



Chapter 5

The Revolution of Artificial Intelligence

For centuries, most of the world's economies grew at a similarly slow rate. However, the “Great Divergence” that occurred with the Industrial Revolution caused industrializing nations to accelerate their growth relative to the rest of the world (Pomeranz 2000). Today, artificial intelligence (AI) is a potentially transformative technology whose effects are often compared with those of the Industrial Revolution.

The world is witnessing clear leaders in AI investment, performance, and adoption metrics across different nations. The Trump Administration is laying the groundwork for American AI dominance by accelerating innovation, infrastructure development, and deregulation while establishing global supremacy through technology exports. If the AI revolution is as transformative as the Industrial Revolution, could this lead to the second Great Divergence? Of course, the future impact of AI is uncertain, so this chapter focuses on empirical data that can be measured today.

The chapter begins by reviewing analyses of the potential for AI-led economic growth, and it then discusses estimates of AI's impact on both gross domestic product (GDP) and the labor force. Recognizing that these effects are uncertain and thus need constant monitoring, the chapter then highlights metrics for tracking the breakneck pace of AI investment, performance, and adoption. It goes on to then discuss how different countries are faring on these metrics. The incredible speed of change cannot be overstated; many of these metrics are doubling every few months and increasing manyfold each year. This means that the AI of the future will likely be very different from today's AI. The chapter concludes by reviewing the actions President Trump is taking to ensure that

America continues to lead on AI. As the President has said, “America is the country that started the AI race. And as President of the United States, I’m here today to declare that America is going to win it” (Voice of America 2025).

The Future Outlook

The last 25 years have seen a great convergence as the world’s richest nations have grown more slowly than many developing nations. However, the advent of generative AI, based on large language models (LLMs), will initiate a new wave of profound economic transformation in the United States, promising significant boosts to productivity and growth. As AI technologies become more integrated into the workplace, economists are reevaluating long-term projections for GDP.

Yet this period of innovation is not without its complexities. This chapter focuses on the long-term analysis of structural trends, because not every AI-related investment will be profitable, and the short run always contains the potential for substantial volatility.

Background on Artificial Intelligence

The last few years have seen a rapid explosion in both AI capabilities and jargon, so it is useful to begin with a review of several key terms being used in the AI space. “Artificial intelligence” can refer to a wide variety of computer systems, from chess-playing computers like Deep Blue to generative AI apps like ChatGPT. For most of AI’s history, AI was only capable of making decisions among a relatively small set of options. The recent surge in AI interest has coincided with the rise of “generative” AI, which is so called because it is able to “generate” text, images, or video. “Large language models” are types of generative AI that can create text (Cloudflare n.d.). They are “large” because of their trillions of parameters, and are “language” because they are trained on large amounts of text written in natural languages (Cloudflare n.d.; Stryker n.d.). Agentic AI is a subset of generative AI that goes beyond mere content creation and can execute actions in order to accomplish goals (Google Cloud n.d.).

One framework for understanding the intelligence of an AI model considers this intelligence on two dimensions: (1) its ability to perform different tasks, from writing essays, to identifying objects in pictures, to writing computer code, to solving math problems; and (2) how AI’s capabilities on that task compare with human-level intelligence. Today’s AI systems have “specialized” (or “narrow”) intelligence because, though they may be superhuman at a particular task (e.g., no human can multiply as fast as a calculator can), AI is not able to perform all the tasks a human can. Humans are capable of performing a wide variety of

different tasks. Thus, we say that humans have “general” intelligence, while current AI (including both ChatGPT and Agentic AI) has “specialized” intelligence.

Artificial general intelligence (AGI) would be a hypothetical AI that could perform all the intellectual tasks that humans can (Google Cloud 2026), but the exact definition of AGI is hotly debated, and some definitions only require that AGI perform “many but not all” human tasks. Artificial superintelligence, sometimes just called “superintelligence,” is AI with an intelligence that surpasses that of humans (Google Cloud 2026). The boundary between AGI and superintelligence is similarly contentious, partly because these terms encompass different aspects of AI: “AGI” and “specialized AI” describe the generality of tasks an AI can perform, while “superintelligence” describes AI’s capabilities for these tasks. However, a “mere” AGI is already superintelligent if it can perform all human tasks, but at computer speeds. Accounting for semantic disagreements, it is worth noting that OpenAI, Anthropic, xAI, Meta, and Google all aim to create AGI or superintelligence (OpenAI 2023; Heath 2024; Musk 2025; Perkel 2025; Dragan et al. 2025).

This brings up an important caveat to this chapter’s analysis: the limitations of analyzing the economics of AI. As noted by Hanson (2001, 6), AI that could perform all human tasks would lead to absolutely explosive growth and to a very different world than that seen today. Thus, the implications of AGI (both economic and otherwise) are an important topic deserving further study, but are generally outside the scope of this chapter’s analysis, which focuses on “narrow” or “specialized” AI.

The Impact of Artificial Intelligence on GDP

Economists often think of the productive power of an economy as coming from three factors: the quantity of labor, the quantity of capital, and total factor productivity (TFP). TFP is a measure of an economy’s efficiency and technological progress. Rising TFP indicates that an economy is producing more goods and services from the same amount of labor and capital, or the same output with fewer inputs. This improvement in efficiency is a key driver of long-run economic growth and higher living standards (Zymek 2024). For rich countries like the United States, whose capital stocks are already very high, economic growth mainly comes from increasing TFP (Shackleton 2013; Prescott 1998; Wolla 2013).

The productivity gains from increasing TFP are eventually translated into higher overall economic output, or GDP. However, the effect of a new technology occurs with a time lag, as businesses must first successfully adopt the new technology and then adapt their operations (Tang, Wang, and Xu 2022). Many of the productivity gains of the 1990s emerged from technological investments that occurred in the 1970s and 1980s. Similar technological investments that occurred during the Great Depression bore fruit during the 1950s and 1960s

(Ferguson and Wascher 2004). As a result, while TFP is an important indicator, it is not a leading indicator of AI’s impact on the U.S. economy. Instead, research-and-development (R&D) spending on AI and the output of AI firms serve as early indicators of technological progress (Blanco, Gu, and Prieger 2016; Chou, Chuang, and Shao 2014). For example, AI-related R&D occurs well before the resulting innovations are widely adopted and then have a macroeconomic effect.

A variety of recent studies have attempted to quantify the effects of AI on GDP levels. These studies have produced a broad range of estimates (table 5-1): AI could increase GDP by 1 percent or up to more than 45 percent. This wide range reflects the high degree of uncertainty surrounding the economic characteristics of AI. However, it is worth noting that in the first half of 2025 alone, AI-related investment increased GDP by an annualized rate of 1.3 percent, harkening back to the scale of railroad investment during the Industrial Revolution (CEA calculations; Hausman, Pereria, and Pereria 2014) and seemingly ruling out the lowest few estimates.¹ Midrange estimates for the effects of AI on GDP include those from a variety of companies such as Oxford Economics (with an increase of 1.8 to 4 percent after 8 years), McKinsey (an increase of 2.4 to 4.1 percent in the long run) and Goldman Sachs (an increase of 7 percent after 10 years). High estimates include those by PricewaterhouseCoopers (an increase of 8 to 15 percent after 10 years) and by a Bank for International Settlements paper by Aldasoro and colleagues (2024; an increase of 20 to 45 percent after

Table 5-1. Estimates of AI’s Impact on GDP Level

Study	Impact on GDP level (%)	Time horizon	Region
Acemoglu (2024)	0.9 to 1.6	10 years	U.S.
Penn Wharton Budget Model (2025)	1.5	10 years	U.S.
Oxford Economics (2024)	1.8 to 4	8 years	U.S.
McKinsey (2023)	2.4 to 4.1	Long run	Global
Alonso et al. (2022)	4.7 to 19.5	Long run	U.S.
Goldman Sachs (2023)	7	10 years	Global
PricewaterhouseCoopers (2025)	1 to 15	10 years	Global
Aldasoro et al. (2024)	20 to 45	10 years	U.S.
Hanson (2001)	≥45	See note	Global

Sources: The sources are the works listed under “Study”; also see the reference list.
 Note: Except for Hanson (2001), the works listed here give the impact of AI on GDP levels, not GDP growth rates.

¹ Note that this 1.3 percent value for AI is the impact of AI investment on the level of GDP, even before any productivity gains from that investment are reaped. CEA staff could not locate this exact statistic for railroads during the Industrial Revolution, but U.S. investment in railroads grew from 0.2 percent of GDP in 1830 to 0.9 percent in 1839, and to a maximum of 2.6 percent of GDP in 1854 (Hausman, Pereria, and Pereria 2024).

10 years, for their approaches assume that all sectors of the economy will be at least somewhat affected by AI). Alonso and colleagues (2022) have a wide range of estimates (4.7 to 19.5 percent), reflecting uncertainty over whether AI will substitute more for skilled or unskilled labor (the latter of which would yield the divergence and therefore the high-end growth estimate for the United States). For comparison, a 2010 study by the Information Technology and Innovation Foundation indicated that the information technology (IT) revolution boosted U.S. GDP by about 14 percent (Andes et al. 2010; FRED 2025a).² These estimates all assume that AI can partially but not completely substitute for human labor: in the case where AI could do all human tasks, capital becomes a substitute for labor and annual economic growth increases to 45 percent (Hanson 2001, 6).

International Economic Growth Before AI

Even before AI, different countries may have been on different growth paths, with the United States exhibiting accelerating growth in potential GDP while growth in Europe and China was slowing. For Europe versus the United States, this was largely due to structural factors such as stronger U.S. productivity growth (especially in technology) and America's better business environment. For China, after decades of rapid growth, growth is now slowing to be more like that of emerging markets (Gourinchas 2025).

AI-led growth may be especially important for China, because its once-rapid growth has slowed in recent years to a level much more comparable to that of other emerging markets. Similar to China, AI-led growth may be especially important for Europe. Although the rise of China is one oft-repeated geopolitical story of the 21st century, another one that is less discussed, although perhaps no less important, is the decline of Europe. The European Union fell from 27 percent of world GDP in 1980 to just 14 percent in 2025 (IMF 2025). This was not just because of high growth rates in emerging markets but also because Germany and many other EU countries have a growth rate lower than that of other advanced economies. This trend continues with AI, where the EU lags behind the United States and China on various AI metrics. For example, cumulative private AI investment in the United States exceeded \$470 billion between 2013 and 2024, compared with roughly \$50 billion across all EU countries combined (Haag 2025).

Recognizing the critical role AI can play for future growth, the United States and many of its Allies have banded together through "Pax Silica," America's international partnership on AI supply chains (U.S. State Department 2025a). Pax Silica members range from major upstream semiconductor equipment manufacturers like Japan to downstream data center investors like Qatar (U.S. State Department 2025b). This varied group is united by a forward-looking

² The Information Technology and Innovation Foundation indicates \$2 trillion. Its report was published in early 2010. The 2009 U.S. GDP was \$14.5 trillion (FRED 2025a).

view of AI and technology. Thus, it is not surprising that Pax Silica members are growing more than twice as fast as their peers, with a 2.5 percent average real GDP growth rate between the release of ChatGPT in 2022:Q4 and the latest data, for 2025:Q3, versus 1.1 percent on average for Group of Seven countries (CEA calculations).

The Impact of AI on Labor, and Jevons' Paradox

Current evidence presents a mixed picture of AI's employment effects. Brynjolfsson, Chandar, and Chen (2025, 16) show that employment is falling for early-career workers in AI-exposed occupations like computer coding and customer service. Other studies have found no correlation between AI exposure and current unemployment rates (Gimbel et al. 2025). Still others have found that, while employment fell in sectors where AI can directly substitute for human labor, AI exposure actually increases employment in sectors reliant on AI-capable tasks (Johnston and Makridis 2025, 17). Notwithstanding the current impact of AI, overall unemployment is currently at a rate of just 4.4 percent as of December 2025 (U.S. Bureau of Labor Statistics 2025).

In the short run, if AI increases labor's efficiency, that reduces the amount of labor needed to create a given amount of output, potentially decreasing employment. But historical precedent suggests that efficiency gains can often increase (rather than decrease) total utilization of that resource—a phenomenon known as Jevons' Paradox (Alcott 2005). Jevons' Paradox occurs if a technological advance reduces the amount of a resource (like labor) needed for a specific application. This actually causes overall usage of that resource to increase, as usage expands to new applications. For Jevons' Paradox to occur and thus employment to increase with AI adoption, three conditions must be satisfied: first, AI must meaningfully boost worker productivity; second, the resulting cost savings must translate into lower prices; and, third, the lower prices must increase consumer demand faster than efficiency gains reduce per-unit labor needs (Rosalsky 2025).

Although these may seem like strong conditions, Jevons' Paradox has been observed occurring in many different fields. Jevons first described the paradox in 1865, when the increasing efficiency of coal in iron engines actually increased the demand for coal, iron, and other resources (Alcott 2005). In agriculture, increases in irrigation efficiency may increase water consumption (Tang, Wang, and Xu 2020). Improvements in energy-efficient lighting have increased both the number of light bulbs demanded and the amount of electricity used in lighting (Dibal 2025). Jevons' Paradox even occurs in topics seemingly unrelated to production: increases in road capacity will increase the numbers of drivers on the road (Duranton and Turner 2011; Hymel 2019; Hsu and Zhang 2014). And specifically for AI and jobs, a similar situation may be occurring for radiologists, a job once predicted to be replaced by AI (Agten and Ruthven 2025), but which is now seeing historically high employment rates (Mousa 2025).

In the longer term, the key issue is comparing and contrasting AI to prior disruptive technologies. Historical analogies suggest that disruptive technologies (e.g., steam power, electricity, computers, and the Internet) ultimately lead to greater employment and earnings (Hötte, Somers, and Thedorakopoulos 2023). AI could be the exception if the technology either develops agency (so that it can work as independently as humans do) or if it dramatically increases worker productivity without generating new labor demand (Ayres 1990; Cooper, Gumbel, and Lund 2018; Donaldson 2018; Feigenbaum and Gross 2024). But the general precedent of past technological changes is that they create a variety of new fields. In 1860, 43 percent of U.S. employment was in agriculture, compared with 1.2 percent in 2015 (Elvery 2019). During these years, a huge range of new professions have been created, many reliant on the new technology. Now the majority of current workers are in jobs created since 1940, ranging from “wind turbine technician” to “software developer” to “textile chemist” to “mental-health counselor” (Autor et al. 2024).

Key Metrics to Track

One of the most important points to focus on, in order to understand AI, is the speed of progress and change. The amount of computer power spent to train an AI model has not merely doubled every year; since 2010, it has increased at an average of about fourfold per year (Roldán and Sevilla 2024). Similarly, top AI companies are seeing revenues triple each year and are projecting future growth more rapid than that seen in the rapid-growth phases of Google, Amazon, or Microsoft (see figures 5-5 and 5-7 below). Because many of these metrics are doubling every few months and increasing manyfold each year, this means that change resulting from AI can be very rapid.

Similar to how market analysts monitor housing starts or manufacturing output to predict broader economic health, a specific set of indicators can reveal AI’s growing influence in the U.S. economy. AI’s impact on GDP materializes in changes in total factor productivity, so this measure is highlighted first. But, because TFP is a lagging indicator, other metrics that serve as leading indicators of AI’s impact are also considered. These metrics, which track the expanding scale of AI-related investment, the accelerating capabilities of AI, and the rising adoption of AI, collectively serve as economic barometers for the AI revolution. These metrics are interrelated: investment is a direct signal that companies are pouring resources into advancing the technology, which results in increased model performance and lower unit costs. Better capabilities at lower costs spur AI usage, which is reflected in the revenues of AI firms. This section begins by discussing these metrics in the United States and then turns to a cross-country analysis.

Total Factor Productivity

AI is so important for growth because of its potential impact on total factor productivity. But previous technological revolutions have had a complicated relationship with productivity. Computers were once only mainframes that took up entire rooms, but now they can fit in your pocket. In 1987, the economist Robert Solow (1987) famously quipped “You can see the computer age everywhere but in the productivity statistics.” Explanations for the apparent lack of impact of computers ranged from time lag, to the inability of economists to measure their real benefits, to claims that computers do not actually increase productivity (Triplett 1999). As a result, while the impact of AI on TFP may be the key question, other metrics must be relied on as well.

Investment

Investment in the AI ecosystem has been massive, both for the models themselves and for the surrounding infrastructure.


AI models. AI models exhibit a predictable tendency for model performance to increase as developers increase the number of parameters in the model, the size of the training data set, and the amount of computer power used to train the model. Known as “scaling laws,” these empirical relationships have enabled model developers to increase the performance of their AI models not solely by relying on fundamental scientific breakthroughs, but simply by throwing more resources at the problem. Similar empirical relationships are seen in other fields, such as Moore’s Law, in which the number of transistors on an integrated circuit doubles every two years. Because scaling laws are not laws of nature but are observed empirical relationships, they could someday end. But they have characterized the current deep-learning era, where the amount of compute spent on training computer models has increased by more than 1 billion-fold since 2012 (figure 5-1) (Epoch AI 2025).

A staggering amount of investment has been required to meet these needs. Global corporate AI investment reached \$252 billion in 2024. Generative AI alone is up 19 percent year over year, reaching \$34 billion (Stanford Institute for Human-Centered Artificial Intelligence 2025). This investment has been concentrated in the United States, which had \$94 billion in private sector AI investment in 2024 (figure 5-2) (Our World in Data 2025b).

From 2016 to 2024, the energy and amortized hardware costs to train (build) an AI model have grown at an average rate of 2.4 times a year, while cloud compute costs (figure 5-3) have grown at an average rate of 2.5 times a year (Cottier et al. 2024). As a result of almost a decade of annually doubling costs, Grok 4, an AI model published in July 2025, cost about \$490 million to train (Edelman, Emberson, and Saunders 2025).


Sustained investment in model training despite increasing costs indicates a commitment to developing more capable and complex AI systems. Many

Figure 5-1. The Amount of Computing Power Used to Train AI Models



Source: Epoch AI (2025).
Note: Training compute (FLOP) refers to the number of floating point operations used to train an AI model.

Figure 5-2. External Funding for Privately Held AI Companies




Sources: Quid (2025); U.S. Bureau of Labor Statistics (2025); Our World in Data (2025a).
Note: This figure only includes private companies whose private funding is at least \$1.5 million.

investment indicators are available globally, so for additional discussion of investment, see the cross-country comparisons below .

AI infrastructure. Beyond just investment in model training, investment in data centers and related equipment surged in 2025 due to the increasing

Figure 5-3. Cloud Compute Cost to Train AI Models



Source: Epoch AI (2025).
Note: Estimated cloud compute costs are the product of a historical cloud rental cost and the number of training chip-hours. These costs are for the final training run of models.

proliferation of AI technology. Investment in information processing equipment and software in the United States increased at an annual rate of 28 percent during the first half of 2025, up from 5.5 percent annual growth in 2024 (FRED 2025d). Alternatively put, in 2025:Q2, this investment was already more than \$125 billion higher (in annual terms) than it was at the end of 2024. Information processing equipment and software comprise one quarter of all U.S. investment (FRED 2025b, 2025c). The fact that an already-large category is growing at a very rapid rate means that AI is now driving an investment-based surge in U.S. GDP, as opposed to a surge driven by consumption or unsustainable government spending.

Performance

Continuous investment in AI has increased the performance of AI models in terms of their capabilities to solve different tasks, the length of tasks they can successfully perform, and reducing the cost per “token” produced by an AI model.³ Two measures of performance are considered: benchmark scores and cost per token.

Improving benchmark scores. Benchmarks are sets of standardized tasks designed to evaluate specific AI capabilities, such as reasoning, coding, and language understanding. As LLMs become more powerful, they achieve


³ AI models break the text you give them down into individual tokens, so tokens are a measure of the amount of data processed by AI models.

near-perfect scores on older benchmarks, a phenomenon known as “benchmark saturation” (Bansal et al. 2021). For example, from 2023 to 2024, AI performance on the computer coding benchmark SWE-bench jumped from 4 percent to 72 percent. Similar phenomena have occurred in benchmarks for graduate-level question answering, advanced math, and a variety of other academic subjects (Stanford Institute for Human-Centered Artificial Intelligence 2025).

However, while frontier AI are vastly better than humans at many exams and tasks, the best current AI agents often struggle with stringing together longer sequences of actions. As a result, they are currently unable to carry out substantive projects by themselves and are unable to fully substitute even for low-skill computer-based work like a remote executive assistant (METR 2025). But this means that the length of tasks that models can complete is a helpful lens for understanding AI capabilities. The length of tasks that AI are able to successfully complete is also increasing, doubling every 7 months for the past 6 years (figure 5-4). This means that AI is becoming better able to manage larger and larger projects on its own and thus complete increasingly more complicated tasks.

Falling cost per token. A “token” is the basic unit of input to an LLM, for example, a single word or number. A decrease in the cost per token makes AI more affordable. This can occur because of smaller, more efficient models (software) or better hardware. Depending on the model, prices are falling at least 9 times a year, and up to 900 times a year (Cottier et al. 2025).

Figure 5-4. Lengths of Software Engineering Tasks That AI Can Complete with a 50 Percent Success Rate




Source: METR (2025).

Adoption and Usage

As a result of improved AI capabilities and falling costs, the use of AI has spread throughout the economy. This can be tracked through revenues of frontier AI companies, the use of AI in the production of goods and services, and an increasing share of Americans using AI in the workplace.

Revenue. AI companies have seen rapid, but not unprecedented, growth, but their future growth could exceed all historical examples. Startups can often see explosive growth, and OpenAI, Anthropic, and Google DeepMind each had over 3 times annualized revenue growth as of the second half of 2024 (figure 5-5) (Emberson, Owen, and Snodin 2025).

Figure 5-5. Revenue from Sales of AI Products to the Public, by Company




Sources: Epoch AI (2025).

Note: Google DeepMind revenue data are estimates based on mobile app usage and web traffic.

This is much faster than the market average: Standard & Poor’s 500 companies had a blended year-over-year earnings growth rate of 10.3 percent in the last quarter of 2024 (Consello 2025). But AI company growth so far is comparable to that of top tech unicorns like Google and Uber during their initial high-growth phases (figure 5-6) (Burnham 2025). So, while this growth is impressive (Amazon only had two years of this level of revenue growth) (Companies Market Cap n.d.), it is not unprecedented.

However, what may be unprecedented is the future growth of AI companies. For example, despite skepticism, OpenAI claims that it will roughly double its revenue each year from 2026 to 2028 (Muppidi 2025; Burnham 2025). In order to try and understand this claim, it is helpful to compare this growth with the historical growth of previous big tech unicorns; such revenue growth from

Figure 5-6. Actual Revenue of OpenAI versus Other Companies’ Historically Rapid Revenue Increases



Source: Epoch AI (2025).

OpenAI would be far higher than the growth rate seen by these previous big tech unicorns (figure 5-7) (Burnham 2025).

Business usage. Use of AI by organizations jumped from 55 percent in 2023 to 78 percent in 2024 (Stanford Institute for Human-Centered Artificial Intelligence 2025). In particular, the use of AI in the production of goods and services increased from less than 4 percent of firms in 2023 to about 10 percent of firms in September 2025 (figure 5-8) (U.S. Census 2025).

The share of businesses with paid subscriptions to AI has grown at an even faster rate, from 7 percent of companies in January 2023 to 45 percent today (Kharazian 2026). A similar story occurs with U.S. workers, about 40 percent of whom are now using generative AI at work (figure 5-9) (Hartley et al. 2025).


Critical minerals. Many minerals are critical parts of the AI supply chain. Silicon is the main building block for most semiconductor chips, while gallium and germanium are two other key components. (Hong 2025) The International Energy Agency (2025) estimates that by 2030, data centers alone could demand over 10 percent of today’s world supply of gallium.

Cross-Country Comparisons

There are many ways to rank countries on AI, with many groups developing their own indices (Ferres et al. 2025, 1-18; Stanford Institute for Human-Centered Artificial Intelligence 2025; IMF 2023) Using the lens of investment, performance, and adoption of AI, this section compares countries to determine which


is leading in AI. In general, the United States ranks first in most metrics, with China second, and the EU third.

Figure 5-7. Projected Revenue of OpenAI versus Other Companies' Historically Rapid Revenue Increases




Source: Epoch AI (2025).

Figure 5-8. Percentage of U.S. Firms Using AI to Produce Goods and Services



Sources: U.S. Census Business Trends and Outlook Survey (2025, biweekly data); Haver data (2025).

Figure 5-9. Percentage of U.S. Workers Who Use Generative AI at Work




Investment

Tracking overall AI investment is difficult, because it is split across a variety of companies throughout the AI supply chain, from chips to data centers to AI labs, and over a variety of public and private sources. Looking at total R&D spending, in 2022, Israel's R&D spending equaled 6.0 percent of its GDP, a higher share than any other country in the world. Israel was immediately followed by South Korea (5.2 percent), Taiwan (4.0 percent), the United States (3.6 percent), and Japan (3.4 percent). By comparison, China spent 2.6 percent, and the EU spent 2.1 percent (National Science Foundation 2025).

For private investment, America's private companies are leading in AI R&D. The United States had \$109 billion in private AI investment in 2024, compared with \$9 billion for second-place China, with the United Kingdom, Sweden, and Canada rounding out the top five (figure 5-10) (Stanford Institute for Human-Centered Artificial Intelligence 2025). Unsurprisingly, then, the United States has about 75 percent of reported venture funding in generative AI startups (Haag 2025).


However, private investment in private AI companies is not the only type of investment. Other countries have sought to catch up, with a variety of special AI investments by governments or sovereign wealth funds (figure 5-11) (Benaich 2025). In addition to expected players like the EU and China, several Middle Eastern countries are investing heavily in AI. China has a large amount of public sector AI spending, with an estimated \$56 billion in 2025 (Kaur 2025). Saudi Arabia's Public Investment Fund established a new AI company, Humain, with a \$10 billion venture fund (Public Investment Fund of Saudi Arabia 2025;

Figure 5-10. Private Investment in AI, by Country, 2024



Sources: Quid (2024); Stanford HAI Artificial Intelligence Index Report (2025d).

Figure 5-11. Number of Notable Announcements of Direct AI Spending by Nations and Sovereign Wealth Funds



Source: Benaich (2025).
Note: This graph focuses on spending from direct government investment and sovereign wealth funds and thus does not include many other traditional sources of AI funding, such as from government funding of general R&D. These totals include target spending amounts for future years. This compilation of AI spending announcements is not exhaustive.


Turak 2025). Likewise, the United Arab Emirates is partnering with OpenAI, NVIDIA, and other U.S. companies to build a variety of data centers as part of the Stargate project (Open AI 2025).

The nature of AI investment could look very different outside the United States, especially in developing countries. For example, when these countries created national telephone networks, many leapfrogged landlines and went straight to mobile phones (Ritchie 2024). A similar phenomenon could occur in developing countries and their use of AI, whereby datacenters might face electrical reliability concerns, and the main AI interaction platform could be smartphones.

Performance

Given the large investments in AI in the United States, it is unsurprising that America leads, with 154 AI systems of size about equal to GPT-3, about half the world total of 331 in 2024 (figure 5-12) (Our World in Data 2025a). However, due to the rapid speed of AI advancement, the performance gap between the best models of each country is relatively small. According to a report by Microsoft, “Only seven countries—the [United States], China, France, South Korea, the [United Kingdom], Canada, and Israel—rank among the top 200 models, and the distance between the frontier (United States) and the last of these (Israel) is now just 11 months” (Cveko 2025).

Figure 5-12. Cumulative Number of Large-Scale AI Systems by Country




Sources: Epoch AI (2025); Our World in Data (2025b).
Note: This graph covers AI systems with training compute exceeding 10^{23} FLOPs of compute. The “other countries” category includes Australia, Canada, Finland, Germany, Hong Kong, Israel, Japan, Russia, Saudi Arabia, Singapore, South Korea, Switzerland, and the United Arab Emirates.

Adoption and Usage

As a result of heavy investment, as of May 2025, the United States has about 74 percent of the world’s compute capacity for AI (figure 5-13) (Heim et al. 2025), and much of the foreign AI hardware was originally made by U.S. companies. For example, almost all Chinese AI models are trained on U.S. hardware (figure 5-14) (Adamson et al. 2025).


Figure 5-13. Percentage of Graphics Processing Unit Clusters (Weighted by Cluster Performance) by Country or Region



Source: Epoch AI (2025).
Note: The United States introduced export controls on key equipment and AI chips to China in October 2022, and strengthened these controls in October 2023 and December 2024 (Heim et al. 2025). As of March 2025, this data set represents about 10 to 20 percent of global graphics processing unit cluster performance.

Usage is more broadly distributed, with Israel and Singapore showing the highest per capita usage of the Claude AI model (Appel et al. 2025, 16). OpenAI shows similar trends, with the United States representing only 19 percent of ChatGPT traffic (Similarweb 2025). In general, middle-income economies have disproportionately high usage of generative AI relative to the size of their economy, and, in aggregate, accounted for 50 percent of global AI usage in 2024, while low-income economies accounted for less than 1 percent (Liu and Wang 2025, 27), partially due to a lack of electricity (Demombynes et al. 2025, 21). In general, AI adoption is highly correlated with GDP, with significantly higher uptake in industrial countries relative to developing countries. Key factors driving this divergence are differential access to “building blocks” of AI, including electricity, data centers, Internet access, language, and digital skills (Microsoft 2025).

Figure 5-14. Chinese AI Models by Country of Hardware Origin



Sources: Adamson et al (2025); Epoch AI (2025).

Note: Chinese AI models refer to models whose developers included a Chinese organization.

The Trump Revolution

The Trump Administration is pursuing a number of policies to improve America's position on each of the investment, performance, and adoption indicators, with many policies improving multiples of these at the same time. Deregulation reduces the cost of constructing data center infrastructure, incentivizing AI investment. The One Big Beautiful Bill Act (OBBBA) makes it easier for Americans to invest, and trade agreements are attracting foreign investment. All this investment leads to the AI performance necessary for American AI to dominate, while U.S. energy dominance provides the electricity to meet the increasing demand from AI.

Investment

The OBBBA (Public Law 119-21, signed July 4, 2025) restored and expanded full, immediate expensing for qualified investment and extended pro-investment business provisions, shifting after-tax hurdle rates in favor of building now (Ceballos et al. 2025). The Federal corporate income tax rate remains at 21 percent (Azebu and Varma 2025). The law preserves and refines a range of internationally focused provisions (Ceballos et al. 2025), and it invested more than \$1 billion in critical minerals across a variety of AI initiatives (BPM 2025). The CEA's analysis predicts that the OBBBA will increase annual GDP growth by more than 1 percent over the four years after its passage, and raise real wages by \$4,000 to \$7,200 per worker (White House 2025e). Of particular relevance to the AI boom, the OBBBA brings back 100 percent bonus depreciation for

IT infrastructure and data center equipment (BPM 2025). The CEA estimates that, altogether, the OBBBA will lead to a 7–10 percent increase in investment, incentivizing the building of data centers, power infrastructure, and chip manufacturing (White House 2025e).

As part of trade deals and other agreements, President Trump has secured trillions of dollars in investment commitments from foreign countries. The European Union, in its trade agreement, specifically committed to buy \$40 billion of U.S. AI chips, and the United Arab Emirates specifically mentioned AI as one of the key sectors that its \$1.4 trillion in U.S. investment will be directed toward.

Performance

The Trump Administration’s AI policy is outlined in its 2025 “AI Action Plan” and several related Executive Orders. The policy focuses on attaining American international dominance in AI by rapidly building data centers, enabling and accelerating innovation, and upholding free speech in AI models (White House 2025c).

To reach these goals, the action plan outlines concrete steps. For example, regarding the rapid buildout of data centers, the plan recommends establishing new categorical exclusions under the National Environmental Policy Act for data centers in order to fast-track the permitting process. To enable and accelerate innovation, the plan recommends direct investment in AI technology by various Federal agencies, along with establishing AI Centers of Excellence, where researchers and startups are incentivized to rapidly deploy and test AI tools in contexts committed to the open sharing of data and results. Regarding the upholding of free speech, the plan recommends that Federal procurement guidelines be updated to mandate that the government will only contract with AI developers that ensure their systems are objective and free from top-down ideological bias (White House 2025c).

Excessive regulation harms economic activity by increasing costs (Crain and Crain 2014, 26), stifling competition (Ashton 2023) and innovation (Aghion et al. 2023, 2895) and raising consumer prices (Chambers, Collins, and Krause 2019, 59). This can reduce growth (Coffey et al. 2020, 15), startup activity, and job formation (Bailey and Thomas 2018, 3), and increase the poverty rate (Chambers, McLaughlin, and Stanley 2018, 143), with a disproportionate impact on small businesses (Crain and Crain 2014, 54).

On July 23, 2025, President Trump signed an Executive Order to accelerate permitting for data centers as well as their underlying energy and manufacturing infrastructure (Federal Register 2025a). Then, on December 11, President Trump signed an Executive Order to reduce barriers at the state level as well (Federal Register 2025b).

As previously studied by the CEA, deregulatory efforts by the Trump Administration aim to ameliorate these issues and deliver benefits to the AI sector along with other sectors throughout the economy. The CEA estimates that these deregulatory efforts are capable of delivering meaningful productivity gains that translate into an extra 0.3 to 0.8 percentage point of GDP growth each year for two decades—and a cumulative increase of about 6 to 17 percent by 2045 (White House 2025d).

Adoption and Usage

President Trump has made energy dominance one of the priorities of his Administration (White House 2025). The Administration has already taken important actions to stimulate domestic production and reduce costs by resuming Federal leasing for energy development, issuing new permits for liquefied natural gas export terminals, and supporting advanced nuclear development, among others. Not counting deregulation, the CEA estimates that policies that support American energy dominance could raise U.S. GDP by at least 0.3 to 1.2 percent by 2035, without accounting for synergies with AI (White House 2025a).

Even more so than traditional data centers, generative AI data centers are especially electricity-hungry, and AI data centers are forecasted to grow from 4 percent of U.S. electricity demand in 2023 to 7 to 12 percent by 2028 (Luna and Pomerleau 2025). To accommodate this increase in demand, the “AI Action Plan” includes concrete policy actions related to building up the energy grid. In particular, it recommends exploring grid management technologies and upgrades to power lines in order to optimize and stabilize the existing grid, in tandem with prioritizing the construction and connection of new power plants using a variety of energy sources to the grid. By April 2025, the U.S. Department of Energy (2025) had already identified 16 potential sites with the energy infrastructure that positions them for rapid data center construction.

Conclusion

The AI revolution, with its parallels to the Industrial Revolution, presents a profound economic inflection point with the potential to significantly increase the GDP of countries that embrace it. The world is witnessing clear leaders in AI investment, performance, and adoption metrics across different nations. The United States, as demonstrated by its comprehensive “AI Action Plan” and related Executive Orders from the Trump Administration, is pursuing a strategy focused on accelerated innovation, infrastructure development, and establishing global dominance through technology exports and deregulation to lay the groundwork for American AI dominance.

