – Quanta Cloud Technology), Wiyynn (Wistron spinoff focused on cloud DC), Inventec, Foxconn / Hon Hai.

²For an overview of the semiconductor sector and related geopolitical issues, see Savini Zangrandi (2020) and Patrignani et al. (2025).

gins and are shielded by substantial barriers to entry. This is especially evident in chip manufacturing, where producing cutting-edge chips requires fabrication plants that not only cost tens of billions of dollars to build, but also demand highly specialized, hard-to-replicate know-how and technical expertise to fine-tune production processes and achieve economically viable yield levels.

2.2 Cloud infrastructure providers

Cloud infrastructure providers own and operate globally distributed data centers. They lease the computing capacity of these data centers to organizations who demand computing, file storage, networking and related services (Infrastructure as a Service - IaaS).

The largest providers – often referred to as hyperscalers – are Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP). Together, they account for 65% of the global cloud-computing market (OMDIA, 2025). Meta is sometimes classified as a “quasi-hyperscaler” because its global data-center footprint is comparable, although its infrastructure primarily supports its own applications. Other, smaller providers include Oracle Cloud, IBM Cloud, and CoreWeave.

These firms supply cloud infrastructure both for conventional services (e.g. web hosting and database management) and AI workloads such as model training and inference. Because hyperscalers do not provide a formal revenue and profit split between AI-related and traditional cloud services, it is difficult to isolate AI activity, especially since conventional services likely still represent the majority of their cloud revenues.

A common pricing model of cloud-infrastructure services is pay-as-you-go, whereby AI compute is billed per GPU-hour (or per second/minute), with separate charges for storage and network usage. These hourly prices are published by hyperscalers on their web sites in real time, which facilitates the assessment of demand and supply conditions in the market. However, pricing conditions are not always common knowledge, as, in addition to on-demand usage, services are often delivered through larger, negotiated contracts with enterprises. Such agreements typically involve committed spend (e.g., a minimum annual spend) and include enterprise-grade features like enhanced security, compliance assurances, administrative controls, data-handling guarantees, dedicated capacity, and dedicated support.

Cloud infrastructure providers incur a range of costs. These include (i) developing and maintaining the software layer that enables resource leasing, (ii) ensuring security, monitoring performance, and maintaining reliability, (iii) covering electricity and other operating expenses, and (iv) most critically, amortizing the substantial physical investment in data centres, which includes not only computing hardware but also land, buildings, and ancillary systems such as power, cooling, and networking. The cost associated to physical investments has become more and more relevant as AI-related services have gained traction. Although AI-related cloud services can appear similar to the traditional cloud services, the underlying infrastructure differs substantially. Traditional workloads run on CPU-centric server racks, whereas AI workloads rely on clusters built around large arrays of costly AI accelerators (GPUs and TPUs). These AI-focused facilities require much higher power density, advanced cooling, high-speed networking, and high-throughput storage. All of these factors drive up the costs of AI infrastructure substantially.

From a geographical standpoint, all of the players mentioned above (Amazon, CoreWeave, Google, IBM, Microsoft, and Oracle) are headquartered in the United States and build data centers worldwide. While the precise distribution of their data centers across countries is not publicly known, most estimates suggest that roughly 50% of global data center capacity is located in the US.³ A substantial share of the remaining capacity is in China, where several large domestic providers operate, most notably Alibaba Cloud, followed by Huawei Cloud, Tencent Cloud, and Baidu Cloud. However, it is unclear how this geographic distribution translates into the revenues and profits reported by US-domiciled subsidiaries, or how these figures are influenced by cross-border intra-group transfer pricing.

2.3 AI labs

AI labs are companies that train frontier AI models and sell services derived from these models. Some of the main labs (and their GenAI models) are OpenAI (ChatGPT), Anthropic (Claude), Google DeepMind (Gemini, Gemma), Meta AI (Llama, Meta AI assistant), xAI (Grok), Mistral, DeepSeek, Z.AI (GLM models), as well as the AI divisions of Chinese tech giants Alibaba (Qwen), Baidu (ERNIE) and Tencent (Hunyuan).

³See, e.g., Synergy (2025), Neufeld (2025), Ciolli (2025), Leichter (2025).

AI labs' revenues come mainly from selling the following two services:

- AI as a Service (AIaaS): AI laboratories expose their models through Application Programming Interfaces (APIs), enabling organizations to submit queries programmatically and receive model-generated responses. This kind of service is typically demanded by corporations that develop and sell AI-based or AI-enhanced software, or by firms that use AI to automate their workflows. For Large Language Models, the service is typically priced on a per-million-token basis (one token roughly corresponds to a single English word generated by a model).
- Software-as-a-Service (SaaS): AI models are integrated into user-facing applications that are sold on a per-seat basis through monthly or annual subscriptions, both to enterprises and to consumers. OpenAI's GPT-based and Google's Gemini chatbots, as well as Microsoft's Copilot products are some examples of popular SaaS offerings.

Another less common source of revenue are licensing contracts, such as those in place between OpenAI and Microsoft, which allow the latter to directly use OpenAI models in its services.

The largest costs of AI labs are associated with recruiting and retaining top talent and with procuring the enormous computational resources that are needed to train and serve models. Several labs (e.g., OpenAI and Anthropic) do not own large data centers and satisfy their computing needs by purchasing cloud infrastructure services (see previous sub-section).

The segment of AI labs remains particularly opaque. Many labs are often large vertically integrated conglomerates, such as Google and Meta, that do not disclose the revenues and profits of their AI research units. Other labs, dedicated almost entirely to model development (e.g., Anthropic, OpenAI), are privately held and therefore their financial performance is not publicly available. Moreover, there are often partnerships and cross-holdings between the two groups of labs/companies.⁴

From a geographic standpoint, the largest labs are controlled by companies headquartered

⁴A few examples. Between 2019 and 2023 Microsoft invested over \$13 billion in OpenAI, securing a claim on roughly 49% of OpenAI LP's future profits and integrating OpenAI's models deeply into Azure for cloud computing and Microsoft 365. Amazon and, with a lower investment, Google are partners/investors in Anthropic, that runs its models largely on AWS as its primary cloud provider.

in the US (Anthropic, Google DeepMind, Meta AI, OpenAI, xAI) and China (Alibaba, Baidu, DeepSeek, Tencent, Z.AI), with the notable exception of Mistral (France).

2.4 General remarks on the AI ecosystem

Although simplified, our depiction of the AI ecosystem highlights several features. First, the supply chain for the manufacturing and assembly of data-center hardware is truly global, while cloud infrastructure provision and AI services are more domestically concentrated.

Second, data centers form the backbone of the ecosystem, absorbing a large share of incremental and recurring capital expenditure, on top of upstream investments such as the construction of semiconductor fabs, and enabling almost all AI-related revenue streams. For cloud infrastructure providers, the data center is the product; for model labs, data centers are probably the largest input cost; for chip and server manufacturers, data centers are the primary deployment location for their products and thus the main source of demand. In this sense, AI production is essentially a data-center-centric industrial system.

Third, within this complex ecosystem, AI affects GDP through multiple observable channels: (i) data-center construction and AI R&D raise aggregate investment; (ii) chip and server production shapes trade flows due to globally fragmented supply chains; (iii) AI service provision increases value added and reallocates it across sectors; and (iv) final demand for AI services influences consumption and government spending. The next sections provide quantitative evidence about these effects.

3 AI Investment and Aggregate Demand

In the US real private fixed investment surged in the first half of 2025, (7.1% in Q1 and 4.4% in Q2); this expansion was followed by a more moderate one in Q3 (1%). These aggregate investment dynamics were largely driven by spending on *Computer and peripheral equipment* and *Software* investment. Over the first three quarters of 2025, these two categories recorded real-term quarter-on-quarter annualized growth rates of 104%, 62%, and 46%, and of 19%, 27%, and 3%, respectively. Taken together, they provided the largest positive contribution to aggregate investment expansion (5.2, 5.4 and 2.4 percentage points, respectively), offsetting

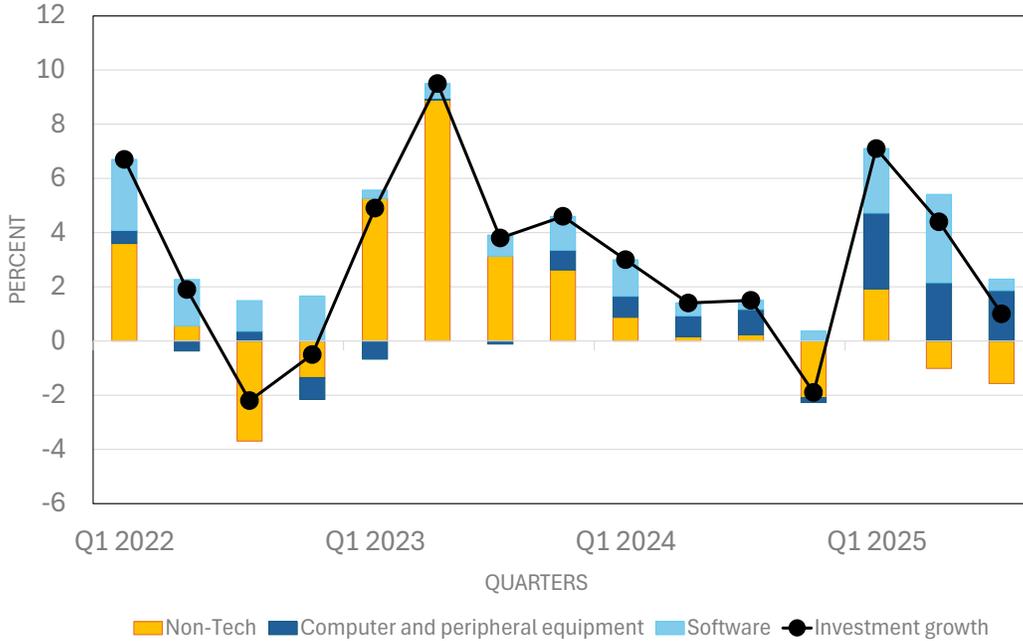
declines both in the residential and non-residential structure categories (Table 1 and Figure 2).⁵

Table 1: Technology-related investments in National Accounts.

| | 2025 Q1 | | 2025 Q2 | | 2025 Q3 | |
|---|-------------|-------------|-------------|-------------|------------|-------------|
| | q/q growth | contr. to I | q/q growth | contr. to I | q/q growth | contr. to I |
| Private fixed investment (I) | 7.1 | | 4.4 | | 1.0 | |
| Nonresidential | | | | | | |
| Equipment | | | | | | |
| <i>Computers and peripheral equipment</i> | 103.7 | 2.8 | 61.7 | 2.2 | 46.1 | 1.9 |
| Intellectual property products | | | | | | |
| <i>Software</i> | 18.8 | 2.4 | 26.6 | 3.3 | 3.1 | 0.4 |
| Info. processing equip. and software | 36.1 | 7.6 | 19.7 | 4.5 | 5.4 | 1.3 |

Note: Quarter-on-quarter growth rates are annualized and computed on real investment values. Contributions to quarterly annualized investment growth are expressed in percentage points. *Source:* U.S. Bureau of Economic Analysis.

Figure 2: INVESTMENT GROWTH AND CONTRIBUTION OF TECH VS NON-TECH ITEMS.

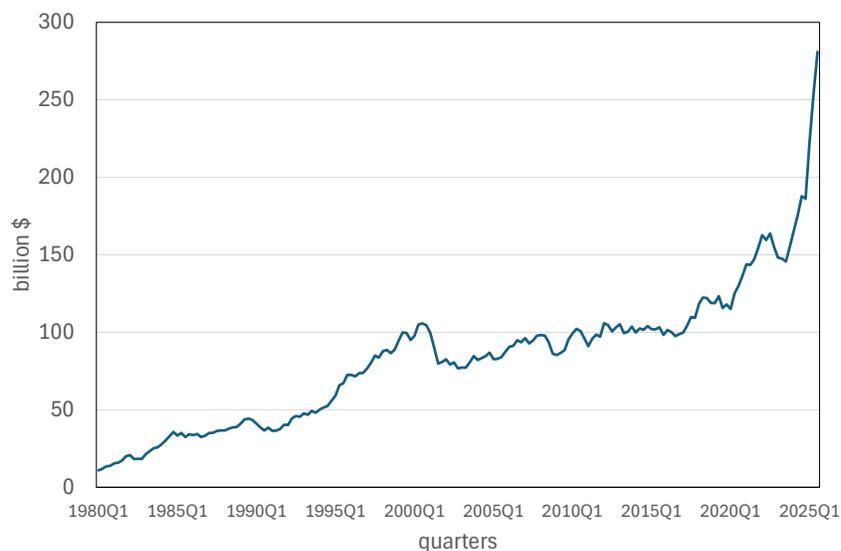


Note: Black dotted line indicates quarter-on-quarter annualized growth (percent); colored bars indicate contribution to growth (percentage points). Last observation: 2025 Q3. *Source:* U.S. Bureau of Economic Analysis.

⁵Software investment expansion, that contracted in Q3, was also less buoyant in the first half of the year, yet – since its level is roughly three times that of hardware – it has made a comparable contribution.

These figures are historically high. In particular, the surge in investment in *Computer and peripheral equipment* over the first three quarters of 2025 was unprecedented (Figure 3).

Figure 3: INVESTMENTS IN COMPUTER AND PERIPHERAL EQUIPMENT



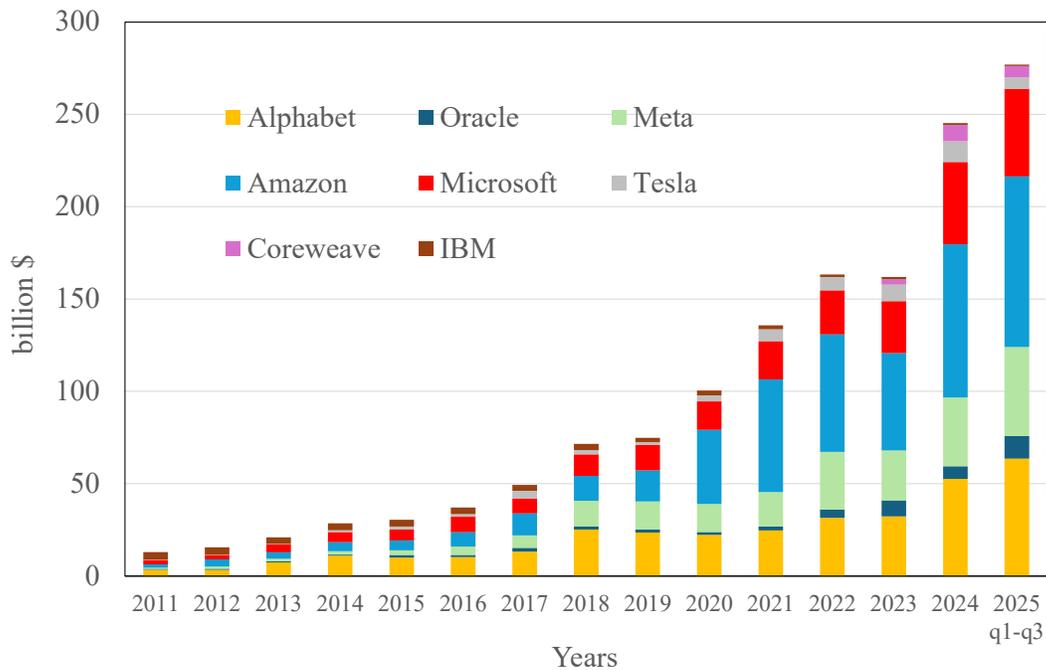
Last observation: 2025 Q3. Source: U.S. Bureau of Economic Analysis.

The soaring investment in this component is most closely linked to the massive expansion in digital infrastructure for AI. The component *Computers and peripheral equipment* indeed includes the purchase of servers (which embed AI chips such as NVIDIA’s GPUs), the costliest item related to the construction of data centers – other expenses being building construction and the installation of networking, cooling and power equipment.

Large tech companies made capital expenditures (capex henceforth) for \$245 billion overall in 2024, and more than \$275 billion in the first three quarters of 2025 (Figure 4).⁶

⁶Here we include both the largest firms selling cloud infrastructure/platform services to other organizations (Alphabet, Amazon, Microsoft, Oracle), two companies that are heavily investing in data centers but mostly for serving their own computational needs (Meta, Tesla), and also Corewave and IBM that have also developed data centers.

Figure 4: CAPITAL EXPENDITURES OF LARGE TECH COMPANIES.



Source: S&P Capital IQ. Capital expenditures up to 2024 are computed for fiscal years.

Most of the forecasts now place overall AI-driven capex by large tech companies around \$300-350 billion for 2025. Companies’ reports reveal indeed that a large chunk of this is related to the construction of data centers: for Microsoft, for example, nearly 100% of the quoted 2025 capital expenditure of around US\$80 billion is explicitly for AI/data-center infrastructure; for Alphabet, “the majority” of the US\$75 billion is for servers/data centers; for Amazon “the vast majority” of over US\$100 billion is directed at “data-center/AI infrastructure”.⁷

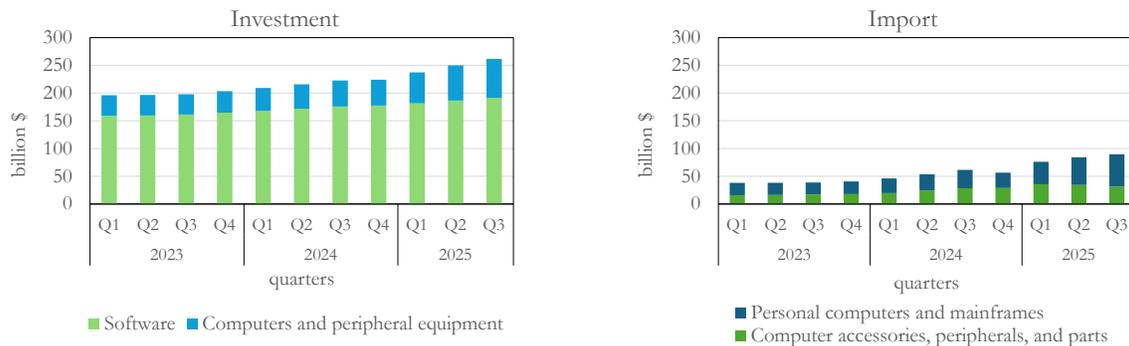
This surge in investment prompted many economists, policymakers and commentators to deem AI investment as the prime engine behind US economic growth, if not the only one. Harvard economist Jason Furman said that in the first two quarters of 2025, 92% of the

⁷Microsoft vice chair and President Brad Smith wrote in a blog post that in fiscal 2025 – which runs to 30 June – his company “is on track to invest approximately \$80 billion to build out AI-enabled data centres to train AI models and deploy AI and cloud-based applications around the world. More than half of this total investment will be in the United States”. On a company’s earnings call, Alphabet CFO Anat Ashkenazi said a majority of that spending would “go towards our technical infrastructure, which includes servers and data centers”. Amazon CEO Andy Jassy said on a company’s earnings call referred that they expected to spend over \$100 billion in capex for the full year, the “vast majority” of which would go toward AI.

increase in demand in the US economy was due to just two categories in GDP – information processing equipment and software. Nobel Prize Paul Krugman seemed to agree that the “AI boom is driving most, possibly all, of the economy’s recent growth”. J.P. Morgan said that AI-related capital expenditures contributed 1.1 percentage points to GDP growth “outpacing the US consumer as an engine of expansion”. A Deutsche Bank newsletter at the end of September read “AI machines – in quite a literal sense – appear to be saving the US economy right now”.⁸

However, an analysis of national accounts suggests that, while AI investment strongly sustained demand, its impact on GDP growth was more limited.⁹ The surge of investments was indeed mirrored by an equally striking increase in imports of technology goods in the first three quarters of the year ((+236%, +47%, and +24%), Figure 5): the flow of *Imported computers and peripherals* was more than half the value of total investment in *Information processing equipment*. The strongest expansion was recorded for the *Personal Computers and Mainframes category*, in which servers incorporating advanced GPUs and other high-end components enter. It should not be a surprise that this infrastructure is largely imported into the US: while chip design and cloud operations are largely US based - the main GPU producer, NVIDIA, is a US company - fabrication, packaging, and much of server assembly occur abroad, as described in Section 2.

Figure 5: INVESTMENT AND IMPORT OF TECH-RELATED PRODUCTS



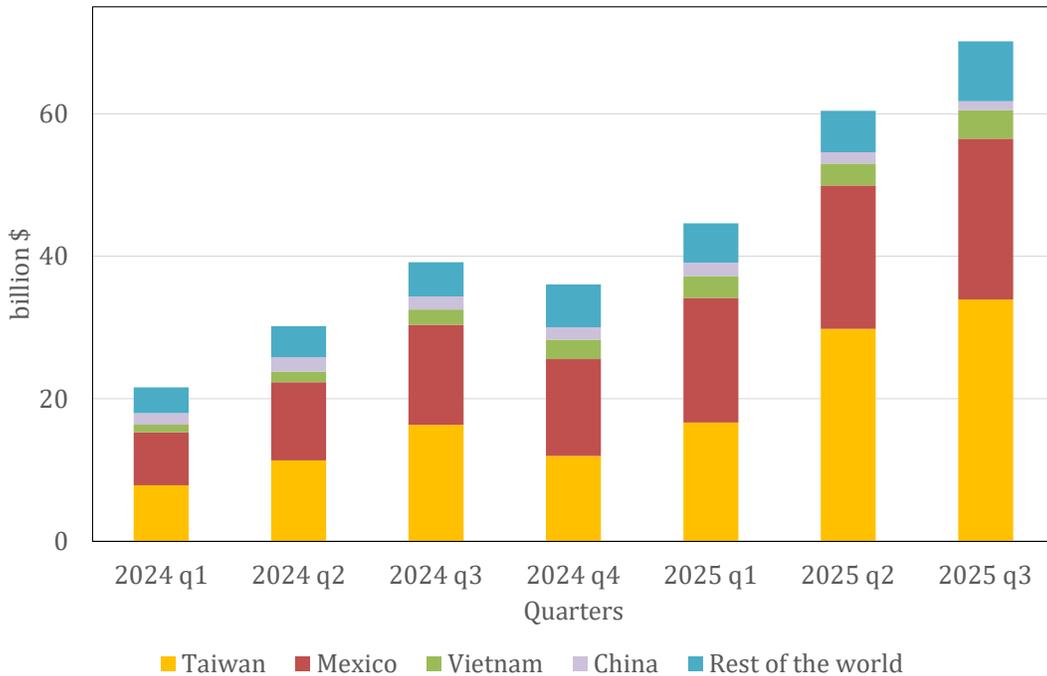
Note: Investment and imports of tech-related products, billions of dollars. Source: Bureau of Economic Analysis.

⁸See Jason Furman (2025); The New York Times (2025); JP Morgan Asset Management (2025).

⁹The role of imports is also acknowledged by Goldman Sachs (2025a,b) and Evercore ISI (2025).

The relevant role played by the imports of servers is confirmed by Trade Monitor data, providing a more granular breakdown of trade flows. Figure 6 shows a proxy of this category of imports: their inflow in the first three quarters of 2025 was almost double the nominal value imported in the same 9 months in 2024. Imports came largely from Mexico and Vietnam. In particular, Mexico’s growing role in the exports of AI-related products stands out, with this category now rivaling its traditionally key auto parts outflows, as also highlighted by specialized analysts (Capital Economics, 2025a,b)

Figure 6: IMPORT OF SERVERS BY COUNTRY OF ORIGIN



Note: Billion dollars. Source: Trade Data Monitor.

In principle, imports do not automatically feed onto investment, as they may be consumed domestically, re-exported, or purchased by the public sector. However, national accounts data provide little evidence that consumption or exports have played any major role in absorbing these AI-related imports. This is particularly evident in the first two quarters of 2025, when growth in both consumption and exports of these products remained subdued.¹⁰ Such weak performance suggests that the recent expansion in technology imports has been driven

¹⁰For technology consumption, we consider the personal consumption expenditure sub-categories “*Personal computers, tablets, and peripheral equipment*” and “*Computer software and accessories*”.

primarily by investment demand.¹¹ Once tech investment is purged by tech imports, its contribution to GDP growth falls markedly. Figure 7 displays GDP growth in the first quarters of 2025, together with the contribution of tech and non-tech items. Three considerations are in order. First, tech investment gave the largest contribution to GDP growth comparing with tech consumption and export, in all three quarters (1,1, 1,1 and 0,5 percentage points). Even under the extreme assumption that all technology imports are fully imputed to technology investment (netting investment contribution one-for-one with that of imports), the contribution remains sizable, given that these investments account for 3.4% of GDP. It is important to acknowledge that some measurement issues might lead to underestimate the contribution of these investments.¹² Second, technology imports in 2025, unlike non-tech ones, did not rebound after q1, suggesting that they were largely insulated from the front-loading that affected products expected to face tariffs. Third, GDP growth has continued to be led by the resilient American consumer (orange bar), who added 0.4, 1.7, and 2.0 percentage points to GDP growth in the first three quarters of the year, respectively.

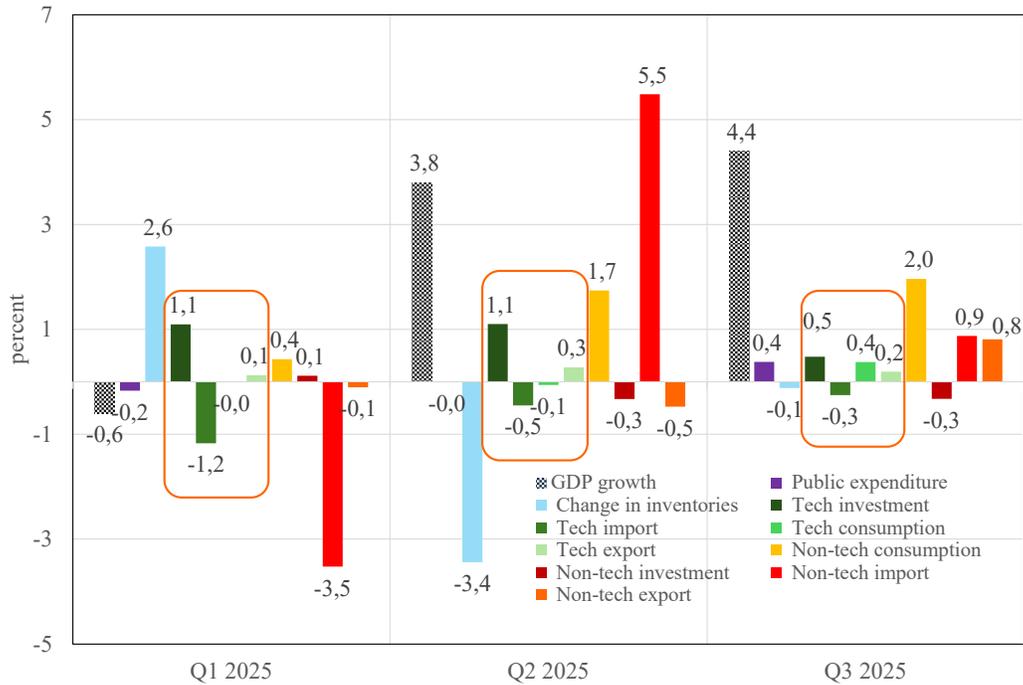
In the rest of 2025 and in 2026 capital expenditure is likely to keep sustaining production at least at the same pace: while US tech companies did not report specific targets for next year's capex, they all announced that they will be higher than in 2025.¹³

¹¹Unfortunately, data allowing us to identify the share of imported technology purchased by the public sector and recorded in government expenditures are not available.

¹²First, official statistics may understate AI investment if some key items, such as GPU servers, are misclassified as intermediate inputs rather than long-lived capital goods. Statistical agencies might struggle to: (i) measure GPU servers' price (due to rapid inflation of GPU costs), (ii) separate them from generic ICT imports – also due to purchases being based on confidential long-term contracts, (iii) account for their massive price-performance changes; (iv) record them in the detailed investment categories. Second, they may understate the US value added generated by AI-related firms whose hardware production takes place abroad, if the profits and other income created in the United States are not fully recorded as exports of services.

¹³Reporting q3 results, Alphabet, Amazon, Meta and Microsoft declared that their 2026 capex will be higher than 2025. For details, see Investors Business Daily (2025).

Figure 7: CONTRIBUTION TO GROWTH, TECH VS NON-TECH



Note: The orange rounded rectangles include the tech components of investment, import, consumption and exports.

4 From Data Centers to GDP: AI Production Channels

While Section 3 focused on the demand-side effects of AI-related capital expenditure and its import content, this section examines the broader question of how the ongoing production and use of AI services contribute to the supply side of the economy. Beyond the build-up of data center capacity, the development and deployment of AI show up in national accounts when agents invest in AI-related research and development (R&D), purchase AI-enabled services, and when these activities raise productivity and value added in technology-intensive sectors (Highfill et al., 2025).

As explained previously, most of the AI services production activity involves data center utilization. For our purposes, the practical consequence is that this activity can be seen, also from an accounting perspective, as a downstream flow going from data centers, which supply computational capacity, all the way to AI services. The flow is illustrated in Figure 8.

Figure 8: FLOW OF AI DATA CENTER SERVICES TO FINAL USE



Note: Light blue boxes indicate cases where the increased production of computational and AI services provides direct positive contributions to GDP. Plum-colored boxes correspond to cases where the increased production is used as an intermediate input and does not provide direct contributions to GDP.

There are two channels for computational capacity – and the AI services that immediately consume that capacity – to surface on GDP. First, in a very straightforward way, they can be sold for final use (consumption, investment, exports or government expenditure), thereby contributing directly to GDP on the resource-use side (the light-blue boxes in Figure 8). A typical example is consumers purchasing ChatGPT subscriptions. Purchases of computational capacity for conducting R&D – provided they are recorded as such in corporate accounts – also fall into this category, since R&D is classified as investment in national accounts; further examples of such transactions are listed in Table 2. This is the channel that appears on GDP with the shortest lag relative to investment.

Table 2: Examples of direct final demand contributions of data center and AI services

| Contribution to final demand | National account categories | Examples (hypothetical) |
|-------------------------------------|---|--|
| Consumption | Personal consumption of other services (NIPA 2.4.5.95) | Consumer buys personal ChatGPT subscription |
| Investment | Gross investment in intellectual property (NIPA 5.2.5U.19,47) | OpenAI trains a new model on Amazon’s cloud (R&D expense) |
| Government expenditure | Government consumption expenditures (NIPA 3.1.21) | Dept. of Defense runs its models on Microsoft’s cloud |
| Exports | Exports of computer services (NIPA 4.2.5B.86) | A foreign firm uses Google’s cloud (on US-based data center) |

Second, computational capacity and AI services can be used as intermediate inputs by domestic producers of goods and non-AI services (the plum-coloured elements in Figure 8). For instance, a car manufacturer might rely on AI-driven production lines whose inference tasks are carried out in an AI data center, or a company may subscribe to ChatGPT for general employee use. In these cases, computational services do not appear directly in GDP as final resource use, although they may substantially alter firms’ cost structures – for example by reducing labor costs and increasing spending on intermediate goods.

To the best of our knowledge, data on the precise distribution of computational services between final demand and intermediate use are currently unavailable. However, because the adoption of AI by non-AI-producing firms is still in its early stages, it is reasonable to assume

that, at present, a large share of computational and AI services is sold directly to final users and thus rapidly surfaces in GDP statistics.

To gauge the size of the potential contribution of AI services to final use, it is useful to consider two additional facts. First, despite very large investments, all available evidence (e.g., Goldman Sachs Research, 2025) points to high utilisation rates of new data centers and to persistent bottlenecks in meeting existing demand for computational capacity. The existing AI infrastructure is thus broadly at full capacity. Second, at current market prices for computational services (GPU rental), the investment payback times at AI data centers is on the order of one year. Appendix A illustrates this with a simple calculation based on current cloud rental prices for NVIDIA’s GB200 NVL72 architecture.¹⁴ Taken together, these two facts, along with the previous assumption that a large share of AI services is sold for final use, suggest that the yearly flow of revenues coming from the sale of AI services produced by these *fully utilised* AI centers is of the same order of magnitude as capex. Under continued rapid deployment, this mechanism has the potential to translate into sizable contributions to GDP growth.

Despite the scarcity of disaggregated data on AI activity, some additional evidence helps to corroborate this picture. The presence of a “revenue” channel is supported by GDP data based on value added, in which the GDP for a given sector reflects not only investment in capital goods, but also the flow of goods and services generated by the existing capital stock. Although AI-related activities cannot be cleanly separated from the broader IT sector – which likely leads to an overstatement of the contribution of the strictly defined AI industry – the relevant aggregates point to a strong contribution from technology-intensive sectors. In particular, value added in the *Manufacturing of computer and electronic products*, and in *Computer systems design and related services* recorded a formidable increase in Q2, likely influenced by AI developments. Importantly, the value-added approach delivers a slightly higher contribution of tech to GDP than the expenditure approach: 0.5 percentage points

¹⁴A modern high-density rack with 72 GPUs and 36 CPUs costs about USD 4.5 million to deploy once land, buildings, electrical systems, liquid cooling and other infrastructure are included, of which roughly two-thirds reflects the purchase of the server rack itself. Quoted annual rental prices for an equivalent rack from major cloud providers range from about USD 6.6 million to USD 10.1 million, while estimated yearly operating costs (electricity, maintenance, networking and other items) are on the order of USD 0.2 million. These figures are consistent with a payback period close to one year for fully utilised facilities.

in Q1 and 0.7 percentage points in Q2 (see Table 3), amounting on average to 40 per cent of GDP growth in the first half of the year (Figure A1),¹⁵ notably more than the figure recovered through the expenditure approach that only captures investment.

Table 3: Contributions to GDP growth by sector.

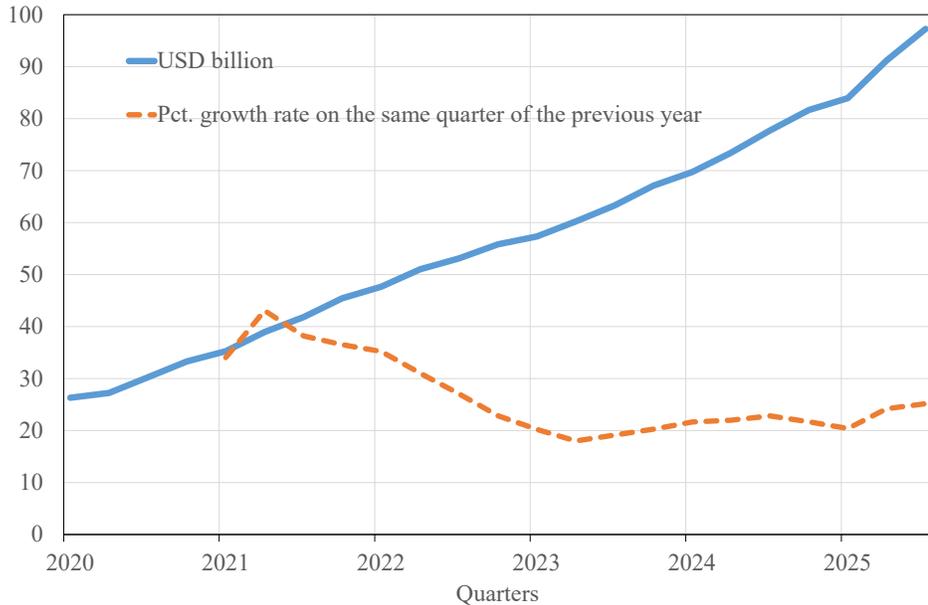
| | 2025 Q1 | | 2025 Q2 | |
|--|----------------|---------------|----------------|---------------|
| | q/q growth (%) | contr. (p.p.) | q/q growth (%) | contr. (p.p.) |
| GDP | -0.6 | | 3.8 | |
| Contribution to GDP growth: | | | | |
| Private Industries | | | | |
| Manufacturing | | | | |
| <i>Computer and electronic products</i> | 3.5 | 0.03 | 17.2 | 0.17 |
| Information | | | | |
| <i>Data proc., internet pub., other info. services</i> | 14.1 | 0.24 | 7.7 | 0.14 |
| Professional and business services | | | | |
| <i>Computer systems design and related services</i> | 12.2 | 0.21 | 23.4 | 0.38 |
| Total contribution | | 0.48 | | 0.69 |

Note: Quarter-on-quarter growth rates are annualized and computed on real value added values. Contributions to quarterly annualized investment growth are expressed in percentage points. “Data proc., internet pub., other info. services” stands for “Data processing, internet publishing, and other information services”.

Moreover, the revenues generated by the three largest public clouds that rent AI servers continue to grow at rates well above 20 per cent, despite already reaching impressive levels (around USD 350 billion over the past four quarters; Figure 9). These revenues represent a large share of the public cloud market (about 65 per cent according to OMDIA, 2025), but by no means the majority of the total data center market, which also includes large private clouds (e.g. those operated by Apple, Meta and Tesla) and smaller data centers owned by other corporations. Many of these private clouds are used to produce AI services sold to final users. Together with smaller public clouds, they likely make the aggregate volume of computational and AI services that may affect GDP much larger than the USD 350 billion figure cited above.

¹⁵At the time of writing, the Bureau of Economic Analysis had not yet released the data of Gross Domestic Product by Industry for Q3.

Figure 9: QUARTERLY DATA-CENTER REVENUES OF THE 3 LARGEST PUBLIC CLOUDS

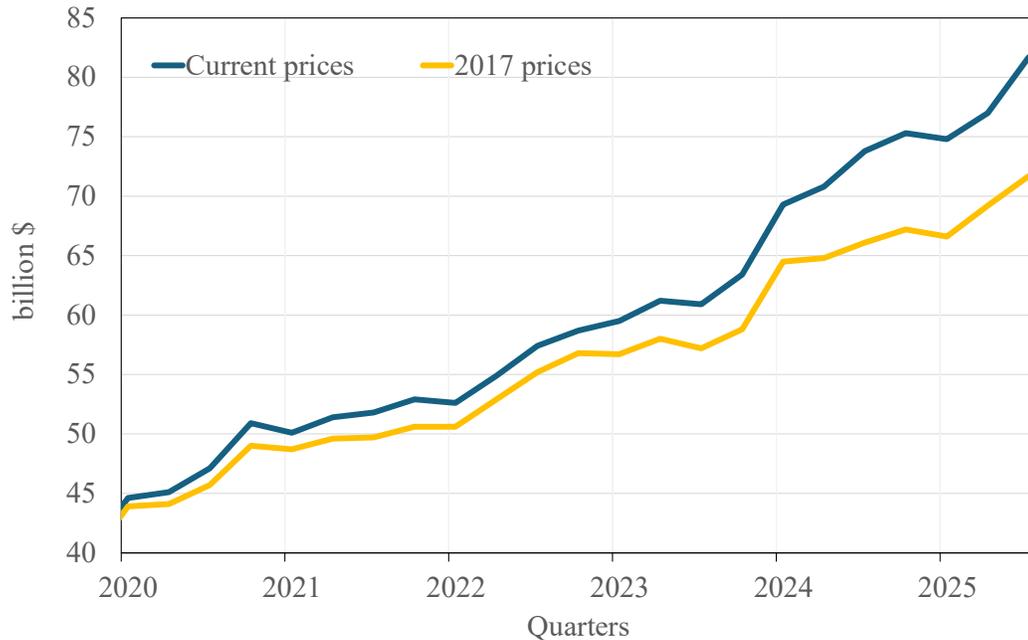


Note: The three largest public clouds are Amazon AWS, Google Cloud and Microsoft Cloud. Revenues generated by these clouds are reported as a separate item in their respective owner companies' financial statements. Source: quarterly financial reports of Amazon, Alphabet (Google), Microsoft and own calculations.

Finally, national accounts aggregates that provide more granularity also point to growing export flows of AI-related services. In particular, “Exports of computer services”, which includes cloud services, reached an annualized¹⁶ rate of USD 82 billion in the third quarter of 2025, up 11 per cent from the same quarter of the previous year (Figure 10). Although these data again bundle AI with broader digital services, they are consistent with the view that the production of computational and AI services is becoming an increasingly important driver of US domestic and external demand.

¹⁶Annualized rates reported by the Bureau of Economic Analysis roughly correspond to quarterly amounts multiplied by 4.

Figure 10: EXPORTS OF COMPUTER SERVICES



5 What to expect going forward

AI shows up in national accounts first as spending, and only later via productivity boost. It works mainly through expenditure-side channels: higher investment in software, data, and AI-related capital, increased consumption of AI-enabled services. Deeper effects connected to productivity gains and resource reallocation materialize over time, as organizations adapt, processes are redesigned, costs fall persistently, new products are enabled, and demand adjusts. It is these second-round changes operating like a positive supply shock—rather than the initial wave of AI investment—that will ultimately determine whether AI leaves a durable imprint on growth. This is the key question on AI and the macroeconomy.

However, also the phase of rising demand, with its staggering spending volumes, poses pressing macroeconomic questions. A first issue revolves around the depreciation of the AI hardware: servers and other computing equipment may have relatively short economic lifetimes because the continuous arrival of more powerful and efficient hardware makes ex-

isting capital obsolete fairly quickly.¹⁷ Technological obsolescence may be compounded by accelerated physical wear, driven by extreme thermal stress that would increase GPUs failure rates. Investors and commentators have pointed to these short lifetimes as a potential source of fragility. It would increase the risk that capital cannot be fully amortized through operating cash flows before the next technology cycle, if the revenue growth fails to translate into adequate returns on invested capital (ROIC).¹⁸ If insufficiently reflected in balance sheets, this dynamic could temporarily inflate reported profits, sustain excessive valuations, and heighten the risk of abrupt corrections once the effective depreciation is recognized.¹⁹ Moreover, frequent replacement would mechanically raise capex-to-revenue ratios and depress free cash flow, limiting the share of operating profits that can ultimately be distributed to investors—even in high-growth environments.

Still, some of these concerns may be overstated. As documented above, short economic life cycles are usually offset by equally short payback periods: the revenues and profits generated by a new data center typically exceed the initial outlay after only a few quarters, a pattern that appears to hold also for the latest generation of facilities, as discussed in the Appendix. In addition, GPUs may retain economic value well beyond their technological frontier, as older chips can be redeployed across less demanding tasks over time, effectively extending their amortization horizon. Taken together, these factors substantially mitigate concerns related to ROIC and accounting depreciation.

By contrast, the nexus between accelerated reinvestment cycles and free cash flow appears more robust. Although data center capex and the associated revenues have large effects on gross income, their impact on net income is likely to be much smaller when viewed over a sufficiently long horizon, given that investment costs must be incurred with a higher frequency than in more traditional investment frameworks. Furthermore, when data center investments are ramped up in response to a permanent increase in demand for computational services, this should not be seen as a temporary spike: it entails a sustained and rapid pace of replacement investment over time. It should also be noted that, to the extent that hardware production remains highly fragmented along globalized value chains, a substantial share of

¹⁷Although replacement cycles may nowadays be longer due to the slowing down of Moore’s law.

¹⁸See for example this discussion by Krugman (2025) and this Bloomberg interview to Chanos (2025)

¹⁹See this WSJ article by Weil (2025).

these large investment waves is likely to continue accruing to foreign suppliers and partners, rather than translating into commensurate gains for the domestic economy. As a result, the macroeconomic spillovers of the AI-related investment would continue to be more limited than suggested by headline capex figures for income generation at home.

The second critical issue revolves around the pace of development of AI services' demand. A distinctive feature of the current AI revolution, compared with earlier technological transformations such as the diffusion of personal computers in the 1980s and the Internet in the 1990s, is the extraordinary speed of adoption. As of end 2025, only three years have passed since the launch of ChatGPT (at the end of 2022), which marked the first major deployment of advanced generative AI services to a mass audience. Yet, according to most estimates (e.g., Microsoft AI Economy Institute, 2025), the global user base of these services has already surpassed one billion people. By contrast, following the birth of the World Wide Web and the release of the first web browsers in the early 1990s, it took roughly 15 years for Internet services to reach a comparable level of penetration.

This unprecedented pace of adoption has created significant challenges for AI providers, which must satisfy booming demand while securing sufficient computational resources. The wave of large scale AI investments observed in 2025, particularly in data center capacity, can be interpreted as a direct response to these pressures.

The outlook of future demand is therefore the key variable for AI firms' decisions. Yet it is surrounded by an extremely high degree of uncertainty. As Anthropic's CEO Dario Amodei noted in an interview with the New York Times (The New York Times Events, 2025), when demand for a service expands by around 1,000% annually for two consecutive years (in 2024 and 2025), it becomes nearly impossible to forecast the following year's (i.e., 2026) growth rate with economically meaningful precision. AI firms are therefore compelled to design investment and procurement plans under uncertainty. Risks are two-sided. On the one end of the spectrum, there is the much debated risk of overinvestment: if AI capacity were to exceed actual demand, the economics of AI production could deteriorate, potentially triggering market corrections and episodes of financial stress. In this gloomy scenario, emphasized by several commentators,²⁰ smaller and less established players, such as independent AI labs,

²⁰See, e.g., Brooker (2025), Laudani and de Rozario (2025), Letzing and Sung (2025).

could be more vulnerable, while larger conglomerates that can count on diversified revenue streams would likely be better positioned to absorb the shock, although their margins, which are currently supported by scarcity and buoyant demand, could shrink significantly. Among the large players, some that are almost exclusively focused on AI chip production could see their profits contract more rapidly. Repercussions on stock prices could be material and potentially induce second-order macroeconomic effects.

On the other end, there is the – often overlooked – opposite scenario of underinvestment. This configuration also presents substantial danger for AI firms. If planned computational resources ultimately were to fall short of a particularly strong expansion in demand, the resulting flow of AI services could be of lower quality²¹ and potentially unreliable²². Alternatively, there could be upward adjustments in selling prices. Either of these outcomes would likely drive users away and firms could risk ceding ground to competitors, including Chinese firms.

Which of these risk scenarios will materialize is difficult to assess, not least because limited transparency and the scarcity of detailed data on the evolution of demand for AI services constrain independent analysis and forecasting. Economic intuition suggests that the rapid, ongoing improvement in AI capabilities,²³ combined with the fact that enterprise adoption is still far from complete, could continue to sustain robust growth in demand. If so, the macroeconomic implications of AI will become even more significant than they are today and will warrant close scrutiny by policymakers, including dedicated efforts to collect more granular data on the AI industry.

6 Conclusions and policy implications

This paper has provided a first pass at mapping the macroeconomic footprint of the current AI wave in the United States. Bringing together a mechanical macro-accounting framework

²¹E.g., user queries could be routed to weaker models.

²²E.g., API requests could be throttled during peak hours.

²³Kwa et al. (2025), for example, show that a measure of the complexity of tasks that AI systems can tackle with satisfactory accuracy doubles every 7 months. Other metrics show similar dynamics. See, e.g., the model scores regularly published by the independent AI-benchmarking company Artificial Analysis on their web site artificialanalysis.ai.

and a stylised description of the AI production ecosystem, we have traced how AI-related activity already appears in US GDP through investment in digital infrastructure and the provision of computational and AI services.

Three broad messages emerge. First, AI investment has played a major role in sustaining demand and activity in the first three quarters of 2025, but it has not been the sole engine of US growth. Furthermore, once the high import content of AI-related hardware is taken into account, the contribution of technology-related investment to GDP growth, while still substantial, is more limited than suggested by headline spending and by some popular narratives.

Second, data centers are the backbone of the AI ecosystem. They should therefore be central to any macroeconomic analysis of AI's effects. On the one hand, the construction and equipment of AI data centers provide a considerable stimulus to aggregate demand and GDP through the investment channel. On the other hand, data centers enable a flow of computational and software services that almost immediately and significantly feed into final consumption and other components of aggregate demand.

Third, given high utilization rates and current pricing of GPU-based services, simple calculations suggest that the revenues generated by these facilities could soon support GDP increases of similar magnitude to the underlying capital expenditures. Sectoral value-added data and trade in computer services already point in this direction.

These findings have several implications for policymakers. For monetary policy, they caution against attributing US economic resilience solely to AI, as the resilient US consumer has remained a central driver of the economic dynamics. At the same time, they show that AI-related investment and production can meaningfully affect both demand and supply conditions. For fiscal and industrial policy, they highlight the importance of the location of high-value segments of the AI value chain, and the potential role of infrastructure, skills and regulatory frameworks in affecting location choices.

The analysis has clear limitations. It abstracts from longer-term productivity effects, from labor-market adjustments, and from distributional consequences of AI adoption; furthermore, it relies on sectoral aggregates that imperfectly isolate AI from the broader technology sector, and on partial information on firms' investment and revenue plans. Addressing these gaps

will require better statistics – including improved classification of AI-related hardware and services in national accounts, and greater use of micro-level data – as well as richer modelling of AI diffusion and its interactions with macroeconomic dynamics.

Nevertheless, the perspective developed here can serve as a framework for interpreting incoming data and for structuring future analysis. As AI technologies continue to evolve and diffuse, a clear understanding of how they map into the standard macroeconomic accounts will be essential for timely and well-informed policy decisions.

7 Acknowledgements and disclaimers

We are grateful to Riccardo Cristadoro and Giovanni Veronese for their insightful comments. The views expressed in this paper do not necessarily reflect those of the Bank of Italy or ESCB. The authors would like to thank Francesco Paolo Conteduca for sharing and support on the Trade Monitor Data.

References

- Brooker, N. (2025). Is there an AI bubble and will it pop next year? *Financial Times*. Accessed: 2025-12-15.
- Brynjolfsson, E. and Hitt, L. M. (2000). Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives*, 14(4):23–48.
- Capital Economics (2025a). AI boom fuelling Mexico’s exports. 27 November 2025.
- Capital Economics (2025b). Mexico: riding the US data centre boom. 13 August 2025.
- Chanos, J. (2025). Odd Lots: Jim Chanos on Who’s Getting Caught Swimming Naked. Odd Lots podcast, Bloomberg.
- Ciolli, J. (2025). America’s AI data center dominance, explained in one chart. *Business Insider*. Accessed: 2025-12-15.
- Corrado, C., Hulten, C., and Sichel, D. (2005). Measuring capital and technology: An expanded framework.
- Evercore ISI (2025). Closer Look At The Data Suggests AI / Tech Is A Big Part Of Growth Story In 2025 But Not The Whole Story. *Macroscope*. 7 October 2025.
- Goldman Sachs (2025a). US Daily: A Primer on the Impact of AI on the GDP Statistics. *Goldman Sachs Economics Research*. 13 September 2025.
- Goldman Sachs (2025b). Why AI Spending Is Not Boosting GDP. *Global Economics Comment*. March 2025.
- Goldman Sachs Research (2025). Is There Enough Data Center Capacity for AI? Artificial Intelligence insight article. Accessed: 2025-12-17.
- Highfill, T., Wasshausen, D., and Prunchak, G. (2025). Concepts and Challenges of Measuring Production of Artificial Intelligence in the U.S. Economy. BEA Papers 0134, Bureau of Economic Analysis.

Investors Business Daily (2025). AI Stocks Dispel Bubble Talk. Capital Spending Is Booming, With No End In Sight. 31 October 2025.

Jason Furman (2025). X post. 27 September 2025.

JP Morgan Asset Management (2025). Is AI already driving US growth? 12 September 2025.

Krugman, P. (2025). Talking With Paul Kedrosky. Substack. Accessed: 2025-12-XX.

Kwa, T., West, B., Becker, J., Deng, A., Garcia, K., Hasin, M., Jawhar, S., Kinniment, M., Rush, N., Von Arx, S., Bloom, R., Broadley, T., Du, H., Goodrich, B., Jurkovic, N., Miles, L. H., Nix, S., Lin, T., Parikh, N., Rein, D., Sato, L. J. K., Wijk, H., Ziegler, D. M., Barnes, E., and Chan, L. (2025). Measuring AI Ability to Complete Long Tasks. In *Proceedings of the 39th International Conference on Neural Information Processing Systems (NeurIPS 2025)*. Preprint available on arXiv. Accessed: 2025-12-15.

Laudani, P. and de Rozario, M. (2025). AI bubble? Opinions divided on tech’s trillion dollar question. *Reuters*. Reporting by Paolo Laudani and Mathias de Rozario in Gdansk; Editing by Matt Scuffham and Milla Nissi-Prussak. Accessed: 2025-12-15.

Leichter, R. (2025). Number of Data Centers by Country (November 2025). Last updated: November 04, 2025. Accessed: 2025-12-15.

Letzing, J. and Sung, M. (2025). What we mean when we talk about an artificial intelligence ‘bubble’. Economic Growth; Centre for the New Economy and Society. Published Oct 7, 2025; updated Oct 16, 2025. Accessed: 2025-12-15.

Microsoft AI Economy Institute (2025). AI Diffusion Report: Where AI is Most Used, Developed, and Built. Technical report, Microsoft AI Economy Institute. Published November 11, 2025; accessed 2025-12-15.

Neufeld, D. (2025). Technology Mapped: Data Center Capacity Around the World. Published July 30, 2025. Accessed: 2025-12-15.

OMDIA (2025). Global cloud infrastructure spending rose 21% in Q1 2025. OMDIA Newsroom. 12 June 2025.

Patrignani, L., Zangrandi, M. S., and Schiavone, A. (2025). Between sanctions and subsidies: reshaping the semiconductor ecosystem. *Questioni di Economia e Finanza (Occasional Papers)* 934, Bank of Italy, Economic Research and International Relations Area.

Savini Zangrandi, M. (2020). Semiconductor Wars. *mimeo*. October 2020.

Synergy (2025). Hyperscale Data Center Count Hits 1,136; Average Size Increases; US Accounts for 54% of Total Capacity. Press release, RENO, NV. Accessed: 2025-12-15.

The New York Times (2025). The Next Economic Bubble Is Here. 23 October 2025.

The New York Times Events (2025). Anthropic C.E.O.: Massive A.I. Spending Could Haunt Some Companies. YouTube video.

Weil, J. (2025). The Accounting Up roar Over How Fast an AI Chip Depreciates. *The Wall Street Journal*.

A The economics of modern AI data centers

A typical last-generation AI data center is built around NVIDIA’s NVL72 architecture: high-density pre-assembled server racks (i.e., small shelves packed with computer equipment) that contain 72 GPUs and 36 CPUs. Each of these racks occupies less than 1 square meter of floor space and costs about 3 USD million. The purchase price of these NVIDIA’s racks represents the majority of the total capital expenditure needed to build a modern AI data center. Table A1 provides an estimate of the break-down of the total capital expenditure needed to build an AI data center, per server-rack. Table A2 reports the rental prices of server racks, as quoted by some cloud infrastructure providers. Table A3 reports estimates of the yearly operating costs of a rack. Putting together these estimates, we can infer that, at current on-demand prices, the estimated payback period of an AI data-center investment is likely to be around 1 year, with a considerable margin for error due to several factors that are difficult to measure but should not radically change the estimate.²⁴

²⁴To be precise, the average yearly rent of an NVL72 rack at current on-demand prices largely exceeds the sum of the capital expenditure and the yearly operating cost. To infer the investment payback period precisely, the following information should also be known: 1) organizational overheads and the cost of developing and maintaining the cloud-management software used to rent the instances (the incidence of these costs decreases as scale increases); 2) the proportion of idle time (which should be small given reports of still binding capacity constraints); 3) the proportion of GPUs that are unavailable due to failures (although there are usually warranties in place, and NVL72 systems have sophisticated mechanisms to bypass failed components); 4) the proportion of capacity rented through long-term contracts and the average discounts applied to these contracts; 5) the additional revenue generated by network, storage and orchestration charges; these revenues are not included in the headline hourly costs reported here because they are highly user- and workload-specific and because they may be partly produced by separate hardware whose costs are not accounted for in Table A1.

Table A1: Data-center capital expenditure (per NVL72 server rack)

| | USD million per rack | Per cent |
|--------------------------------|----------------------|----------|
| NVIDIA GB200 NVL72 server rack | 3.00 | 67% |
| Land | 0.12 | 3% |
| Building | 0.24 | 5% |
| Electrical systems | 0.53 | 12% |
| Liquid cooling and other HVAC | 0.38 | 8% |
| Other | 0.24 | 5% |
| Total | 4.50 | 100% |

Source: Savills, Cushman & Wakefield, Dgtl Infra, Data Centre Dynamics, Anastasi In Tech, own calculations.

Note: estimated costs are obtained by averaging estimates from various sources and applying correction factors for liquid-cooled high-density equipment when estimates refer to traditional data centers. These figures refer to US-based data centers and do not include capital expenditures for dedicated power generation plants (i.e., power is assumed to be bought from third parties).

Table A2: Rental prices (USD) of NVL72 server racks

| Cloud provider | Price per hour | Price per day | Price per year |
|----------------|----------------|---------------|----------------|
| EC2 | 762 | 18,288 | 6,675,120 |
| CoreWeave | 756 | 18,144 | 6,622,560 |
| Oracle | 1,152 | 27,648 | 10,091,520 |

Source: web sites of cloud infrastructure providers.

Note: prices as quoted by cloud providers in October 2025.

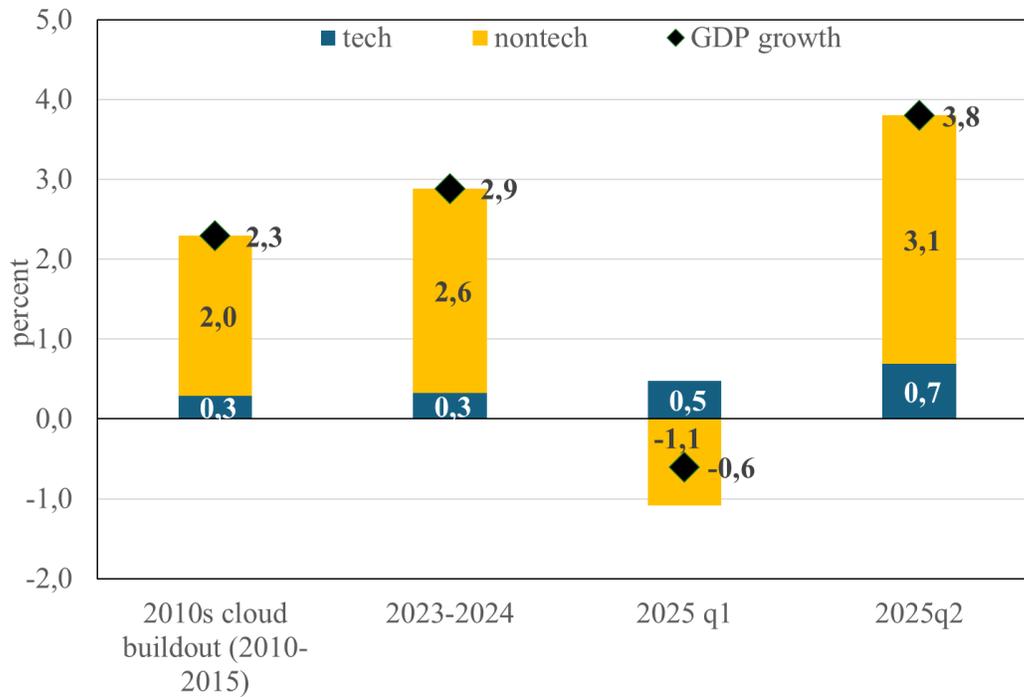
Table A3: Estimated yearly operating cost of an NVL72 server rack

| Cost component | Yearly cost (USD) |
|----------------|-------------------|
| Electricity | 105,000 |
| Maintenance | 25,000 |
| Network | 20,000 |
| Other | 40,000 |
| Total | 190,000 |

Source: authors' estimates based on multiple sources.

B Value added approach

Figure A1: CONTRIBUTION TO GROWTH, TECH VS NON-TECH -- VALUE ADDED APPROACH



Note: GDP growth rates are average quarterly annualized GDP growth over three different periods (2010q1–2015q4, 2023q1–2024q4, 2025q1–2025q2). Contributions to GDP growth of tech and non-tech components (blue and orange bars), in percentage points, are averages of the contribution to quarterly annualized GDP growth in the same periods.