August 2026

This marketing communication is for professional investors only.

White Paper

AI and investing: navigating the tech evolution

AI is reshaping markets, productivity, infrastructure demand, and sector leadership. What does this mean for long-term investors? This article breaks it down.

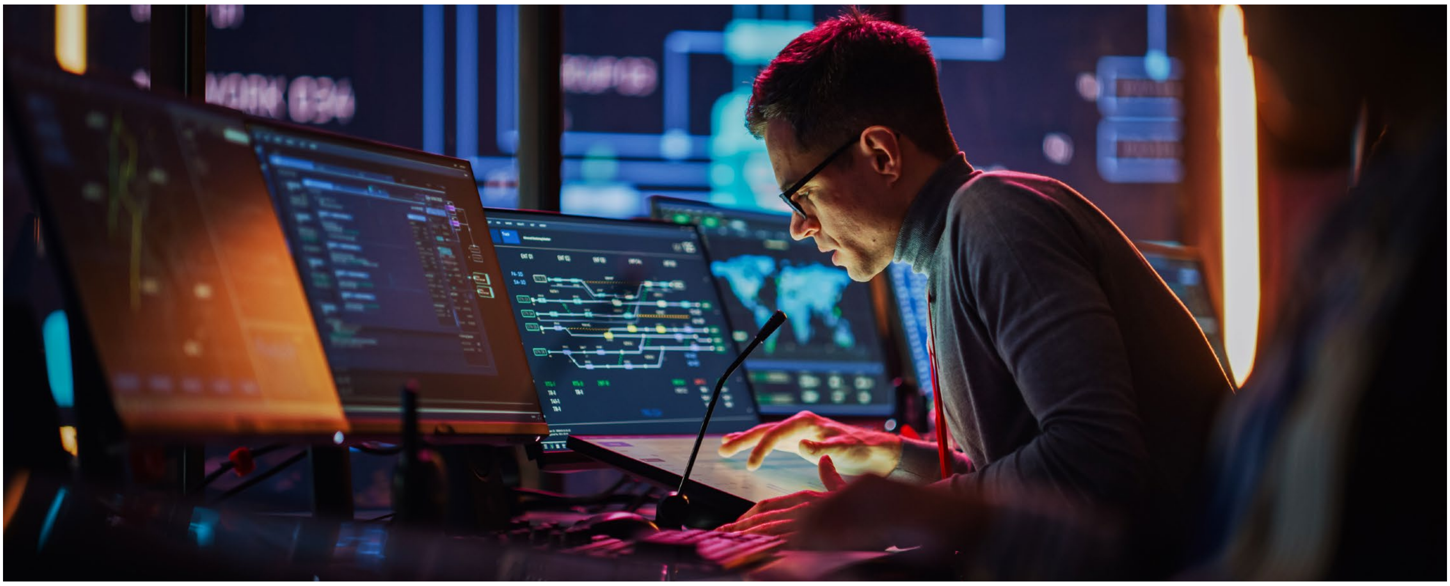

Artificial intelligence has moved from technological curiosity to economic force with incredible speed. It has transformed leadership in financial markets, how capital is deployed, how businesses operate, and how investors assess the long-term growth potential of a variety of companies. For long-horizon investors, we believe the question is no longer whether AI matters, but how its evolution will reshape the investment landscape over time. Here, we explore the wave of AI excitement among investors and the economy, and how markets are pricing this story.

How AI became a powerful technology

It may seem that AI came out of nowhere, but underlying today's rapid progress is decades of research and many dead-ends. AI was first defined as a formal field of research in the 1950s, and its slow progress was enabled by limited computing technologies of the time. Advancements in semiconductor technology in the 1960s and 70s, however, helped drive improvements in AI capabilities alongside greater computing resources.

The design of systems over time reflected the resources available at the time. In the 1980s, "knowledge-based AI" sought to pre-specify all human knowledge (see the Cyc project[1], for instance) into systems that could use steps of logic to process information - a monumental task that struggled with real-world results. In the 1990s, the increasing power of computing equipment enabled the rise of machine learning - systems that can acquire and apply knowledge to a wide range of problems. They grew in popularity and became the mainstream in artificial intelligence research, leading to widespread commercial use in the 2010s through applications like product recommendation, spam filters, and voice assistants. Indeed, machine learning was made commercially successful by enormous scale in computing, data, and better-defined models.

[1] The Cyc project is a long-running artificial intelligence effort that started in 1984 with the goal of building a comprehensive knowledge base of everyday facts and rules about the world, enabling machines to draw logical inferences and reason in a more human-like way.

A particular form of machine learning, called "deep learning," is behind today's AI models. Deep learning is a subset of machine learning inspired by how the human brain works. Though the concept was pioneered as early as the 1940s, deep learning AI models require such significant computing power that computers of the day struggled to deliver meaningful results from them. In the 2010s, scale in cloud computing helped solve these large computing problems, leading to deep learning-powered AI models that could read handwritten text, recognize objects in pictures, and more. Today's generative AI systems, including ChatGPT, Claude, Gemini, and other name-brand models are built on this same deep learning technology. But their success depends on enormous capital demands for their development, deployment, and operation.

$730 bn
projected spend on AI projects

The enablers of AI

In the few years since generative AI became a mainstream topic, adoption has proceeded rapidly with more than two billion monthly users combined across major AI platforms.[2] Capabilities of these AI systems have also progressed rapidly, with milestones now measured in weeks rather than years and enabled in large part by a surge in AI-related infrastructure spending. For instance, US hyperscalers expect to spend $730 bn this year, largely on AI projects.[3] Gartner, a tech advisory company, expects 2026 AI-related spending to exceed $2.5 trillion globally. Today, most companies have articulated some kind of AI ambition, in our analysis.

Most of this spending comes back to one thing: Computing demand. The training of today's highly capable AI models requires significant computing power, and these needs have continued to grow rapidly. According to Epoch AI, leading edge models have seen their compute needs roughly double every six months over the past decade. After a model is trained, it requires ongoing computing costs (referred to as inference) every time the system is used. Here, too, compute demand continues to grow, particularly with the rise of agentic AI and growing adoption.

As computing demand grows, hyperscalers, neoclouds[4], and other data center builders are looking to secure and deliver as much compute as possible. The data center build-out is about meeting this demand, and it comes in four core verticals:

Commodities	Semiconductors	Data center infrastructure	Energy

Building and powering data centers requires significant raw resources. While there are various commodities that go into the manufacture of semiconductors, computing equipment, and the related infrastructure, we highlight two key areas of focus: Copper and rare earths. Data centers require enormous amounts of power delivery and cooling equipment. Copper is the commodity of choice for its conductive properties in both cases. Indeed, one gigawatt scale data center uses as much copper as 500,000 electric vehicles, and worldwide data center capacity is expected to grow from around 150 GW as of the end of 2025 to around 330 GW by 2030.[5] Similarly, the manufacture of semiconductors requires rare earths with limited supplies (and often concentrated supply chains).

2 Based on user data for ChatGPT, Gemini, Qwen, Doubao, DeepSeek, Perplexity, Grok, and Claude. This selection focuses on standalone chatbot apps ("AI as a service") and excludes embedded offerings ("AI as a layer") such as Microsoft's Copilot and Baidu's ERNIE Bot. Sources: Invesco Strategy & Insights and company and media reports as of 30 April 2026. User counts are mixed between weekly and monthly active users. Some numbers are unconfirmed, including DeepSeek, Grok, and Claude.

3 Based on Bloomberg consensus forecasts for headline capital expenditures from AI Hyperscalers, a basket that includes AMZN, GOOGL, META, MSFT, and ORCL. We classify these companies as "AI Hyperscalers" in that they have significant, global cloud computing platforms and resources as well as preexisting access to AI expertise. Latest available data as at 31 May 2026. A report from CreditSights, Technology: Hyperscaler Capex 2026 Estimates, 10 November 2025 indicates about three-quarters of the capital spending is expected to affect AI-related infrastructure.

4 Neoclouds are specialized cloud computing providers that run data centers optimized for AI workloads, supplying high-performance computing power rather than general-purpose cloud services.

5 Note: AI data centers are anticipated to consume around 32 to 47 metric tons of copper per megawatt (MW) installed. Source: S&P Global, Copper in the Age of AI: Challenges of Electrification, published 8 January 2026. Copper intensity per electric vehicle is assumed at 70 kilograms, based on Goldman Sachs figures as of December 2025. There is no guarantee that forecasts will come to pass.

The second key vertical is semiconductors. Data centers built for AI workloads require purpose-built chips, which are designed by specialized companies and then manufactured at high-tech facilities. Chip designers have been one of the key beneficiaries of the AI investment boom, with companies such as Nvidia and AMD riding the investment wave, and hyperscalers themselves have also turned to designing their own chips. Since the manufacture of these chips is highly capital intensive and highly specialized equipment is required, almost all these chips are manufactured in a small number of facilities. Orders for these chips have been a key component of AI-related spending, and bottlenecks have appeared along the way, such as in memory and packaging.

Chips are fitted into data centers but require a host of infrastructure to be useful, which comprises the third vertical of interest. This includes networking, cooling, and grid interconnection equipment. As data center construction has ramped up, many of these inputs have come under multi-year delivery delays due to order backlogs.[6] They have also created winners in certain hardware and capital goods manufacturers.

15%
forecast annual growth rate of electricity consumption through 2030

Most salient, in our view, has been the fourth vertical: energy. The International Energy Agency predicts that data center electricity consumption will grow by 15% per year globally through 2030, four times faster than consumption from all other sectors and rising to 3.0% of all electricity use.[7] In developed markets, data centers are expected to contribute 20% of the total growth in electricity demand over this same period. Reflecting this, local grid capacity served by utilities have come under strain in recent years. Last year, ERCOT (Texas's grid operator) reported that it received requests for more than 226 GW of power by 2030, **nearly 100 times more** than it approved in 2022.[8] Data center builders have increasingly turned to a "bring-your-own-generation" (BYOG) model where part of the data center build includes solving for the data center's energy needs, such as through natural gas-fired turbines or nuclear energy. Both are subject to very long lead times (5-7 years), and we expect this area will continue to yield attractive investment opportunities as data centers are built.[9]

From infrastructure build out to fundamentals

Exposure to AI in financial markets has so far been captured through a narrow set of infrastructure providers—semiconductors and data center related businesses as detailed above. This trade has been propelled by two forces: a rapid build-out of computing capacity to facilitate greater AI capabilities and use cases, and growing adoption from consumers and businesses.

We expect this to continue, but with greater nuance than in recent years. Indeed, in recent months, markets have increasingly differentiated between companies based on balance sheet strength, capital discipline, and paths to monetization. Where AI names previously moved largely in tandem in a sort of "AI beta trade," more recently we believe markets are expressing greater focus on fundamentals—particularly capital structure—as much as technological leadership.[10]

At the same time, AI infrastructure faces physical constraints. Our analysis shows that power availability, grid capacity, and local opposition are emerging as real bottlenecks. These frictions do not end the AI build out, in our view, but they are likely to change its economics—and with them, the profile of long-term returns.

6 Power infrastructure equipment has been one of the primary bottlenecks. For example, GE Vernova and Siemens Energy report time to delivery for new orders of combined-cycle gas turbines of 5 to 7 years, as of 2025 figures.

7 Source: International Energy Agency, April 2025.

8 Source: The Economist, 15 January 2026.

9 Note: AI data centers are anticipated to consume around 32 to 47 metric tons of copper per megawatt (MW) installed. Source: S&P Global, Copper in the Age of AI: Challenges of Electrification, published 8 January 2026. Copper intensity per electric vehicle is assumed at 70 kilograms, based on Goldman Sachs figures as of December 2025. There is no guarantee that forecasts will come to pass.

10 Sources: Bloomberg L.P. and Invesco Strategy & Insights, based on the rolling 20-day pairwise correlations of Alphabet, Amazon, Meta, Microsoft, and Oracle. As of 5 June 2026.

Productivity and growth

The long term macroeconomic implications of AI rests on productivity, and indeed we think effective implementations of AI are likely to result in significant value capture going forward. Generative AI can assist, automate, and augment work, particularly in information intensive tasks. Evidence from early use cases—such as software development—suggests meaningful productivity gains are possible, although estimates still vary widely. Over time, we expect AI will replace certain categories of tasks and transform many aspects of work. However, we think a substantial role nevertheless remains for labor.

If AI adoption follows prior technology cycles, productivity gains are likely to be incremental and uneven at first, then broader and more impactful over time. For developed economies facing aging populations and possible labor shortages, AI's role as a labor augmenting technology is particularly relevant. Over the long run, higher productivity should be a tailwind to average growth rates and exert disinflationary pressure, in our view. For investors, this likely points to AI as an ongoing, multi-year macro force rather than solely a near-term earnings catalyst.

Disruption and dispersion across sectors

As AI capabilities improve and costs fall, attention is expanding from who builds the technology to who might be disrupted by AI. Software, professional services, financial intermediation, and other information heavy industries are increasingly seen to be exposed as AI tools threaten to transform or even substitute some business models. Price action in the first quarter of 2026 is an example of some of these fears playing out, where software companies were under selling pressure as a result of recent advancements in AI.[11]

At the same time, not all sectors are equally vulnerable. Activities requiring trust, accountability, explainability, or highly variable real-world interactions are likely to retain a meaningful human role, in our view. We anticipate that companies possessing high-quality proprietary data, embedded switching costs, and operating in complex regulatory environments are also likely to see more limited risk of disruption because of AI adoption, at least in the near term.

Disruption raises the prospect of greater valuation dispersion within and across sectors. For long-term investors, it argues in favor of diversified exposure that can accommodate both winners and unexpected beneficiaries.

Capital allocation beyond technology

One of the most important shifts ahead may be the diffusion of AI opportunity beyond US mega-cap technology. Early leadership has been highly concentrated in the US and semiconductor names, particularly in South Korea and Taiwan. While we think these companies will continue to play a significant role in the investment landscape, alternative paths in the embrace of AI are emerging elsewhere. China is a stand-out example of this, where implementation has followed a low-cost model utilizing open-source models with less computational requirements than Western leading-edge offerings.

National AI strategies are likely to continue to define the playing field, in our view. Emerging questions around overinvestment, regulatory intervention, geopolitical fragmentation, and social backlash all present meaningful risks. Shifts in labor bargaining power, changes in income distribution, and evolving norms around data ownership and intellectual property may also drive differences in regional AI strategies. These factors are likely to influence both economic outcomes and political responses over time.

11 Based on Bloomberg L.P. data and Invesco Strategy & Insights analysis, as of 31 May 2026.

What this could mean for long term investors

We believe AI is still in its early economic stage. The first chapter rewarded speed and scale. We anticipate the next phase will reward selectivity and real economic impact, even as the data center build-out looks set to continue. In our view, returns are likely to be more uneven among AI players, less index driven, and increasingly tied to how AI is integrated into businesses, sectors, and societies.

For long-horizon investors, we consider the following strategic implications:

- We focus on long-term productivity and growth effects, particularly those economies that stand to benefit from adoption of the technology.
- We expect greater dispersion in returns as disruption broadens beyond the technology sector.
- We manage concentration risk and remain alert to second order consequences.

AI is moving fast—and it is beginning to test assumptions embedded in markets and business models alike. Navigating this evolution will require patience, perspective, and a willingness to look beyond the obvious beneficiaries of today's technology cycle, in our view. As with past instances of technological disruption, we look to be attentive to how the investment landscape may change and those companies best placed to capture value in the new environment.