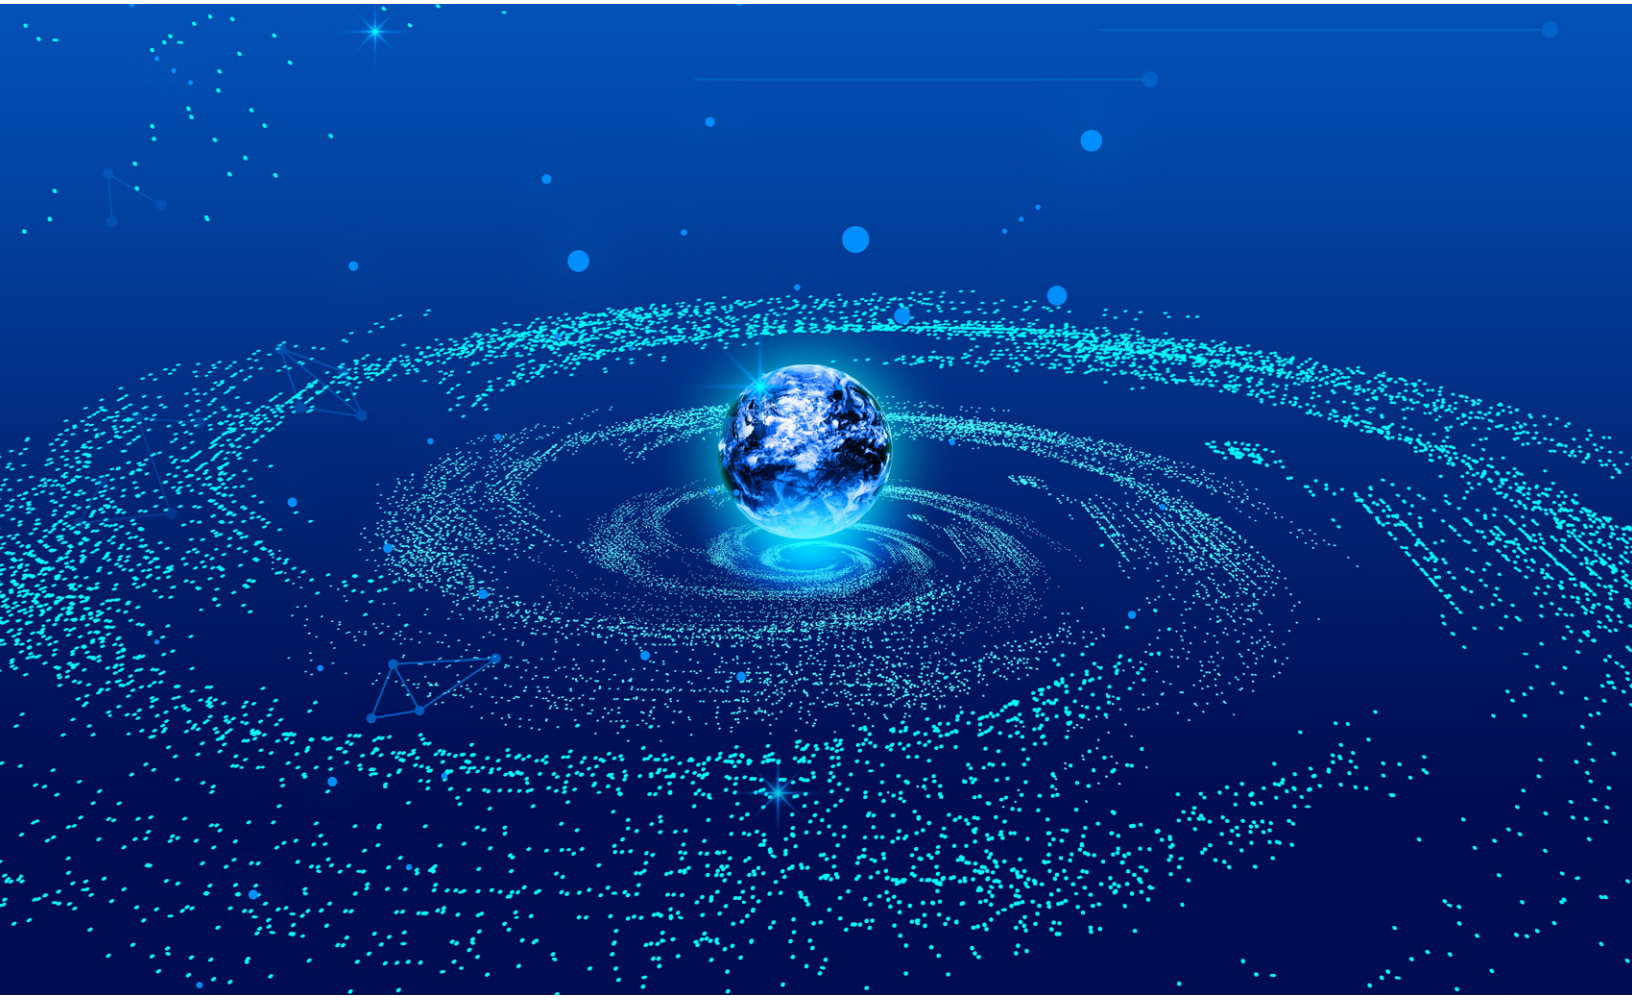
McKinsey
Global Institute

The AI economy: Interconnected forces, feedback loops and speeds of change

AI is shaped by forces far beyond the technology itself, and its effects reach just as widely. Mapping how the pieces interact helps us see what could influence its course and where hurdles and opportunities may emerge.



September 2026

McKinsey Global Institute

The McKinsey Global Institute was established in 1990. Our mission is to provide a fact base to aid decision making on the economic and business issues most critical to the world's companies and policy leaders. We benefit from the full range of McKinsey's regional, sectoral, and functional knowledge, skills, and expertise, but editorial direction and decisions are solely the responsibility of MGI directors and partners.

Our research is currently grouped into five major themes:

— Productivity and prosperity: Creating and harnessing the world's assets most productively

— Resources of the world: Building, powering, and feeding the world sustainably

— Human potential: Maximizing and achieving the potential of human talent

— Global connections: Exploring how flows of goods, services, people, capital, and ideas shape economies

— Technologies and markets of the future: Discussing the next big arenas of value and competition

We aim for independent and fact-based research. None of our work is commissioned or funded by any business, government, or other institution; we share our results publicly free of charge; and we are entirely funded by the partners of McKinsey. While we engage multiple distinguished external advisers to contribute to our work, the analyses presented in our publications are MGI's alone, and any errors are our own.

You can find out more about MGI and our research at www.mckinsey.com/mgi.

MGI directors

Shubham Singhal (chair)
Chris Bradley
Tanguy Catlin
Kweilin Ellingrud
Sylvain Johansson
Nick Leung

MGI partners

Arvind Govindarajan
Mekala Krishnan
Jan Mischke
Jeongmin Seong

At a glance

AI's advance will leave few of us unaffected. It has been less than four years since the launch of ChatGPT, yet generative AI is already changing how people think and work—and where it goes from here is deeply uncertain. Will it improve our lives or take our jobs? Lead to unprecedented growth or upheaval? Who will benefit, and who might lose?

Answering big questions about AI means seeing the whole picture, not just focusing on the technology. To help do that, we mapped the connections between AI and physical resources, institutions, businesses, markets, geopolitics, and people—all of which shape AI and change in response to it.

Risks and bottlenecks arise because different parts of the AI system are moving at vastly different paces. AI capabilities are advancing exponentially by some measures. Progress in building data centers, power grids, chip factories, and mines is more linear. Redesigning how a business and its people work spans years. As one constraint eases, others can be exposed.

In the AI era, leaders must therefore learn to manage at multiple speeds. Progress may slow when infrastructure or organizations cannot keep up, while risks can grow when adoption outruns safeguards, governance, or public acceptance. Sometimes the answer is to accelerate what is lagging; other times it may be to slow down. A systems view can help us navigate these choices.

Introduction

AI may be the most consequential technology of our time. But how far will it spread, how quickly, and with what effects? The answers matter to almost everyone. Workers want to know whether predictions of mass job losses are credible or a scare story; one widely reported, admittedly extreme, forecast says AI could eliminate as many as half of entry-level white-collar jobs within five years.[1] Investors want to know whether the enormous sums being spent on infrastructure will pay off—or leave behind the excess capacity and losses seen after earlier railway and dot-com booms.[2] Policymakers ask how tightly to regulate AI's risks without putting their economies at a disadvantage; China, the European Union, and the United States are following different paths.[3] For CEOs and CTOs, the question is more immediate: We have rolled out the technology, so why aren't we seeing it in our results?

None of these questions can be answered in isolation. They require analysis that cuts across disciplines and domains, just as the technology does. Whether jobs disappear depends on whether AI can do the tasks that make up a job as well and as cheaply as a person can, and whether new tasks emerge fast enough to replace the ones it takes over. Investment will pay off only if demand for AI services grows fast enough to fill the computing capacity being built for it—the chips, power, and data centers. Regulators face a trade-off: Move too slowly and risks may go unmanaged; move too quickly and lighter-touch rivals may pull ahead. And whether a company's rollout produces results depends not only on the technology but also on its economics and, importantly, on whether the organization adapts to capture its full value.

To understand how these forces shape one another, we need to see the whole system and the connections among its parts. Drawing on our prior research and work with clients, we propose a general framework that connects factors often considered separately and reveals relationships that are otherwise easy to miss.

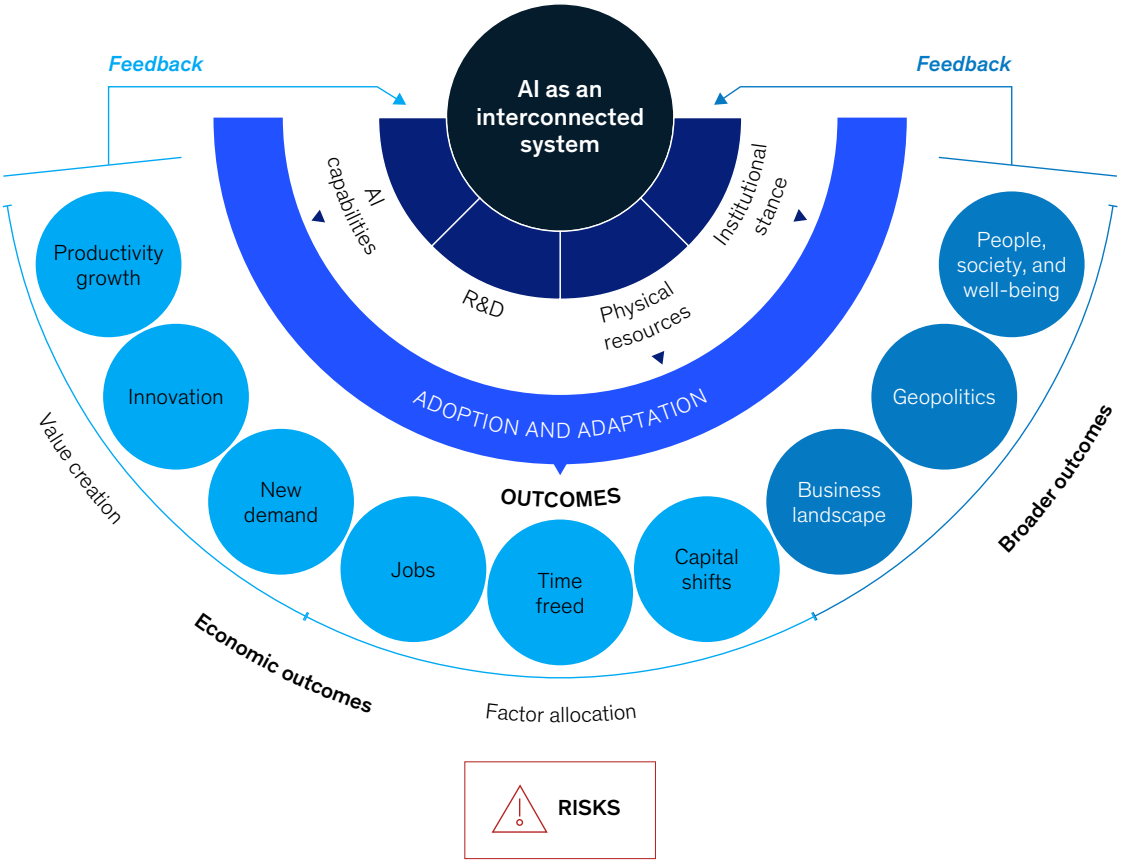
The paper begins by showing how AI's components can fall out of step with one another and the bottlenecks and risks that follow when they do. It then brings those components together to show how they fit into a single system and how changes in one part can affect others. The framework is a starting point for thinking through the big questions AI raises: how rapidly improving capabilities may translate into higher productivity and economic growth, for example, or how infrastructure, organizational change, or public opinion may affect adoption. A systems view can help us understand the forces shaping AI's path and where our choices can make a difference.

Chapter 1 looks at how fast different parts of the system are changing and what can happen when they fall out of sync. Use the map below to get oriented (Exhibit 1). Online and in the digital PDF, click on any element for a closer look; print readers can find the same material on pages 19–24.

Click on any element for a closer look.

Exhibit 1

In the AI system, changes in one part can ripple through the rest.



Source: McKinsey Global Institute analysis

McKinsey & Company

1. Mismatched timelines: Bottlenecks, risks, and opportunities

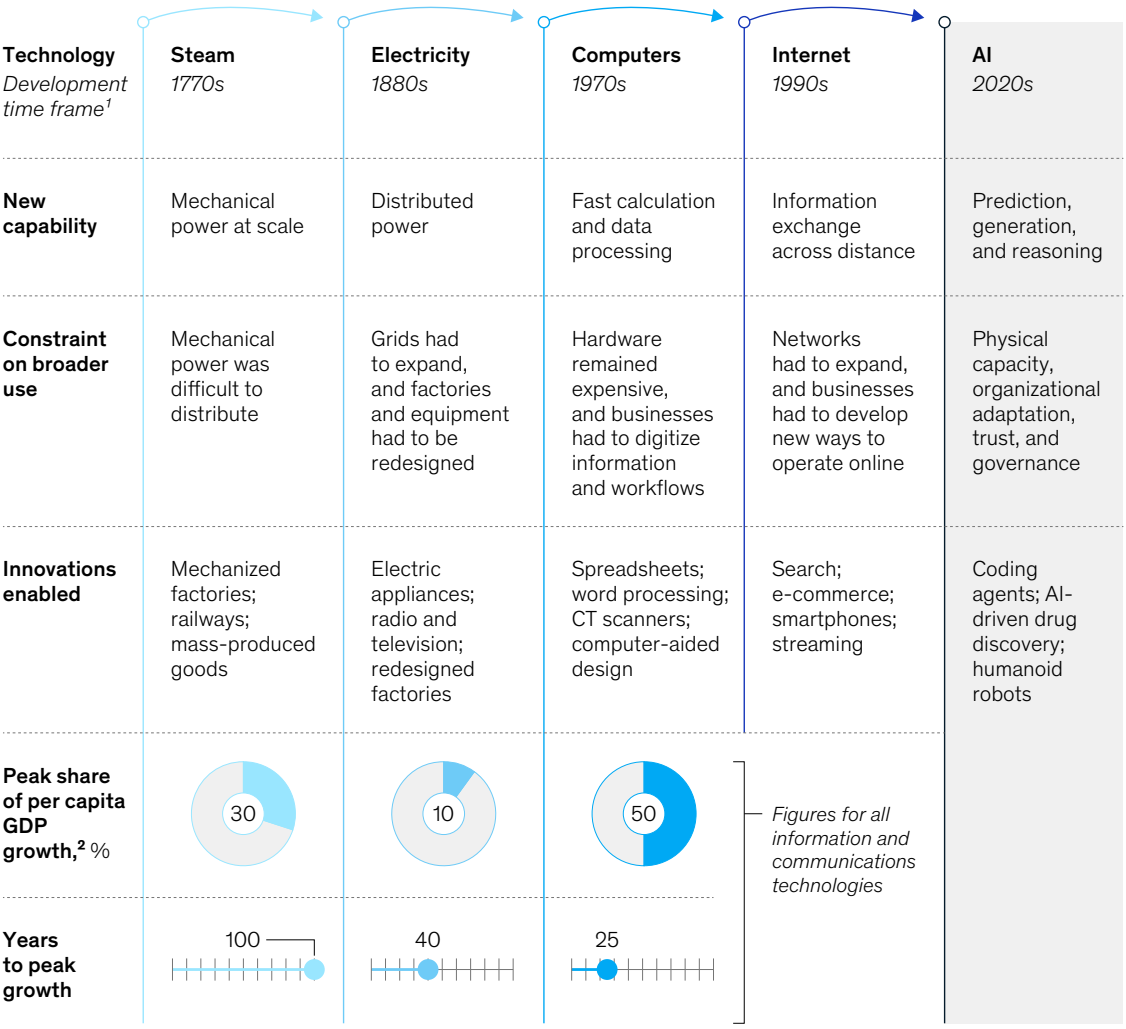
General-purpose technologies can transform economies and improve lives, but realizing their full potential takes more than the invention itself. It requires complementary innovation, new infrastructure, new rules, capital, and people and organizations willing to change how they work. The pieces rarely arrive all at once, and the gap between a new capability and its full payoff has historically taken decades to close (Exhibit 2).

In the early decades of electrification, it was difficult simply to generate enough power to run factories and industrial machinery. Later, as that hurdle eased, the limits of the power grid became apparent. The big productivity gains arrived only once factories were physically redesigned around the electric motor, a step that came later still.[4] The internet's path was much the same: First fiber-optic networks had to be built, then new business models were needed, and finally companies had to reorganize around the technology.[5]

Exhibit 2

General-purpose technologies can outpace the systems needed to realize their potential.

Innovations, constraints, and value creation associated with general-purpose technologies



Note: Figures are rounded for clarity.
[1]General-purpose technologies typically develop through many innovations over time. The development time frame approximates when each technology began to be commercially deployed. For example, the first public electric utilities emerged in the early 1880s.
[2]Maximum estimated share of economy-wide per capita GDP growth attributable to each general-purpose technology, based on the United Kingdom for steam and the United States for electricity and for information and communications technology.
Source: Nicholas Crafts, "Artificial intelligence as a general-purpose technology: an historical perspective," *Oxford Review of Economic Policy*, Volume 37(3), 2021, pp. 521–536; McKinsey Global Institute analysis

McKinsey & Company

Mismatches

AI faces a similar challenge, with many necessary parts coming together on vastly different timelines and with different shapes of change. Model capabilities are improving exponentially, doubling every few months by some recent measures.[6] Physical infrastructure builds out more linearly, often taking years. Change at companies may occur more unevenly and can take longer, as businesses develop new products and ways of working.

When the parts move out of step, it can cause bottlenecks that slow down technological advancement or deployment—such as when demand for AI chips outpaced the ability to manufacture them in 2023, discussed below. For organizations, a period of rapid adoption can just as easily be followed by a slowdown or reversal if a product fails or causes unintended harm, or if deployment of the technology does not deliver productivity gains soon enough. At other times, a mismatch can create risks, such as when corporate governance or cybersecurity lags behind the adoption of powerful new tools.

Early in the generative-AI build-out, one of the first hurdles was access to frontier chips and computing power. As those bottlenecks eased, attention shifted to electricity supply and grid connections.[7] The next constraints are more likely to involve applications, workforce skills, and organizational workflows than hardware (Exhibit 3).

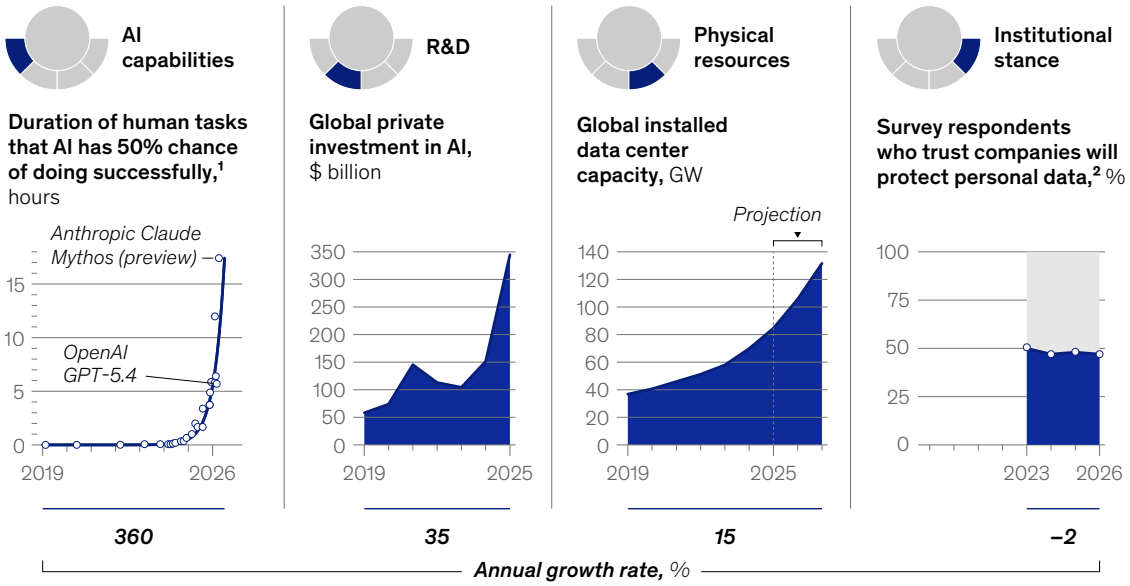
Bottlenecks

By 2023, only a year after OpenAI launched ChatGPT, demand for AI chips had outgrown the capacity to package them with the memory they require.

Nvidia, which designs most of the processors that frontier AI runs on, had built exactly the chip for the job, the H100. The problem was fusing it with memory in a single package, so that the hardware could work fast enough. Only one company could do this at the scale AI now demanded: TSMC, the Taiwanese firm that manufactures chips for most of the world's major technology companies, including Nvidia and Apple. But in 2023, demand for that packaging outpaced even TSMC's ability to keep up.[8] Buyers reportedly faced lead times of up to 11 months for a finished H100-based server.[9] The fix was more capacity, and TSMC built faster than it ever had.[10]

Exhibit 3

AI's foundations move at fundamentally different speeds.



[1]At the time of writing, no data were available for models released after April 2026. No data are reported for open-weights models. METR reports that task duration measurements above 16 hours are less reliable due to the limits of the task suite.
[2]The same question has been asked for 4 consecutive global panels; however, the number of countries polled has ranged from 30 to 32 across samples.
Source: METR; Stanford HAI *AI Index report* (2026); McKinsey Data Center Supply Model; Ipsos AI Monitor (2023–26); McKinsey Global Institute analysis

McKinsey & Company

Other hurdles emerged in a completely different part of the system: electricity. Building and connecting enough power capacity—a challenge for utilities and grid operators, not chipmakers—is another constraint, one familiar from the early era of electrification. A single data center can consume as much power as a small city, and when local grids cannot handle that much new demand, the utility has to build new power lines and substations before the facility can turn on—construction that requires its own permits and land, behind a long queue of other projects waiting for the same thing. Several hyperscalers, the handful of companies building AI at the largest scale, have run into this wall. On an October 2025 earnings call, Microsoft's CFO told investors that the company's cloud computing platform, Azure, which is used to handle AI workloads, had been short on operating capacity for several quarters. Electricity was the obstacle. Microsoft had AI chips it could not deploy because it lacked powered, ready-to-occupy data center space.[11] The median new US power plant now takes about five years from interconnection request to commercial operation, up from less than two years in 2008.[12]

Data center construction itself is now shaping up as a bottleneck, especially in the United States, which accounts for nearly half of global installed data center capacity.[13] AI's compute, power, and water costs per query have fallen as models get more efficient, but total demand is rising faster than efficiency gains or new supply can offset.[14] Construction is expanding rapidly and still can't keep pace. Many facilities are pre-leased even before they open.[15]

In some places, local communities are pushing back. Concerns about rising electricity prices and large, unsightly facilities have triggered backlash, leading to permitting delays and, in some cases, construction freezes. At least 75 data center projects worth some $130 billion were blocked or delayed in the United States by community opposition in the first quarter of 2026 alone—matching all of 2025 in three months, according to Data Center Watch.[16] In July, New York's governor signed the first-ever statewide permitting freeze on large data centers.[17]

Other potential constraints are emerging even as the data center drama plays out. Supplies of copper and other critical minerals needed for chips and wiring are showing signs of stress, but new mines and processing plants take years to permit and build. Unless current plans are accelerated, the world could face a sizable copper shortfall by the mid-2030s, according to the International Energy Agency.[18] AI will be competing with many other industries for those materials. And physical resources are not the only constraint: Data center developers also report shortages of engineers and other workers, who take time to train.[19]

AI's capabilities, meanwhile, are advancing on an entirely different timescale. The complexity of tasks that AI can reliably handle—measured by how long the same tasks take a human expert—has doubled roughly every seven months since 2019, and roughly every four months since 2023.[20] This is partly because AI increasingly contributes to the research that improves it; existing models can now write the code, generate the data, and run the evaluations used to train the next generation of models. Technologies available today could, in theory, automate more than half of current work hours in the United States, Europe, and Latin America.[21]

So, the tools are here. But the reorganization is not. In a McKinsey survey of 1,719 respondents across 97 countries, conducted in May and June 2026, nearly nine in ten said their organizations regularly use AI in at least one business function. Yet outside a small group of AI high performers, representing 6 percent of the sample, only about one-quarter reported redesigning workflows around AI rather than inserting AI into existing ones.[22]

Risks

Governance is also lagging. A separate McKinsey survey of about 500 organizations on AI trust and governance, conducted in December 2025 and January 2026, found that only about 30 percent had reached level 3 or higher on a four-level maturity scale.[23] Strategy remains underdeveloped as well. Only 22 percent of operations leaders report a fully developed and implemented enterprise AI strategy, even though more than half name strategy as the single biggest driver of AI returns.[24]

The problem usually stems from fragmentation. Responsibility for AI is split across IT, legal, data, HR, risk, and strategy, with no single owner. Day-to-day, oversight often defaults to the teams closest to the technology—IT or engineering—simply because no one else has claimed it. At the same time, 70 percent of global CEOs say they have the greatest influence on their company's AI strategy.[25] The gap between who actually runs AI and who is nominally accountable for it creates its own risk. Often, no one is positioned to answer basic questions about autonomous AI systems—which actions they may take, which decisions need human approval, who is accountable when something goes wrong.

Similarly, roughly one in three organizations still have no formal process to check AI systems for security risks before deploying them.[26] Weak links may include smaller players, such as local governments, which may lack the resources for advanced security but remain connected to larger systems, creating vulnerabilities that extend beyond their own networks.[27]

As adoption increases, so does the risk of cyberattack. More AI systems in more places give criminals more surface area to exploit. The technology is also becoming a more capable tool for hackers. In April 2026, the UK AI Security Institute (AISI) found that Anthropic's Claude Mythos Preview could autonomously execute multi-step attacks on test networks that would take human professionals days of work. Just two years earlier, the best available models in AISI's evaluations could barely complete beginner-level cyber tasks.[28] The risks extend beyond businesses and governments to individuals. AI-generated phishing and fraud are harder to spot because AI can produce personalized, convincing messages at a scale no human scammer could match.

Just as with a large governance failure—that causes or allows major financial or physical harm—a single large AI-enabled cyberattack could trigger public backlash and stricter regulation that slows adoption well beyond the companies directly involved.

Opportunities

Major breakthroughs could also change several of these timelines at once. A leap in model architecture—or in an adjacent field like quantum computing—could loosen several of AI's constraints at the same time, the way the invention of the transistor simultaneously advanced computing, communications, and electronics. The benefits could flow the other way, too. AI-driven breakthroughs in fields like energy or biotechnology could unlock innovation and growth well beyond the AI industry itself.

2. AI as a system: How the pieces fit together

Many of today's hardest challenges and most important developments must be understood as complex systems. This is the case for AI. Its effects come from the interaction of many interdependent parts, not from any one part alone. Capability, R&D, physical resources, institutions, and adoption each shape the others (see sidebar, "Why a systems view matters").

Sidebar

Why a systems view matters

Systems thinking is probably as old as philosophy itself, but it was applied directly to business and technology in the 20th century by people studying how organizations actually behave. Biologist Ludwig von Bertalanffy spent decades arguing that breaking a phenomenon into isolated parts misses essential properties that surface only when the parts are observed together, an idea he finally consolidated in his 1968 book *General System Theory*.[1] The engineer Jay Forrester, in the 1950s, studied why factories lurched between hiring booms and mass layoffs even when customer demand barely moved.[2] The cause was not demand; it was information delay. Orders took time to travel from retailer to distributor to factory, so by the time a factory ramped up hiring in response, the original signal was already stale, and it overshot. The factory would then cut hard to correct, and the cycle would repeat. The swings were caused by the system itself, not the macroeconomy. The management theorist Russell Ackoff, a pioneer of systems thinking, illustrated the idea with a story about cars.[3] Take the world's best automotive engineers, he said, and ask them to identify the single best engine, the best transmission, the best brakes, and so on, from every car in production. Now put those "best" parts together. The result is not the best car in the world—it may not run at all, because the parts were never designed to work together. The same logic applies to AI. Optimizing one piece in isolation—faster chips, more capital, a new rule—will not improve outcomes if the real constraint lies elsewhere.

1 Ludwig von Bertalanffy, *General System Theory: Foundations, Development, Applications*, George Braziller, 1968.
2 Jay W. Forrester, *Industrial Dynamics*, MIT Press, 1961.
3 Russell L. Ackoff, "Towards a system of systems concepts," *Management Science*, July 1971, Volume 17, Number 11. For automobile analogy, see Ackoff's 1994 presentation for the W. Edwards Deming Institute.

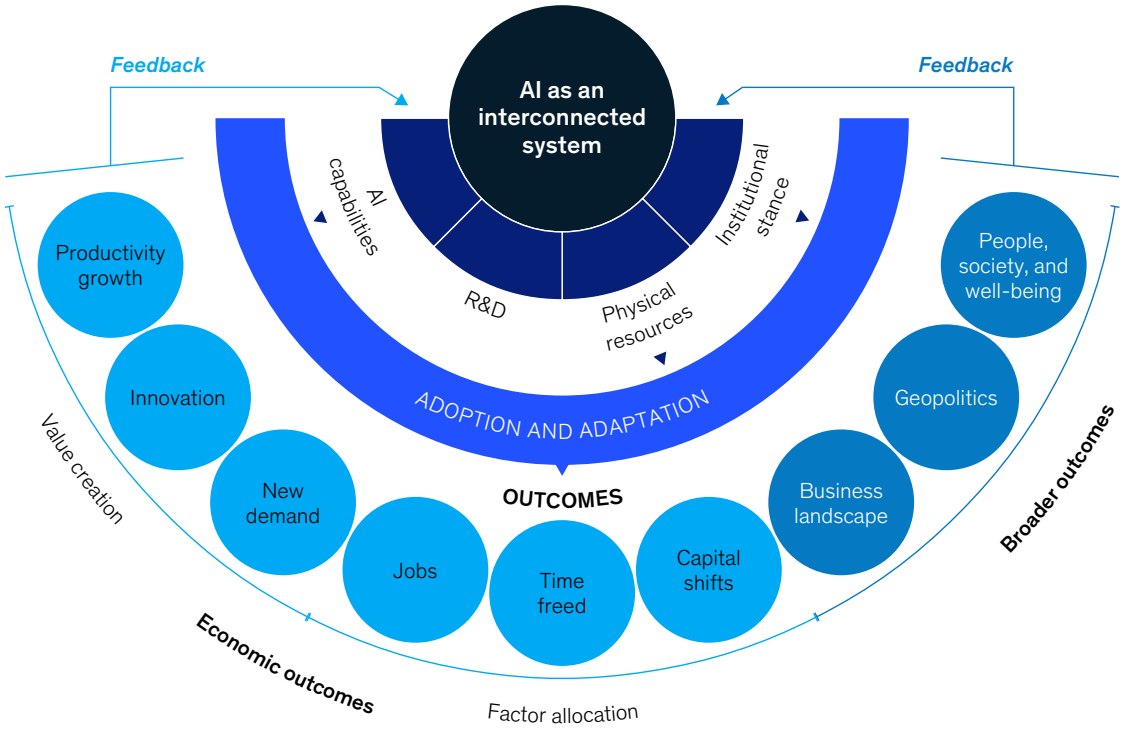
Returning to the framework introduced in Exhibit 1, we divide the system into three analytical layers (Exhibit 4). Foundations shape what AI can do and the conditions under which it develops. Transmission is how people and organizations put AI's capabilities to work. And outcomes are the effects that follow. The framework organizes the factors we think matter most; it is not a full inventory, and the categories overlap. Innovation, for example, is intertwined with new demand.

Better models can encourage adoption; adoption attracts investment; investment expands infrastructure and R&D; which in turn speed up development of yet better models. Conversely, shortages of power, weak governance, or public concern can slow deployment. The system is in motion, with feedback running through every part of it.

The online report includes an interactive version of this exhibit that shows how the elements connect.

Exhibit 4

In the AI system, changes in one part can ripple through the rest.



Source: McKinsey Global Institute analysis

McKinsey & Company

Foundations

AI's foundations—capabilities, R&D, physical resources, and institutional stance—determine what AI can do and how fast it can be built. Each shapes the others: R&D drives capabilities, capabilities require resources, and institutions influence how AI is developed and deployed.

AI capabilities

What AI can do extends beyond models and agents to embodied systems such as robots. Capabilities are the product of R&D, human expertise, data, computing resources, and the capital that funds their development. As they advance, new uses become possible, organizations reconsider what can be automated, and governments reassess AI's risks. Whether the technology's potential is realized depends on reliability, cost, public acceptance, and organizational readiness.

R&D

Research and development on frontier AI is mainly concentrated in a small group of companies in the United States and China.[29] It improves model capabilities directly, and it makes physical resources go further by raising the efficiency of chips and data use. Progress is tied to private and public investment, computing power, data, electricity, talent, and government policy. Profits can fund the next round of research, and expectations of future profits attract investment; export controls can limit access to resources and affect what researchers work on or how they innovate; breakthroughs can shift markets and geopolitical advantage.

Physical resources

The list of physical resources needed to build AI data centers and keep them running is long: Chips, copper wires, cooling systems, land, electricity, and water are just some of the required ingredients.[30] Strong AI demand, combined with this wide range of inputs, is reshaping global trade and supply chains, and sometimes provoking local disputes over utility costs and the visual impact of vast new facilities. At the global level, some key resources are geographically concentrated and can become instruments of geopolitical leverage, as China's restrictions on rare earth exports make clear.[31]

Institutional stance

Institutions—the rules, norms, and organizations that govern AI—shape the technology at every stage. At the outset, export controls and rules on data, infrastructure, and safety can influence what gets researched and built, where, and at what cost. As organizations put AI to use, privacy and liability rules, professional standards, and company policies can shape how quickly and widely it is adopted. As the technology spreads, its effects feed back into institutions themselves, prompting new regulation and governance practices, influencing social norms, and building or eroding public trust. Many basic questions remain open: how much testing should be required before deployment, what incidents companies must disclose, and who is liable when the technology causes harm.

Transmission

Think of technological capabilities, physical resources and institutional stance as the brains, brawn, and rules that shape AI's development. R&D supplies the nutrients. All together, they determine what AI can do and at what cost. The next question is whether—and how—people and organizations put those capabilities to work.

Adoption and adaptation

Consumers, businesses, and governments are all adopting AI, but at different speeds. Individuals can pick up new tools right away, and millions experimenting with them almost immediately can help fuel innovation faster than institutional adoption often allows. Businesses move more slowly and often unevenly. Redesigning workflows around a new technology takes time, and not everyone inside an organization wants it to happen—workers may resist for fear of losing their jobs, for instance. Adoption is also uneven across companies of different sizes.[32] Greater financial resources may help large companies invest in AI, while smaller firms can benefit from greater agility; midsize companies, with less of either advantage, may struggle to keep pace. As organizations scale AI, successes attract further investment in infrastructure and R&D, while failures expose weaknesses in governance, security, and skills that shape the next round of regulation and organizational change.

Outcomes

Putting AI to work changes economies, societies, and institutional stances. Some effects appear quickly, such as higher productivity in specific tasks. Others emerge gradually as organizations redesign work, investment shifts, new products appear, and governments respond. As with earlier general-purpose technologies, some of AI's largest effects may come later, from innovations that are difficult to anticipate today. All of these outcomes, in turn, feed back into AI's development and adoption.

Productivity growth

AI has already made many individual tasks more efficient, but turning those gains into higher productivity across a company or economy requires more. Companies must often redesign entire workflows around new technologies. To achieve economy-wide gains, those changes must then spread across many companies and industries. Growth must also be part of the equation. Without it, improvements in labor productivity could reduce employment and, in turn, consumer spending, which accounts for more than half of global GDP.[33] Lower spending could reduce returns on capital, discourage further investment, and generate backlash.

The broad benefits of new technologies can take time to materialize. This is why transformative technologies can seem ever-present and yet leave little trace in macroeconomic data. That disconnect is familiar enough to have a name: the productivity paradox. Robert Solow famously captured the idea early in the digital era with his 1987 quip that "you can see the computer age everywhere but in the productivity statistics."[34] The paradox resurfaced with the internet and is being debated again with AI.

Innovation

AI is already speeding invention and discovery. AI systems that predict protein structures have compressed into hours parts of biomedical research that once took years, while models screening vast molecular libraries have identified new antibiotic candidates.[35] Progress in materials science is accelerating in much the same way.[36]

Another important arena of innovation is commerce. AI could help suppliers better understand buyer intent, improving matches between products and users and spurring new solutions to meet hidden or latent needs. The internet transformed commerce by reducing search costs and enabling marketplaces that connected vast numbers of buyers and sellers globally. AI could take this further, shifting commerce from the internet era's focus on attention toward a greater focus on intention.

History shows that technologies have often evolved far beyond their original uses. Steam power was first developed to pump water out of mines but later transformed shipping and enabled railroads. Electricity was initially sold as a replacement for gas lighting before powering homes and enabling factories to reorganize production. Consider also how electricity paved the way for entirely new products and industries. By making refrigeration possible, for example, it transformed how food could be produced, transported, sold, and consumed. AI could similarly enable new products, businesses, and business models that are difficult to anticipate today.

New demand

AI may generate demand through many channels. Building the AI economy creates a need for AI models and applications, robots and infrastructure, cybersecurity, and support for organizational change.

AI may also create markets for services that take on tasks people currently do for themselves, such as drafting letters, planning trips, helping children with homework, or managing household finances. Advances in robotics could extend this to physical household tasks such as cleaning, cooking, or yard work.[37]

Automation itself could spur demand by simply making things cheaper and more accessible, echoing the logic of Jevons paradox, which shows why greater efficiency can increase total consumption rather than reduce it.[38]

Demand may also increase if productivity gains yield higher incomes and more free time, giving people more money and time to spend on goods and services. One well-known example of this dynamic played out in the US auto industry about a century ago, when the introduction of mass production was followed by higher wages and shorter working hours, culminating in a five-day workweek. Henry Ford explicitly argued that workers needed both the money and leisure to consume what industry produced—including automobiles.[39]

Jobs

What happens to jobs is the most worried-over consequence of AI. Yet history suggests that major technological change creates new jobs, even as it eliminates the need for others. About one-third of current US employment, for example, is in occupations that did not exist before earlier general-purpose technologies enabled them—truck drivers, car mechanics, electricians, software developers, cybersecurity analysts.[40] Whether AI ultimately creates more jobs than it eliminates will depend on whether innovation leads to new kinds of work and generates enough new demand fast enough to offset losses from automation.

Within jobs, AI is already changing how work gets done as processes are reorganized around human-machine collaboration.[41] Unlike earlier forms of automation, which mostly affected physical labor, AI is also changing knowledge work in fields such as computing, accounting, and law. As AI takes on more routine tasks, productivity may rise, freeing people to do more complex tasks. Skills where people may retain an advantage over AI, such as judgment, curiosity, and conflict resolution, could become more valuable. But the division of work between people and machines will continue to shift as AI becomes more capable. For some entry-level workers, the technology may already be making the first rung on the career ladder harder to reach.[42] For other lower-skill workers, there is evidence that AI can narrow the gap with more experienced workers.[43] AI-driven shifts in employment could favor higher-wage occupations over lower-wage ones, although some workers may struggle to move into the new roles.[44] How broadly workers benefit from AI—or feel threatened by it—will help shape public support and, in turn, adoption and regulation.

Time freed

By making tasks faster or taking them over entirely, AI can free time from both paid and unpaid work. That time may be redirected toward more work, new tasks created by AI, learning, entrepreneurship, unpaid activities, or leisure and rest. Greater efficiency may also increase demand for existing work, absorbing some of the time initially saved.

Some freed time may be unwanted. People may work fewer hours because less of their labor is needed, even if they would prefer to work more. How freed time is ultimately used will shape both output and demand. Sustained economic growth will require people to have the income and willingness to buy the additional goods and services the economy can produce.

Capital shifts

The focus of AI investment has changed as the technology's bottlenecks have shifted. When there was a shortage of computing capacity, capital poured into chips and data centers; as electricity and grid connections emerged as constraints, investment moved toward even larger infrastructure. As the need for capital has grown, the financing mix has broadened well beyond venture capital to include public equity markets, private credit, and public debt markets. Circular financing, in which AI companies invest in one another or make commitments that help finance one another's growth, is also becoming common. These arrangements can accelerate the build-out but also make it harder to distinguish demand generated within the AI ecosystem from demand coming from outside it—and create interdependencies that can cause losses at one company to spread to others.

The scale of spending is enormous. McKinsey estimates that global investment in data centers will reach $7 trillion cumulatively between 2025 and 2030, much of it AI-driven.[45] More than 40 percent is expected in the United States. In 2026 alone, US imports of AI-related goods are on track to reach around $750 billion—about as much as its imports of cars, oil and gas, and pharmaceuticals combined.[46] That spending could come at the expense of investment elsewhere: Total private investment as a share of GDP has remained flat even as investment in data centers has soared.[47] There is also a risk of overbuilding. Earlier railway and internet booms saw investment race ahead of demand, leaving excess capacity and losses for many investors even though the infrastructure ultimately generated substantial economic benefits.[48]

Business landscape

AI is changing relationships among companies, customers, and suppliers. Some businesses are bringing work in-house as AI reduces the need for outside help and increases the value of connecting data, workflows, and decision-making. Others are becoming more dependent on suppliers that control critical data, infrastructure, or access to customers. Meanwhile, AI is lowering barriers to entry for smaller firms and startups, allowing them to handle coding, marketing, customer service, and administration with fewer people or outside providers. Shifting boundaries can turn suppliers and customers into competitors.

As models improve and spread, access to the best technology may provide only a temporary edge because today's leading model can be quickly overtaken or matched. AI can also commoditize capabilities, making skills that once distinguished companies widely available. That could shift advantage toward things AI cannot easily replicate or bypass—critical workflows, customer relationships, or even administrative requirements such as documentation needed for medical billing. Companies that control these choke points may capture outsize value even when they play only a small role in producing the underlying product or service. As a result, businesses that develop or make the best use of AI may not always capture the greatest economic gains from it.

Geopolitics

Companies are not the only ones competing for an AI advantage. The United States and China are racing to develop the technology, with the outcome likely to affect the balance of power between them. So far, the United States narrowly leads in frontier models and has more high-end compute, while China leads in open-weight models and industrial AI deployment.[49] Whoever harnesses AI most effectively is likely to gain a real edge. Faster growth and innovation can strengthen a country's economic influence and provide more resources for defense, while AI itself can directly enhance military capabilities.

Control over chips, compute, models, and standards could become a source of leverage, much as dependence on foreign oil and gas can leave countries vulnerable to pressure from suppliers. A country reliant on AI built and run elsewhere could expose sensitive data to outside access, and public services and critical infrastructure to disruption. In the worst case, these imbalances could spill into trade, cyber, or military conflict.

People, society, and well-being

AI's effects extend to billions of people everywhere, changing how they think and the kinds of thinking they do. The technology can put forms of expertise that were once expensive or difficult to access—from tutoring and translation to help navigating health, legal, or financial questions—within reach of many more people. But relying on it for routine thinking may weaken the habits that build reasoning skills, a particular concern for children still learning to think for themselves.[50] For workers, AI may take on routine tasks, freeing people to spend more time on work that calls for judgment, creativity, problem-solving, and human interaction. But as people produce more in less time, work may also become more cognitively demanding, increasing the risk of overload and burnout.[51]

How people feel about the technology matters, too; it shapes how governments regulate it and how businesses use it. Surveys of the general public suggest that people in advanced economies tend to be wary of AI. In the United States, some 70 percent of survey respondents say AI will make their personal information less secure. Roughly half of American adults under 30 expect AI to have a negative effect on society.[52] Views are more positive in developing and middle-income countries, where people are more likely to expect AI to improve their country's economy and the job market.[53] A similar geographic divide appears in a study of active Claude users. Anthropic interviewed more than 80,000 users of its chatbot across 159 countries and found that people in Africa, South America, and much of Asia were more optimistic than those in Europe and the United States, with those in developing regions more often seeing AI as a source of economic opportunity.[54]

3. The value of a systems view

A systems view can help identify where action could have the greatest effect. For businesses, it can reveal what is holding back results; for investors, how opportunities and risks are shifting; for policymakers, where rules or public investment could help; and for individuals, which skills and opportunities may become more valuable.

For business leaders

— *Look beyond efficiency for lasting advantage.* AI can reduce costs and make existing work more productive, but competitors with access to the same technology may achieve similar gains, driving down prices and eroding any initial advantage.[55] Long-term outperformance requires companies to combine AI with advantages that rivals cannot easily replicate, creating competitive moats.[56] These can include proprietary data, scale, strong customer relationships, and access to scarce physical assets, combined in a business model that delivers greater value to both customers and suppliers.

— *Build new things.* The first uses of major new technologies are rarely the most consequential. Electricity began as a replacement for gas lighting before transforming industries and enabling entirely new forms of economic activity. AI may likewise give rise to new products, businesses, and markets that are difficult to anticipate today. Using the technology to automate existing work is a first step, but leaders should use the resources freed up to build new sources of growth. The potential impact extends beyond individual organizations: A small number of standout companies have driven a disproportionate share of economy-wide productivity growth, largely through shifts in strategy and moves into new businesses that reshape their industry segments, rather than by making existing businesses more efficient.[57]

— *Learn to learn faster.* AI capabilities can improve exponentially, much faster than many companies can change their technology, workflows, or organization. New businesses may have an edge because they can build around AI from the start, without having to adapt or throw out legacy systems. But established companies have an advantage, too: their scale gives them more customers, workers, and operations across which to test new approaches, gather feedback, and learn what works. To make that advantage count, they need to build the capacity to adapt their systems and ways of working quickly enough to keep pace with evolving AI capabilities.

— *Give people a reason to support change.* Whether employees, customers, and communities see AI as making their lives better or worse will shape their willingness to accept and use it. Leaders therefore need to deliver benefits that people can see and experience rather than simply explaining what they are trying to achieve with the technology. If people experience better products, services, or ways of working, they may be more likely to support change; if they feel that the harms outweigh the benefits—that their jobs, livelihoods, and privacy are at risk—resistance may mount. A systems view shows that despite technological progress, businesses may struggle to create value if societal resistance is high.

— *Develop a view of how AI could reshape your industry.* AI could change who has an advantage, how companies compete, and where profits are made. It may lower barriers to entry, shift power among companies, customers, and suppliers, or change the advantages of scale. Leaders should consider how these forces could play out in their industry—and what they mean for strategy.

— *Watch the whole system.* AI's trajectory will depend on more than improvements in the technology itself. Changes in adoption, infrastructure, investment, regulation, and economic outcomes can reinforce or constrain one another. The framework in this report identifies metrics for tracking these changes across the AI system.

For investors

— *Follow the bottlenecks.* Constraints can create investment opportunities because they generate demand for solutions. When computing power was scarce, opportunities emerged around chipmakers and frontier-model companies. As constraints shift toward electricity, infrastructure, and putting AI to productive use inside organizations, new opportunities may emerge around solving those problems.

— *Expect different growth paths.* Exponential growth in AI capabilities or usage will not translate into exponential growth across every business that uses or supplies the technology. Model capabilities can improve rapidly, while physical infrastructure expands much more slowly and companies may take years to translate AI into productivity and profits. Investors should be clear about which of these growth paths their assumptions depend on and not mistake rapid growth in one part of the system for rapid growth in another.

— *Look past adoption.* A company can adopt AI, even on a grand scale, without changing its long-term prospects. For signs of lasting advantage, investors should ask: Are productivity gains supporting growth rather than cost reduction alone? Are new products, businesses, and markets emerging rather than simply improvements to existing ones? Is the company overcoming legacy technology and organizational constraints? And is it shifting resources from the existing business toward new opportunities?

— *Look for correlated risks.* The performance of investments in the AI system may be highly correlated even when the companies involved operate in entirely different industries. Investments in chips, data centers, power, and AI software, for example, could all be affected by the same slowdown in AI demand or investment. Investors should account for these connections when assessing how diversified their exposures really are.

For policymakers

— *Address system constraints.* What holds AI back will differ by country or region and will change over time, from electricity and digital infrastructure to permitting, skills, and institutional capacity. Policymakers should identify the constraints that matter most in their economies and focus their efforts on relieving them. As the constraints change, policy should move with them.

— *Promote competition.* AI may lower the cost of creating new products and services. To help make that opportunity broadly available, policymakers should ensure that rules do not unnecessarily favor incumbents or make it harder for new entrants to compete and grow. They should also broaden access to the skills, tools, and capital people and companies need to build with AI.

— *Make policy responsive.* Regulators need to protect the public from harm. But AI capabilities and applications can evolve faster than laws and regulations. Not every new use or risk can be anticipated and some predicted risks may never materialize. Rules therefore need to be adjusted as evidence changes rather than trying to settle every question in advance. Requiring the industry to self-regulate is a good first step. Staying close to businesses, workers, researchers, and the public to monitor impacts across the full system can help policymakers see when more substantive regulatory adjustment is needed.[58]

— *Build trust and confidence.* AI could support economic growth and improve living standards for decades, but realizing that potential may depend on public willingness to accept the changes it brings. Policymakers need to address legitimate concerns about AI's potential effects on jobs, privacy, and safety while also making the case for what AI could make possible. As noted above, responsive policy

can help build trust. Confidence also depends on whether people believe they can share in the economic opportunities AI creates. Proactive policies that expand access to skills, technology, and the resources needed to start and grow businesses can help.

— *Foster investment.* Decide which AI capabilities must be developed at home and where reliable access through partners is enough. Public investment, incentives, and partnerships can help strengthen computing infrastructure, energy, research, and talent without trying to replicate the entire AI ecosystem domestically. Some governments are already pursuing versions of this approach: the United Kingdom has established a Sovereign AI Fund, while the European Union's InvestAI initiative aims to mobilize public and private investment in AI infrastructure.[59]

For individuals

— *Cultivate AI fluency.* As AI becomes a bigger part of work and everyday life, people will need to become increasingly fluent in using the technology. AI fluency also means knowing when to use the technology and when to rely on human judgment. Skills such as creativity, empathy, adaptability, and critical thinking may become more valuable as the division of labor between people and AI changes. That boundary will continue to shift, requiring people to keep learning and adapting.

— *Be thoughtful about what you share.* AI tools can become more useful when people give them personal information as context, but that can also create risks. Before sharing sensitive financial, health, or personal information, consider where it could end up, who might gain access to it, and how it could be used.

— *Expand your ambitions.* AI gives individuals access to capabilities that once required specialized expertise, substantial resources, or large teams. This can change not just how people do things, but what they can realistically attempt: learning a new field, creating a product or service, starting a business, or pursuing an idea that previously might have been out of reach. AI can allow people to think bigger about what they might accomplish—and give them the ability to pursue those ambitions.

Conclusion

History suggests that the economic impact of new technologies may be arriving faster with each successive wave. Steam power took about a century from commercialization to its peak impact on economic growth, electricity about 40 years, and computers and the internet about 25. It is too early to know whether AI will compress that timeline further. But the possibility adds urgency because businesses, workers, and governments may have less time to adjust.

AI's trajectory will not be determined by the technology alone. Models will improve, costs will fall, and new uses will emerge. But what follows will depend on how capabilities interact with physical resources, institutions, organizations, markets, and people—all moving at different speeds. A constraint in one part of the system can slow progress elsewhere; removing it can reveal another. And as AI spreads, the changes it produces will feed back into the system itself.

That is why the whole system matters. The task is not to predict a single AI future, but to understand how the pieces interact, where mismatches may emerge, and how the choices people make can alter what happens next.

Appendix: The AI system in detail

The following pages take a closer look at the elements introduced in Exhibit 1, with key data and context on each and metrics to watch as the system evolves.

FOUNDATIONS

TRANSMISSION

ECONOMIC OUTCOMES

BROADER OUTCOMES

FOUNDATIONS

AI capabilities

Capabilities describe what AI systems can do.

16+ hours

The leading model succeeds half the time at tasks that take humans more than 16 hours to complete.

What to know: AI capabilities also extend into the physical world. Embodied AI systems such as robots are gaining dexterity and becoming better able to navigate and manipulate objects, although their performance becomes less reliable as they move from settings designed for robots into more varied real-world environments.[1] Large language models, meanwhile, are getting better at solving harder problems as well as completing longer tasks. For example, the leading score on a graduate-level science benchmark reached 95% in 2026, up from 39% roughly three years earlier.[2]

Trajectory: Exponential—the complexity of tasks AI can reliably complete has shown a sustained pattern of doubling since 2019.

Metrics to watch: Scores on difficult AI benchmarks such as Humanity's Last Exam; the length of tasks AI agents can reliably complete; robots' success rates on complex tasks in changing physical environments; and the cost of using AI at a given level of performance.

Source: METR

FOUNDATIONS

R&D

R&D advances AI and the technologies and resources that support it.

$290B

Global private AI investment in 2025—almost 5 times the 2019 level.

What to know: R&D advancing AI extends well beyond models. Research on computing, energy, cooling, materials, and other technologies can also expand what AI can do or make it cheaper and easier to deploy.

Trajectory: R&D often advances incrementally, with occasional breakthroughs producing large jumps in capability or efficiency.

Metrics to watch: AI research publications and citations; computing power used in the largest AI training runs; granted patents in AI-enabling fields; and open-source AI projects on GitHub.

Note: Public investment is also an important source of AI R&D funding in some economies, but comparable global data are not available.

Source: Stanford Institute for Human-Centered Artificial Intelligence, 2026 AI Index report

FOUNDATIONS

Physical resources

Physical resources include the infrastructure, equipment, energy, and materials needed to build and run AI.

15 GW vs 18 GW

Global data center capacity grew by 15 gigawatts in 2025, lagging behind a roughly 18-gigawatt increase in demand.

What to know: The list of physical resources AI requires is long, from chips and electricity to water, land, cooling systems, and minerals. Expanding supply can take years. Greater efficiency can reduce the resources needed for each AI task, but cheaper AI can also encourage more use and increase total demand.

Trajectory: Linear—physical capacity expands on construction and manufacturing timelines, typically much more slowly than AI capabilities.

Metrics to watch: Data center capacity demand growth; data center electricity consumption; capacity awaiting grid interconnection and median wait times; leading-edge logic and memory semiconductor lead times and price indices; and data center permitting times.

[1] "Gemini Robotics 2 brings whole-body intelligence to robots," Google DeepMind, July 30, 2026; Jeremy Morgan et al., "Colosseum V2: Benchmarking Generalization for Vision-Language-Action Models," May 26, 2026.
[2] Epoch AI, "GPQA Diamond," AI Benchmarking Hub, accessed August 28, 2026.

FOUNDATIONS

Institutional stance

Institutions are the rules, norms, and organizations that guide how AI is developed and used.

47%

Less than half of survey respondents trust companies that use artificial intelligence to protect their personal data.

What to know: Institutional stances are broad and often hard to quantify. They include regulations, liability rules, professional standards, company policies, social norms, and public attitudes toward AI. They can shape how the technology develops and spreads—and change in response to its effects.

Trajectory: Institutions tend to change more slowly and unevenly than physical infrastructure can be built, and far more slowly than AI capabilities improve.

Metrics to watch: AI-related laws and regulations; company AI governance policies; professional standards for AI use; and surveys of public trust in AI.

Source: Ipsos AI Monitor 2026

TRANSMISSION

Adoption and adaptation

Adoption is the use of AI, while adaptation is the changes people and organizations make to integrate it into how they work.

89% vs 46%

In 2026, 89% of organizations used AI in at least one business function, but of these only 46% had progressed beyond pilots.

What to know: Individuals are adopting AI tools quickly, including in their professional lives. For example, in early 2026, more than 80% of US physicians reported using AI in their work, more than twice the share in 2023.[3] Organizations tend to move more slowly because they must integrate the technology into existing systems, redesign complex workflows, and sometimes overcome resistance to new ways of working.

Metrics to watch: Share of firms piloting, scaling, and fully embedding AI; AI usage intensity by industry and business function; worker and consumer use; and share of workflows redesigned around AI.

Source: Dan Tinkoff, Lieven Van der Veken, and Michael Chui, with Tara Balakrishnan, "The state of AI in 2026: On the road to ROI," McKinsey, August 25, 2026

ECONOMIC OUTCOMES

Productivity growth

AI's productivity gains may take time to spread from individual tasks to companies and economies.

2.5%

Average annual US non-farm labor productivity growth between late 2022, which marked the launch of ChatGPT, and the first half of 2026. This is slightly above the very long-run average of 2.1% annual growth since 1947. But it is too soon to tell whether this rise is structural or part of normal volatility.

What to know: AI has already made some individual tasks quicker and easier. Broader productivity gains require organizations to redesign workflows around the technology and for those changes to spread across businesses and industries.

Metrics to watch: Total factor productivity growth; output per capita; output per hour worked; and firm- and task-level studies of time use and output.

Sources: US Bureau of Labor Statistics; McKinsey Global Institute analysis

[3] "2026 Physician Survey on Augmented Intelligence," American Medical Association, March 2026. The survey included 1,692 US physicians and was conducted January 15–February 2, 2026; 81% reported using AI professionally, up from 38% in 2023.

ECONOMIC OUTCOMES

Innovation

AI can accelerate discovery and help people create things that were not previously possible.

Up to 80%

AI has shortened drug-discovery timelines by as much as 80%. Pharmaceutical companies using AI could bring 1.5 times as many new drugs to market for the same cost.

What to know: AI is already accelerating parts of scientific discovery, from identifying drug candidates to designing new materials. Over time, it could also help create entirely new products, services, and businesses, much as earlier general-purpose technologies have.

Metrics to watch: New products and services developed with AI; share of companies using AI in product and service development; and patents granted and R&D productivity in AI-enabled fields such as pharmaceuticals and materials science.

ECONOMIC OUTCOMES

New demand

AI can create new economic activity in addition to making existing work more efficient.

$30T

Approximate value of e-commerce sales in 2024—an enormous market enabled by an earlier general-purpose technology, the internet.

What to know: Like earlier general-purpose technologies, AI is likely to help people create entirely new products—and perhaps even new wants and needs. It may create demand in several specific ways. It could help people find goods and services that better match what they want and could lower prices for some goods and services, including those it provides, encouraging people to consume more. If history is a guide, this new activity could drive economic growth over time.

Metrics to watch: Revenue from new AI-enabled products and services; and estimates of consumer value from AI products and services.

Source: UN Trade and Development, "E-commerce: Domestic and international sales" (dataset), last updated May 21, 2026; World Bank, World Development Indicators, 2026

ECONOMIC OUTCOMES

Jobs

AI will reduce the need for some jobs while creating and changing many others.

72%

By 2035, 72% of US workers could see AI reinvent how they spend more than 15% of their working time.

What to know: Most jobs combine tasks that AI can perform with others it cannot. In the near future, jobs are more likely to change than to be replaced outright. As AI takes on more routine work, the skills workers need will change, potentially increasing the value of judgment, creativity, and working effectively with AI.

Metrics to watch: Employment and wage growth in AI-exposed occupations; the share of occupations requiring AI skills; and occupational transition rates.

Source: McKinsey Global Institute analysis

ECONOMIC OUTCOMES

Time freed

AI can free time from paid and unpaid work, but people decide how to use it.

30%

Nearly one-third of US workers who used AI at work in the previous week said it saved them three or more hours; an additional 31% reported saving one to two hours.

What to know: Time freed by AI could be used for more paid work, learning, entrepreneurship, leisure, or other activities. How people use that time will help determine AI's effects on economic growth and well-being.

Metrics to watch: Time saved through AI use in paid and unpaid work; average weekly work hours; and hours per day spent in learning, leisure, and unpaid work.

Source: US Census Bureau, "AI use at work is growing—and saving workers time," August 2026

ECONOMIC OUTCOMES

Capital shifts

Capital is shifting toward the infrastructure, technology, and businesses expected to benefit from AI.

$756B

Combined capital investment and R&D spending by seven leading AI-related companies in 2025, twice as much as in 2022.p

What to know: As AI investment has grown, so have its financing needs. Funding comes from a wide range of sources, including public markets, private credit, governments, and venture capital. Capital tends to flow toward parts of the AI system that need more capacity, financing everything from chips and data centers to electricity generation.

Whether these investments ultimately pay off will depend on the returns they generate. The industries that the McKinsey Global Institute classifies as the AI foundation added about $500 billion in revenue from 2022 to 2025. They are among 18 future arenas of competition identified by MGI; between 2022 and 2024, these arenas delivered a return on invested capital about 1.5 times that of other industries.[4]

Metrics to watch: AI infrastructure investment as a share of GDP; venture, private-credit, and government investment in AI; revenue growth and return on invested capital in AI foundation industries; depreciation of AI chips by generation; and the share of AI investment going to the largest companies.

Source: McKinsey Global Institute analysis

BROADER OUTCOMES

Business landscape

AI could change how companies operate, compete, and create advantage.

1/3

About one-third of survey respondents report that their organization used agentic coding tools to build their own software instead of purchasing it.

What to know: Using the best AI models may not by itself give companies a lasting edge because the technology is widely available and leading models may not stay ahead for long. Companies may increasingly compete on proprietary data, customer relationships, and how effectively they integrate AI into their work. AI may also shift the competitive landscape by changing relationships among companies, customers, and suppliers, and by reducing some advantages of scale. As AI use expands, governance structures are needed to manage the risks.

Metrics to watch: Performance differences between AI leaders and laggards; new-business formation and average employment at entry; distribution of firms by enterprise size; shifts in gross value added and employment by industry; and governance, audit, and cybersecurity readiness across firms.

Source: Dan Tinkoff, Lieven Van der Veken, and Michael Chui, with Tara Balakrishnan, "The state of AI in 2026: On the road to ROI," McKinsey, August 25, 2026

[4] The McKinsey Global Institute defines the "AI foundation" as semiconductors, cloud services, and AI software and services. The seven companies included in the key-data calculation are Nvidia, Oracle, TSMC, Microsoft, Meta, Alphabet, and Amazon. See Kweilin Ellingrud, "The race takes off in the next big arenas of competition," McKinsey Global Institute, March 26, 2026.

BROADER OUTCOMES

Geopolitics

AI is changing geopolitical power as countries compete for capabilities and the resources needed to develop them.

55% vs 23%

Over half of the notable AI models released since 2022 were developed in the United States, compared with 23% in China.

What to know: The United States and China are competing for an AI advantage, each with different strengths. Both countries are also using policy tools to gain leverage, from US restrictions on advanced-chip exports to Chinese controls on critical-mineral exports, with consequences for global trade, alliances, and national security. There is no single measure of AI leadership—the US may lead in frontier models, while China may be ahead in manufacturing deployment—but the technology is becoming an important source of economic and military power.

Metrics to watch: Notable AI model performance by country; data center capacity by country or region; cross-border model downloads and usage; AI hardware trade flows; and trade restrictions on chips, minerals, and other AI-related technologies.

Source: Epoch AI; McKinsey Global Institute analysis

BROADER OUTCOMES

People, society, and well-being

AI is changing how people think and interact with the world around them.

73%

In June 2025, personal queries made up nearly three-quarters of consumer ChatGPT interactions, compared with about half in June 2024.

What to know: AI could change how people think, learn, work, and relate to one another, with consequences for relationships, mental well-being, and how people spend their time. It could spur creativity and enjoyment, but work that requires more oversight, judgment, and context-switching could also contribute to cognitive overload and burnout.[5] These experiences will shape people's views of the technology, influencing how businesses use it and governments regulate it.

Metrics to watch: AI use for health, professional, relationship, and financial guidance; employee burnout and well-being measures; and surveys of public trust, concern, and excitement about AI.

Source: Aaron Chatterji et al., "How people use ChatGPT," National Bureau of Economic Research working paper 34255, September 2025

Risks

AI poses risks in many places across the system.

What to know: In our framework, risks are potential harms, distinct from bottlenecks that can slow AI's development and adoption. They can arise from the resources AI requires, the models themselves, how people and organizations use them, and the broader changes AI produces—from pressure on communities and the environment to model errors, cyberattacks, and fraud.

Metrics to watch: Model errors and safety failures; cyberattacks and fraud involving AI; and environmental and community impacts.

Feedback across the system

Changes in one part of the AI system can feed back into others, amplifying or slowing change.

What to know: The system does not move in one direction. Economic gains can fund more R&D and infrastructure, accelerating progress. At the same time, people's experiences with AI can change business behavior and government policy, which can speed or slow adoption.

[5] *People & Organization Blog*, "Five principles for designing brain-powered organizations," blog post by Jacqueline Brassey, Julie Goran, Alexis Krivkovich, and Kate Smaje, McKinsey, July 13, 2026.

Acknowledgments

This research was led by **Shubham Singhal,** McKinsey Global Institute (MGI) chairman and a senior partner based in McKinsey's Detroit office; **Arvind Govindarajan,** an MGI partner in the Boston office; **Tanguy Catlin,** an MGI director and senior partner in the Boston office; **Maria Jesus Ramirez,** an MGI senior fellow in the Bay Area office; and **Xavier Rigby,** a consultant in the Austin office.

Many McKinsey colleagues and alumni also participated in this research. Camillo Lamanna and Xavier Rigby led the working team. Team members were Claire Cheng, Evelyn Howe, Kene Okoye, Samantha Swartz, Stephanie Lee, and Tyler Berretta. Former McKinsey colleague Anu Madgavkar also helped shape this research.

Many McKinsey colleagues provided valuable input and support. We wish to thank colleagues from the McKinsey Global Institute, including Chris Bradley, Diego Castresana, Fabio Neves, Jan Mischke, Kevin Russell, Marc Canal, Mekala Krishnan, Natalya Goryunova, Qianyi Khor, Suhayl Chettih, and Tiago Devesa. We are also grateful to McKinsey colleagues Aamer Baig, Carolyn Dewar, Clarisse Magnin, Fransje van der Marel, Gemma D'Auria, Kelsey Robinson, Lucy Perez, Marco Piccitto, Maya Horii, Naveen Sastry, Pankaj Sachdeva, Shu Chern Lim, Tarek Mansour, and Vinayak HV.

We also thank our MGI colleagues David Batcheck, social media editor; Mark Staples, editorial director; Nienke Beuwer, Rebeca Robboy, and Suzanne Albert in external communications; Rachel Robinson, director of publishing operations; and Rishabh Chaturvedi, assistant managing editor.

The report was edited and producing by MGI senior editor Jason Clenfield. Chuck Burke designed the report's data visualizations.

This report contributes to McKinsey's ongoing research on AI and aims to help leaders understand the forces transforming work, identify strategic priorities, and prepare for the next wave of growth. As with all MGI research, this work is independent and was not commissioned or sponsored by any business, government, or other institution. We seek to provide business and policy leaders with the facts and insights needed to better understand the forces shaping the global economy. Any errors are our own.

Endnotes

1 The forecast is Anthropic CEO Dario Amodei's; see Jim VandeHei and Mike Allen, "Behind the Curtain: A white-collar bloodbath," Axios, May 28, 2025.

2 On overinvestment and stranded capacity in earlier infrastructure booms, see Carlota Perez, *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*, Edward Elgar, 2002.

3 See, for example, Airlie Hilliard et al., "Artificial intelligence policy worldwide: A comparative analysis," *Royal Society Open Science*, 2026, Volume 13, Number 2.

4 Paul A. David, "The dynamo and the computer: An historical perspective on the modern productivity paradox," *American Economic Review*, 1990, Volume 80, Number 2.

5 Erik Brynjolfsson, Daniel Rock, and Chad Syverson, "The productivity J-curve: How intangibles complement general purpose technologies," *American Economic Journal: Macroeconomics*, January 2021, Volume 13, Number 1.

6 Thomas Kwa et al., "Measuring AI ability to complete long software tasks," METR, March 19, 2025; and "Time horizon 1.1," METR, January 29, 2026.

7 *Energy and AI*, International Energy Agency (IEA), April 10, 2025.

8 As noted by C. C. Wei, TSMC's CEO, in the company's second-quarter earnings call, July 20, 2023.

9 Anton Shilov, "Wait times for Nvidia's AI GPUs ease to three to four months, suggesting peak in near-term growth—the wait list for an H100 was previously eleven months: UBS," Tom's Hardware, February 16, 2024.

10 "TSMC announces the opening of advanced backend fab 6," TSMC, June 8, 2023; and C. C. Wei in the TSMC fourth-quarter 2025 earnings call, January 15, 2026.

11 "Microsoft fiscal year 2026 first quarter earnings conference call," Microsoft, October 29, 2025; and Charlotte Trueman, "Microsoft has AI GPUs 'sitting in inventory' because it lacks the power necessary to install them," Data Center Dynamics, November 3, 2025.

12 Joseph Rand et al., *Queued up: 2026 edition—Characteristics of power plants seeking transmission interconnection as of the end of 2025*, Lawrence Berkeley National Laboratory, May 2026. The wait has grown so long that some companies are turning to on-site gas turbines instead. Among the largest new AI data center campuses, those of 500 megawatts or more, more than a quarter are expected to generate their own power this way by 2030, up from about 1 percent today. See *Perspectives*, "Turbines at the fence line: How gas power is fueling the AI data center race," blog entry by Mallory G. Wynne, Jones Walker LLP, February 27, 2026.

13 "Colocation data centers: The infrastructure race behind AI," McKinsey Global Institute, June 30, 2026.

14 Ben Cottier et al., "LLM inference prices have fallen rapidly but unequally across tasks," Epoch AI, March 12, 2025; and *Energy and AI*, IEA, April 10, 2025.

15 "The cost of compute: A $7 trillion race to scale data centers," *McKinsey Quarterly*, April 28, 2025; Andrew Batson, *North America data center report year-end 2025*, JLL, February 17, 2026; and Joseph Rand et al., *Queued up: 2026 edition—Characteristics of power plants seeking transmission interconnection as of the end of 2025*, Lawrence Berkeley National Laboratory, May 2026.

16 *Q1 2026: Data Center Watch report*, Data Center Watch, June 12, 2026.

17 "First statewide moratorium on new hyperscale data centers launched by Governor Kathy Hochul," Office of the Governor of New York, July 14, 2026; and Jeannie Salo, "Backlash against data centers could cost the US its AI edge," Atlantic Council, July 17, 2026.

18 *Global critical minerals outlook 2026*, International Energy Agency, July 2026.

19 Drew Robb, "Are there enough engineers for the AI boom?," IEEE Spectrum, January 17, 2026.

20 "Task-completion time horizons of frontier AI models," METR, updated May 8, 2026.

21 "Agents, robots, and us: Skill partnerships in the age of AI," McKinsey Global Institute, November 25, 2025; "Agents, robots, and us: How AI reshapes work and skills in Europe," McKinsey Global Institute, May 12, 2026; "Agents, robots, and us: How AI reshapes work and skills in Latin America," McKinsey Global Institute, July 9, 2026.

22 Dan Tinkoff, Lieven Van der Veken, and Michael Chui, with Tara Balakrish, "The state of AI in 2026: On the road to ROI," McKinsey, August 25, 2026.

23 *Tech: Forward*, "State of AI trust in 2026: Shifting to the agentic era," blog entry by Gabriel Morgan Asaftei, Roger Roberts, Abby Sticha, and Cécile Prinsen, McKinsey, March 25, 2026. Maturity was assessed on a four-level scale across five dimensions, including agentic AI governance and control.

24 "2026 AI Impact Survey Report," Grant Thornton, 2026. Based on live interviews with 950 business leaders across ten industries, conducted February 13–March 18, 2026; the enterprise AI strategy figure reflects operations business leader respondents.

25 "Global AI confessions report: CEO edition 2026," joint report from Dataiku and The Harris Poll, May 2026.

26 *Global cybersecurity outlook 2026*, a report from World Economic Forum in collaboration with Accenture, January 2026. Roughly one in three organizations still have no process to assess the security of AI tools before deploying them; the share that do has risen to 64 percent, from 37 percent a year earlier.

27 David Kertai, *Improving state and local government cybersecurity*, Information Technology and Innovation Foundation, April 2026.

28 "Our evaluation of Claude Mythos Preview's cyber capabilities," UK AI Security Institute, April 13, 2026.

29 "Research and development," in *2026 AI Index report*, Stanford Institute for Human-Centered Artificial Intelligence, April 2026. The United States produced 59 of the notable AI models identified for 2025, China 35, and South Korea eight.

30 "Beyond compute: Infrastructure that powers and cools AI data centers," McKinsey, October 29, 2025.

31 "China's new rare-earth and magnet restrictions threaten U.S. defense supply chains," Center for Strategic and International Studies, October 9, 2025.

32 Dan Tinkoff, Lieven Van der Veken, and Michael Chui, with Tara Balakrishnan, "The state of AI in 2026: On the road to ROI," McKinsey, August 25, 2026. In 2026, 54 percent of respondents at organizations with at least $1 billion in annual revenue said AI was scaling across the enterprise, compared with only about one-third at smaller organizations.

33 World Development Indicators, World Bank, accessed September 2026.

34 Robert M. Solow, "We'd better watch out," *New York Times Book Review*, July 12, 1987.

35 "AlphaFold: Five years of impact," Google DeepMind, November 2025; and "Using generative AI, researchers design compounds that can kill drug-resistant bacteria," MIT News, August 14, 2025.

36 Amil Merchant et al., "Scaling deep learning for materials discovery," *Nature*, November 2023, Volume 624.

37 William Stanley Jevons, *The Coal Question: An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal-Mines*, Macmillan, 1865. Jevons observed that greater efficiency in the use of a resource can lower its effective cost and increase overall consumption.

38 William Stanley Jevons, *The Coal Question: An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal-Mines*, Macmillan, 1865. Jevons observed that greater efficiency in the use of a resource can lower its effective cost and increase overall consumption.

39 The Henry Ford Staff, "Henry Ford quotations," The Henry Ford, December 15, 2025.

40 McKinsey Global Institute analysis, forthcoming.

41 "Agents, robots, and us: Skill partnerships in the age of AI," McKinsey Global Institute, November 25, 2025.

42 Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen, "Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence," Stanford Digital Economy Lab, August 2026.

43 AI assistance increased customer support productivity by 15 percent, with the largest gains among less experienced and lower-skill workers; see Erik Brynjolfsson, Danielle Li, and Lindsey Raymond, "Generative AI at work," *Quarterly Journal of Economics*, May 2025, Volume 140, Number 2.

44 McKinsey Global Institute analysis, forthcoming.

45 "The cost of compute: A $7 trillion race to scale data centers," McKinsey, April 28, 2025. Projected investment includes about $3.1 trillion for chips and other hardware, $1.3 trillion for power, cooling, and electrical equipment, $800 billion for construction, and $1.5 trillion for non-AI data centers.

46 McKinsey Global Institute analysis of data from the US Census Bureau through June 2026. For the definition of AI-related goods and a previous examination of AI-related trade, see "Geopolitics and the geometry of trade: 2026 update," McKinsey Global Institute, March 19, 2026.

47 "Catalyzing competitiveness: Where investment happens and why," McKinsey Global Institute, June 30, 2026.

48 Carlota Perez, *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*, Edward Elgar, 2002.

49 Luz Ding, Spe Chen, and Hayley Warren, "US lead in the AI race with China is rapidly narrowing," Bloomberg, August 20, 2026.

50 This risk to critical thinking is among the risks considered in Yoshua Bengio et al., *International AI safety report 2026*, UK Department for Science, Innovation, and Technology, February 3, 2026.

51 "The human advantage: Stronger brains in the age of AI," McKinsey Health Institute, January 15, 2026.

52 Jeffrey Gottfried et al., *Americans and AI 2026: Chatbots, smart devices and views on impact*, Pew Research Center, June 17, 2026.

53 "Ipsos AI Monitor 2026," Ipsos, June 2026. In a survey of 23,532 adults across 32 countries, 55 percent of respondents in China and 54 percent in India expected increased AI use to improve the job market over the next three to five years, compared with 18 percent in the United States.

54 Saffron Huang et al., "What 81,000 people want from AI," Anthropic, March 18, 2026.

55 Antoine Montard, Dago Diedrich, and Tanguy Catlin, "Where AI will create value—and where it won't," McKinsey, April 29, 2026.

56 Dago Diedrich, Evan Williams, Tanguy Catlin, and Tim Fountaine, "From AI table stakes to AI advantage: Building competitive moats," McKinsey, May 15, 2026.

57 Jan Mischke et al., "The power of one: How standout firms grow national productivity," McKinsey Global Institute, May 6, 2025. The analysis covered about 8,300 large firms in Germany, the United Kingdom, and the United States across four sectors from 2011 to 2019.

58 For discussion of a FINRA-style self-regulatory body for AI, see Vinh Nguyen, Elham Tabassi, and Kat Duffy, "The U.S. is about to design an AI regulator. Here's how to get it right," Council on Foreign Relations, July 23, 2026; Mark Thomas, "Designing a FINRA for frontier AI," *Lawfare*, July 30, 2026.

59 "Tech Secretary launches Sovereign AI," UK Department for Science, Innovation and Technology, April 16, 2026; "EU launches InvestAI initiative to mobilise €200 billion of investment in artificial intelligence," European Commission, February 11, 2025.

Copyright © 2026 McKinsey & Company. All rights reserved.