July 2026

Global Macro Insights

Artificial intelligence and the global economy: A framework for investors

Edoardo Cilla, PhD
Director, Senior Strategist

Salman Ahmed, PhD
Global Head of Macro & SAA

Tony Lee
Associate Director Product & Content Marketing

Fidelity INTERNATIONAL

Contents

Executive summary

- At Fidelity International, we are approaching the artificial intelligence (AI) revolution through an integrated framework that combines bottom-up company and sector evidence with top-down macro analysis. The bottom-up work, including insights from our Fidelity Analyst Survey, helps us understand how individual firms are adopting AI and where productivity gains are appearing.
- This paper presents the top-down pillar of that framework, which is designed to assess how AI may affect macroeconomic variables such as real GDP growth, inflation and equilibrium interest rates by looking across the entire AI value chain, from energy and critical minerals to semiconductors, data, models and real-world applications.
- As uncertainty is unusually high, this pillar of our framework is scenario-based rather than providing a fixed forecast. We assess both the strengths and weak links that determine whether countries can produce AI models and whether they can exploit it broadly enough to generate economic growth. This includes financing, energy security, talent, sector structure, workforce adaptability, regulation and environmental constraints. Geopolitical forces, including recent export controls on AI models and advanced chips, and inequality are the key underlying forces shaping both AI production and diffusion.
- We will continue to update our assumptions as new evidence emerges. We will also extend our analysis in a forthcoming paper, focusing on how AI could impact capital market returns through the channels of real rates and corporate earnings, primarily through its impact on labour costs, interest expenses and corporate taxes.

Key findings (base case)

- AI should lift real GDP growth in our base case, but gradually rather than explosively, through stronger productivity, faster innovation and higher investment. However, countries will not capture its full potential if weak links in areas such as energy, capital, skills, adoption or geopolitical fragmentation slow diffusion.
- On inflation, AI is likely to be inflationary before it becomes disinflationary, leaving the net effect broadly neutral in our base case: investment in electricity, data centres and chips can create bottlenecks and add to inflation volatility, while later productivity gains should lower costs. AI should also put modest upward pressure on the equilibrium real rate, R*, with productivity and investment demand acting as upward forces, while aging, inequality and precautionary saving acting as offsets.
- AI's macro impact remains scenario-dependent: real GDP growth, inflation and R* could diverge materially across our low- and high-AI cases. We will therefore monitor bottom-up company and sector evidence alongside top-down macro data to judge whether probabilities are shifting between high-, mid- and low-AI scenarios.
- Each major economy has distinctive strengths and untapped potential, and while the U.S. and China are ahead today, no country is independent. The U.S. has models, capital and cloud; China has scale and industrial depth; Europe has savings and science; Asian countries control robotics, minerals, memory and fabrication.
- AI has constructive implications for investors. We have already incorporated the implications of higher GDP growth into our latest CMAs through stronger aggregate corporate revenues, while AI-related dynamics and other structural forces should keep rates above the ultra-low post-GFC, pre-Covid regime. Bottom-up, dispersion will be significant: aggregate implementation remains early and some valuations already embed high expectations, but the AI value chain extends beyond the most visible technology leaders. Investors should avoid treating AI as only a U.S. equity, chip or technology-sector story, and instead combine selective diversification, valuation discipline and forward-looking analysis of firms able to turn AI into competitive advantage and durable earnings growth.

Our scenario-based framework for assessing the AI revolution

Salman Ahmed, PhD
Global Head of Macro and Strategic Asset Allocation

Various countries, sectors and social groups will experience the AI transition differently, so it is prudent to acknowledge that a wide range of potential macroeconomic outcomes are possible. As such, we established a scenario-based framework for assessing the uncertainty that surrounds AI's macroeconomic effects in our H1 2026 Capital Assumptions, initially setting out three potential scenarios: High-AI, Low-AI and a base case, each of which would result in different assumptions across asset classes. This paper represents the next step in our ongoing and evolving process. It intends to tell us which countries will benefit from AI and how, by establishing a forward-looking, two-pillar framework for assessing its top-down impacts across real GDP growth, inflation and rates:

1. Can a country build AI with resilient strategic autonomy?: This depends on whether a country can mobilize deep public and private capital, secure abundant energy and strategic inputs, access usable data and model infrastructure, and attract the talent and entrepreneurs that convert these into frontier capability. In a fragmented world, geoeconomics determines whether inputs remain available and affordable, or whether they become chokepoints in a wider contest over technology and supply chains.
2. Can a country exploit AI for structural economic growth?: The real test is whether AI reaches everyday firms, workers and public services. This depends on factors like political trends, economic structure, workforce adaptability, environmental impact and governance, as well as whether AI gains diffuse broadly or concentrate among certain countries, sectors and firms.

Our base case, High- and Low-AI scenarios should be viewed as different combinations of production and diffusion capacity. The distributional channel is central: the more unequal the transition, the weaker the demand multiplier, social consent and political durability of AI-led growth. This risk is amplified by the long decline in labour's share of income; in the U.S., this has fallen from 60-65% before 2000 to roughly 55% in recent years (Figure 1). If AI mainly raises profits, capital rents and returns to highly skilled workers, it would reinforce this trend, leaving household demand weaker and raising the likelihood of political backlash.

Figure 1: The U.S. labour share decline

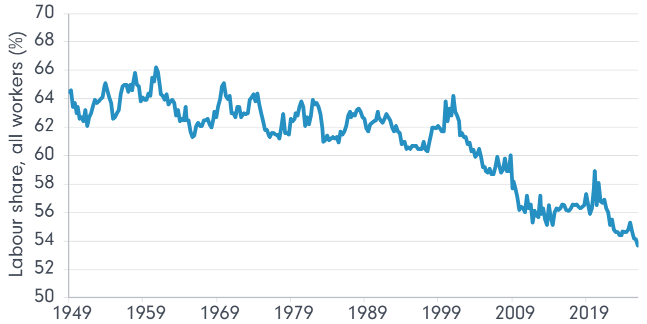

Source: Bureau of Labor Statistics, Fidelity International, 30 June 2026.

In our broader investment framework, we combine the findings of our top-down analysis with bottom-up insights, such as the findings of our Analyst Survey. Bottom-up analysis is essential as it helps inform our top-down assumptions, while top-down forecasts provide the context for framing bottom-up investment theses. It is the iterative accumulation and integration of these multiple perspectives across industries, geographies and asset classes which allows us to assess the most complex investment themes.

The assumptions and findings presented in this paper should therefore be treated as a living framework, rather than a fixed forecasts. We will continue to update them as more evidence becomes available on AI adoption, investment, productivity, employment effects and policy responses, acknowledging that noise surrounding these variables is likely to remain high. The next phase of this work will focus more directly on capital market effects through analysis of competitive dynamics, labour market impacts and potential policy responses (e.g. tax adjustments) on corporate profit margins and asset valuations.

Figure 2: Our forward-looking, two-pillar framework for assessing AI's top-down impacts

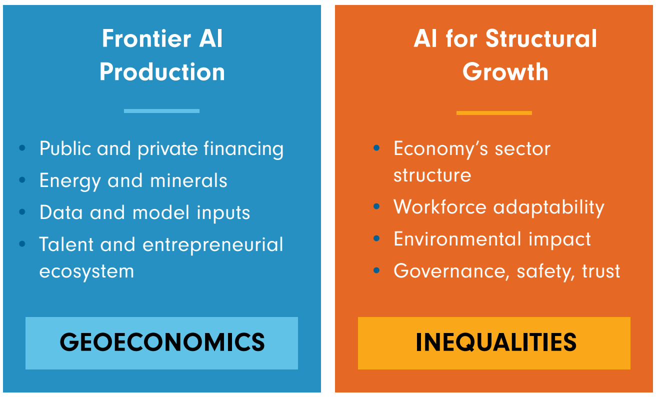

A multidimensional problem

"Periods such as the present are rare because they combine a rupture in the prevailing macroeconomic regime with the emergence of a technology capable of reshaping the foundations of our economy and society."

Edoardo Cilla, PhD
Director, Senior Strategist

In the years before 2020, investors and companies had become accustomed to relatively stable rules, deep global integration, low inflation and real interest rates, and the belief that economic interdependence would help contain geopolitical tensions. Many of those assumptions now look less secure. Trade has become more politicized, fiscal policy is increasingly constrained by high public debt, energy security has returned as a strategic priority and global cooperation has become more difficult. This is all occurring precisely as AI has begun to require more scale, trust and shared infrastructure.

AI has also arrived at a time when the global economy is already under stress on various fronts, so its impact can only be understood through how it interacts with the other key macro forces defining our time:

- **Demographics:** We are in the middle of an unprecedented period of rising old-age dependency. AI-augmented robotics may help raise the productivity of scarce workers, while existing occupations may be disrupted.
- **Public debt** is high, yet AI requires large investment in data centres, power grids, chips, skills and public services.
- **Inequality** has increased in many countries in recent decades and early concerns that AI will exacerbate this show how quickly technological excitement can translate into political resistance if the benefits flow narrowly to capital, platforms and highly skilled workers.
- **Climate** is another important variable, as AI requires large amounts of electricity and water, even though it may also improve grid management, materials discovery and energy efficiency.
- **Geopolitical fragmentation:** AI has developed inside a deeply globalized economy, and its value chain is truly global. However, governments are now trying to localize or restrict the inputs that make AI possible through export controls and technology screening. The Fable 5 export-control episode and a possible 5% U.S. government stake in OpenAI show AI is moving into state power and global diplomacy.

Figure 3: The AI revolution is not just about hardware or software

Investors' attention has been rotating across the AI stack over the recent years.

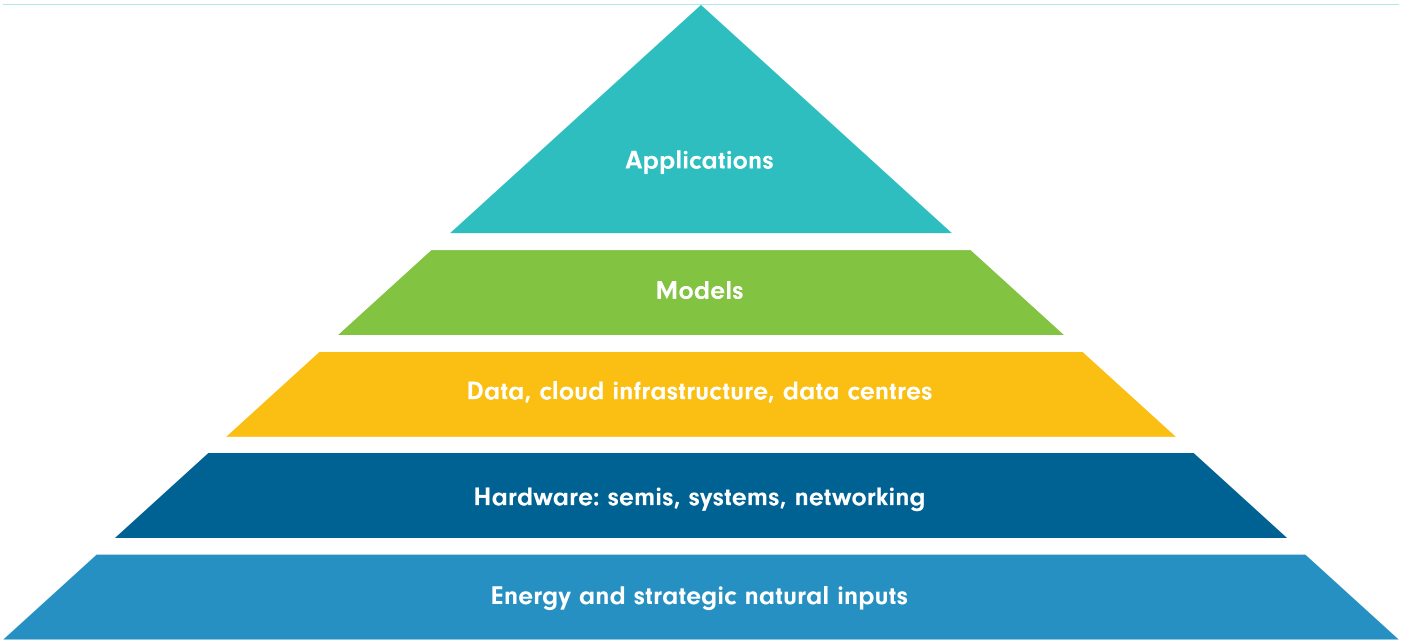

For illustrative purposes only.
Source: Fidelity International, July 2026.

Is this time different?

AI differs from previous breakthrough technologies in three fundamental ways:

1. It can act with a degree of **autonomy**, which means it can perform tasks, make recommendations and interact with other systems, performing cognitive tasks that were once considered uniquely human.
2. **Adoption** is much faster because software spreads rapidly through cloud platforms, enterprise tools and consumer applications.
3. It can contribute to rapid **idea generation**, which makes it relevant not only for executing existing processes more efficiently, but also for research, innovation and the creation of new products.

These factors make the AI revolution potentially more powerful than previous waves of innovation, but also make its economic effects harder to gauge through historical comparisons alone. A forward-looking approach is therefore essential for assessing AI's impacts across economies, from labour markets and health care to defence, industrial policy and income distribution within and between countries.

Moreover, AI is more than a chip story, as its value chain includes energy and natural inputs, semiconductor hardware, cloud infrastructure, data, model and, finaly, applications. It is not just a U.S. story. The semiconductor chain alone links U.S. chip design, European lithography equipment, Korean high-bandwidth memory, Taiwanese fabrication and packaging, and Chinese strength in critical mineral processing.

The result is a highly uncertain picture which has encouraged a wide range of forecasts, from those presenting AI as the start of a new productivity golden age to those warning of labour-market disruption social instability, geopolitical conflict, and excessive concentration of power. Both perspectives capture part of the truth, although neither is sufficient on its own.

Key findings: Macroeconomic impacts

Isolating the macroeconomic impacts of AI is difficult, as it has arrived at a time of rapid change driven by diverse factors like post-Covid adjustment, fiscal strain, demographic change, geopolitical fragmentation, protectionism and energy shocks. Nonetheless, our framework is designed to help us simplify complexity, with country-level analysis allowing us to gauge the likely effects of AI on key macroeconomic variables across jurisdictions:

- Real GDP growth
- Inflation
- Interest rates

Our broader framework incorporates bottom-up insights from across our extensive global research platform, including through initiatives like our periodic Analyst Survey, which allow us to build a better understanding of AI's net macroeconomic effects by capturing shifts in multiple underlying variables simultaneously.

A positive macro force

Overall, in our base case, AI should be considered as a positive macro force by investors. It is likely to lift real GDP growth, which is good news for companies because a larger and more productive economy should support stronger long-term revenue growth. Nominal rates are also likely to remain higher than they were in the pre-pandemic period, because inflation, equilibrium real rates and the term premium all settle at higher levels, also thanks to AI build.

Impact on global real GDP growth: Positive but gradual

Real GDP growth depends on two forces: labour productivity, meaning how much each worker produces, and the number of workers. AI can raise the first component through various activities, including faster coding, cheaper analysis, better customer service, more efficient logistics, accelerated research and redesigned business processes. However, its effect on the second component is much less predictable.

The key point is broader than whether the effect on employment is positive or negative in aggregate; what also matters is how the impact is distributed across workers, sectors and regions. AI can complement some workers, create new activities, replace some tasks and reduce labour demand in certain industries. However, an AI boom whose benefits accrue only to a narrow group of firms, capital owners and high-skilled employees, while displacing others, creating wage pressures and weakening career progression could result in the overall economy experiencing a weaker consumption multiplier than the productivity numbers imply. Aggregate consumption, the primary driver of demand in most large economies, could even weaken.

Here, historical comparisons are useful only up to a point, as AI differs from previous breakthrough technologies. Uncertainty is therefore resulting in an extremely wide range of forecasts, some of which seem designed more to attract attention more than they should be considered serious macroeconomic analysis. Our approach to assessing these complex dynamics is to combine multiple forward-looking perspectives, pairing top-down macro evidence with bottom-up firm and sector insights. The findings take the form of a range of potential scenario outcomes which we can evolve over time as new information becomes available.

AI transition currently a capex boom, rather than a productivity boom

In the U.S., AI is already visible in the investment cycle. Real GDP grew at an annualized 1.6% in the first quarter of 2026, below its 2.0% estimate of longer-run trend growth, while business investment, particularly in AI infrastructure, was the largest contributor. That data suggests that the economy is experiencing an infrastructure and capital-expenditure boom at present, rather than an economy-wide productivity boom. This gap between adoption and realized productivity is not surprising, as a firm can buy AI tools quickly but it takes longer to redesign workflows, change incentives, train staff, integrate data and decide which decisions can safely be delegated to machines.

In particular, two important variables to track are labour productivity and total factor productivity (TFP) (Figure 4). Labour productivity normally grows faster because it also reflects capital deepening and changes in labour quality, not just efficiency gains. In the U.S., since 2022, labour productivity has risen (2.5% annualized) more than TFP (1.4% annualised), consistent with stronger investment and better tools. At the same time, the gap remains broadly in line with history, so it is not yet evidence of an abnormal productivity regime, although early signs are encouraging.

Figure 4: The case of the U.S. - total factor productivity versus labour productivity (3-year rolling growth, annualized)

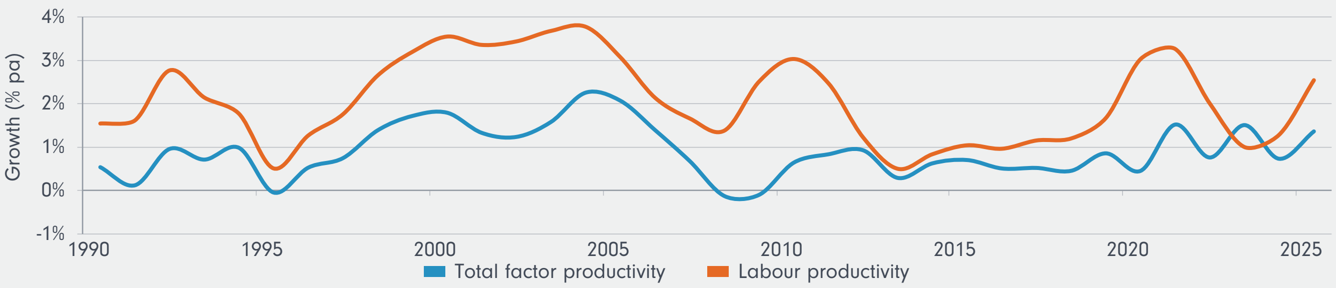

Source: FRED, Fidelity International, 30 June 2026.

Countries cannot simply free-ride on foreign AI adoption.

The OECD's "AI meets trade" paper is especially relevant because it links AI adoption to international trade, sector structure and cross-country spillovers through a micro-macro model. The authors estimate that, in their central scenario, AI could raise per capita real income growth by 0.1 to 0.95 percentage points annually across OECD and G20 economies over the next decade. The crucial lesson is that countries cannot simply free-ride on foreign AI adoption. Cheaper AI-intensive imports may help laggards, but maintaining competitiveness, especially in highly AI-exposed sectors, requires strong domestic adoption. In our broader framework, this reinforces the point that economies will struggle to exploit the full GDP potential of AI if they do not control at least part of the AI production and diffusion stack, including capital, data, skills, energy and firm-level implementation.

Productivity gains might not translate one-for-one into output.

Yotzov, Barrero, Bloom and co-authors surveyed almost 6,000 chief financial officers, chief executive officers and other executives across the U.S., the United Kingdom, Germany and Australia. Over the next three years, they expected AI to raise productivity by 1.4%, increase output by 0.8% and reduce employment by 0.7%. This is a useful microeconomic warning: even when firms expect a productivity gain, they do not expect it to translate one-for-one into output, partly because employment adjusts.

Karger and co-authors reach a similarly moderate conclusion using a different survey method. They surveyed economists, AI-company employees, AI policy researchers and highly accurate forecasters on the future effects of AI on the U.S. economy. The median respondent expected strong AI capability improvements by 2030, but still forecast annual GDP growth of about 2.5%, above conventional medium-run and long-run baselines but far from explosive levels. Even in a High-AI scenario, where AI systems exceed human performance on many cognitive and physical tasks, experts forecast annualized GDP growth rising to around 4% by 2050, a major shift but still historically recognizable.

Positive but not explosive effects (ten-year strategic investment horizon)

Base case: Moderate-GDP growth uplift AI scenario

Combining the findings of our own pillar-by-pillar analysis and those of credible external literature that leverages differentiated methodologies, our base case is that the net effect on growth is likely to be positive but gradual. AI raises global productivity growth by around 80 basis points, although this does not flow mechanically into real GDP growth. The employment channel offsets around 40 basis points of this gross productivity impulse, leaving a net positive effect on global trend real GDP growth of around 40 basis points, with significant dispersion across regions. This assumption also captures the distributional channel, as some workers and sectors benefit more quickly than others. Nevertheless, this is good news for our capital market assumptions overall, as higher real GDP growth translates to higher corporate revenue growth in our models.

High-GDP growth AI scenario

AI adds around 1.5 percentage points to trend growth, which would be materially positive but still not an explosive take-off. That outcome would require faster diffusion of AI agents, stronger scientific and business-process innovation, broader adoption beyond large technology firms and a smoother employment transition than we currently assume. Labour disruption would still exist, but new task creation and better labour-market adaptation would keep the negative impact on employment much lower, at around 0 to 10 bps, allowing most of the productivity impulse to pass through to GDP.

Low-GDP growth AI scenario

AI adds essentially nothing to trend growth relative to our existing long-term baseline: it helps economies defend a roughly 2% long-term global growth rate, but headwinds associated with energy bottlenecks, regulation, trust, labour-market adaptation and firm implementation prevent an additional growth dividend. This cautious case echoes Acemoglu's warning that AI may deliver only modest productivity gains if it is used mainly to replace existing labour tasks rather than to help workers perform new tasks more effectively, in which case growth could disappoint while labour substitution fuels stronger social and political backlash. In this scenario, the labour-market drag would be much larger, offsetting much or all of the productivity gains.

Geographic GDP growth differences

- Better-positioned economies, especially the U.S. and China, will receive a larger boost of roughly 55 basis points in our base case. The U.S. benefits from frontier models, cloud infrastructure, venture capital, large technology firms and a service-heavy economy. China benefits from scale, engineering depth, manufacturing diffusion and policy support, although chip restrictions and capital allocation risks matter.
- Less-exposed developed markets will receive closer to 30 basis points. Europe, Japan and other developed markets can still gain meaningfully, but the dividend is likely to be smaller unless adoption accelerates and domestic AI ecosystems strengthen.

Impact on global inflation: Inflationary before disinflationary

To understand the impact of AI on inflation, it is important to place it inside a wider macro regime shaped by three key forces:

1. Geopolitical fragmentation
2. Politicization of economic policy, of which fiscal dominance is an important component
3. The AI revolution itself

These forces should be analyzed together because they are mutually reinforcing. A useful starting point is a hybrid New Keynesian Phillips Curve, where the level of inflation is dependent on past inflation, inflation expectations, the output gap and cost or supply shocks (Figure 5).

For example, AI is becoming a geopolitical weapon, since leadership in chips, models, cloud infrastructure and data can influence military capability, industrial competitiveness and strategic dependence. Meanwhile, geopolitical fragmentation turns AI inputs into contested assets, and political pressure encourages deeper state intervention in markets, pushing governments towards more active industrial policy. Politicization of economic policy is mainly inflationary because fiscal support, industrial subsidies, tariffs, migration restrictions and pressure on central banks can keep demand stronger, costs higher and expectations less firmly anchored. Over our ten-year strategic investment horizon, fragmentation and politicization will likely therefore put net upward pressure on inflation by raising costs, shocks and persistence.

Temporal dynamics of the AI impact

The effect of AI is likely to be more mixed as it has the most two-sided effects, with lags operating in opposite directions:

- **Short run:** AI can raise inflation by increasing demand for chips, data centres, electricity grids, cooling systems, construction, specialized labour and financing. So far, it has generated some demand pressure and created local bottlenecks, but recent inflation surprises have still been driven more visibly by energy, commodity and supply-side shocks. The scale of investments leads us to believe that the effects will become increasingly visible over the next few years.
- **Longer run:** AI can become disinflationary through productivity gains, but their distribution among businesses, workers and consumers will be crucial in determining the magnitude of the impact. Three things are important to consider:
  1. **Margins:** If companies use AI mainly to widen margins, strengthen market power or invest in more capacity, the pass-through to consumer prices may be slow. Conversely, if competition forces firms to pass cost savings on to customers, the disinflationary effect will be stronger.
  2. **Inequality:** Distribution matters because AI gains may accrue disproportionately to capital owners and high-income households, shifting purchasing power towards groups with higher saving and investment propensities. If AI reinforces this trend, it could dampen broad consumer demand, limiting inflation in products but favouring inflation in asset prices.
  3. **Labour-market effects:** These are equally ambiguous, as AI may reduce wage pressure if it substitutes for labour, but it may also raise demand for scarce technical skills and complement workers in sectors where human judgement remains central.

Figure 5: Different forces affect different components of inflation to varying degrees

	Inflation expectations	Output gap / demand pressure	Cost / supply shocks	Net effect
Geopolitical fragmentation	Medium ↑	Medium ↑ / ↓	High ↑	Upward
Politicization of economic policy	High ↑	High ↑	Medium to High ↑	Upward
AI revolution	Low to Medium ↑ / ↓	High ↑ / ↓	High ↑ / ↓	Mixed

For illustrative purposes only.
Source: Fidelity International, July 2026.

Inflationary before disinflationary (ten-year strategic investment horizon)

Leveraging this framework, we built three inflation scenarios by varying the assumptions on the main inflation drivers, especially the output gap and the cost-shock channels. In all three scenarios, AI first adds to inflation and later subtracts from it, but the timing and magnitude of these two lags determine whether the cumulative ten-year effect is positive, negative or broadly neutral. We will continue to update our assumptions as new evidence becomes available, not underestimating the potential for one effect to eventually dominate the other.

Base case AI scenario

The net effect is close to zero, because the inflationary impulse from investment, energy demand and bottlenecks is gradually offset by productivity gains, automation and lower unit costs. However, even if the cumulative effect is neutral, AI is still likely to increase inflation volatility, which supports our view that term premia should remain higher than in the post-global-financial-crisis and pre-pandemic period.

High-inflation AI scenario

AI investment proves larger and more persistent than expected, power-grid and data-centre constraints bind for longer, chip and construction bottlenecks remain material, industrial policy supports demand, and firms retain a large share of AI-related efficiency gains through higher margins rather than passing them on to consumers.

Low-inflation AI scenario

AI diffuses more rapidly across services and manufacturing, labour displacement is larger, competition forces firms to pass cost savings on to customers, and cheaper inference reduces the cost of AI deployment. Even in this scenario, the disinflationary effect is unlikely to appear immediately, as firms first need to reorganize workflows, invest in complementary systems and translate technical capability into measurable productivity gains (Figure 6).

Figure 6: Scenarios of the marginal impact of AI on inflation

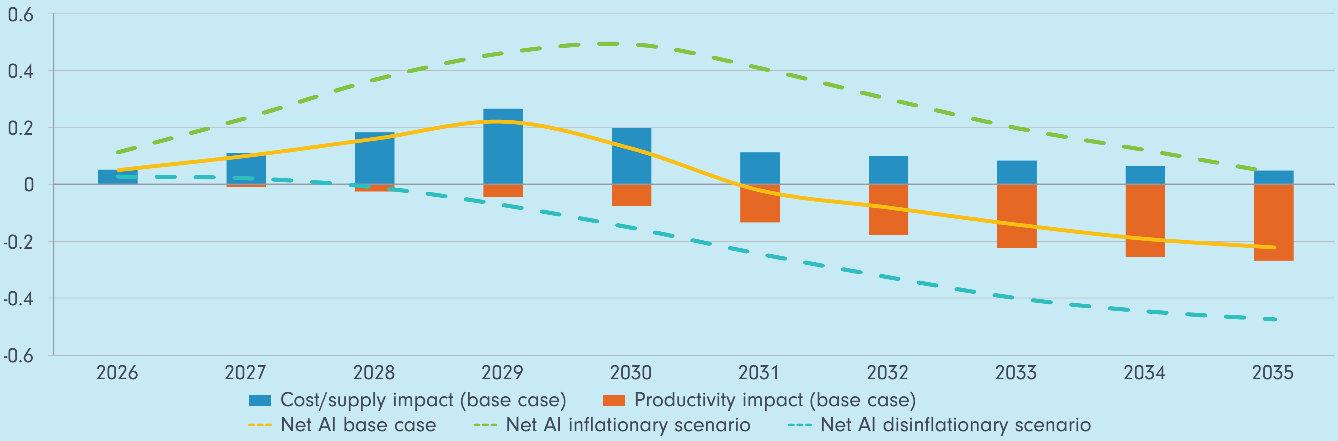

For illustrative purposes only. Assumptions are based on proprietary modelling, reflecting the views of investment professionals at Fidelity International.
Source: Fidelity International, June 2026.

Geographic differences

- The **U.S.** should have the largest short-run inflationary impulse because it is the centre of AI investment, cloud infrastructure and data-centre construction, but it also has the largest potential productivity payoff.
- The **euro area** should experience a smaller direct investment shock, although inference still requires data centres and electricity, while slower diffusion and fragmented markets may delay the productivity benefit.
- The **U.K. and U.S.** will also likely see a smaller impact than the U.S. as its services-heavy economy is highly exposed to AI productivity gains, but its infrastructure constraints could limit speed.
- **Japan** may experience a relatively more disinflationary AI channel if automation helps ease aging-related labour scarcity, especially in labour-intensive services where worker shortages have been putting upward pressure on wages and prices.
- **China** could combine state-led AI investment with manufacturing productivity and overcapacity, creating both domestic bottlenecks and global disinflationary pressure through cheaper AI-enabled goods.

Impact on global equilibrium real rates: Likely mixed

The AI transition's effect on interest rates is central to its economic and investment implications. The debate has sometimes been framed too simply, with some arguing that AI-driven productivity gains should lower inflation and therefore give central banks more scope to cut rates. Rather, we only expect this dynamic to kick in in the medium/long-term.

When it comes to rates, two things are important: R*, the (unobserved) real short-term interest rate expected to prevail when an economy is at full strength and inflation is stable, and the actual real rates experienced by companies and consumers in the economy.

It is important to stress that AI is only one of several forces shaping R* in this new macroeconomic environment: central-bank reaction functions, inflation credibility, politicization, fragmentation, communication strategy and balance-sheet policy will also play a role, because they shape expectations, risk premia, capital flows and the extent to which debt sustainability constrains monetary policy.

R* higher: Productivity, investment and credit effects

R* will likely be influenced significantly by AI, as it is structurally driven by macro variables that AI affects, such as productivity, demographics, inequality, fiscal policy, global capital flows and preferences for safe assets. Over recent decades, it has declined mainly because of demographic change and weaker productivity growth, the latter being one of the forces that now could push it higher through the AI revolution. In fact, the strongest upward channel is productivity. If AI raises trend productivity, it increases the expected return on investment and the opportunity cost of saving, which should lift R*.

AI will also likely continue to drive a large investment boom in data centres, chips, power grids, software, security and organizational redesign, boosting demand-led inflationary pressures while also raising demand for loanable funds during the adoption phase.

The R*–labour share nexus

The downward channels are equally important. If AI creates uncertainty about jobs, wages and the distribution of income, households may save more for precautionary reasons, while firms may delay investment until they understand which business models will survive. If AI is mostly capital-augmenting and income gains therefore accrue mainly to owners of capital, inequality could rise; because richer households save a higher share of income, aggregate demand may weaken more than the productivity numbers suggest.

This links directly to the nexus between R* and labour share: in the U.S., the labour share fell from around 60-65% before 2000 to about 55% in recent years, while estimates of R* fell from about 2.5-4% to around 1%. The implication is that AI raises R* more convincingly if it raises productivity in a way that supports labour income, rather than only capital rents.

Figure 7: Productivity gains may lift R* while forces behind lower labour share may suppress it

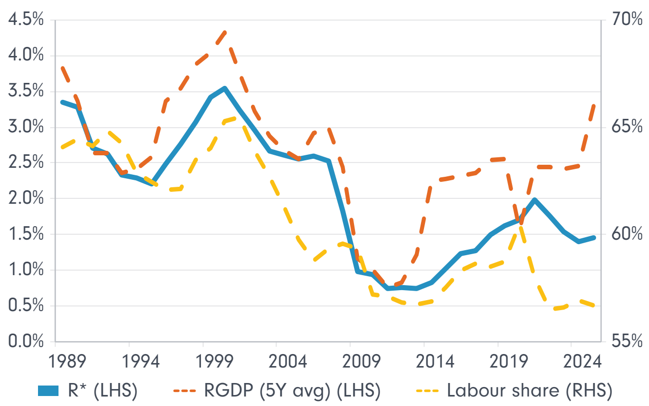

For illustrative purposes only.
Source: Fidelity International, Federal Reserve Bank of New York, U.S. Bureau of Labor Statistics, June 2026.

R* impact: Neutral to slightly positive effect (ten-year strategic investment horizon)

Our analysis of the effects of the AI transition leads us to expect a neutral to slightly positive effect on R* over our ten-year strategic investment horizon, on a net basis once all factors are taken into account. However, this could mask significant dispersion across jurisdictions, especially between AI-producing and AI-consuming economies.

Geographical effects

- In the U.S., the balance of risks points to a higher R*, at least during the investment phase, because the country combines frontier model production, hyperscale cloud, venture capital, data-centre construction, deep equity markets and a large stock of AI-related intangible assets.
- China also faces upward R* pressure from domestic AI adoption, industrial policy and compute investment, although aging, high precautionary saving and property-sector weakness may offset part of that effect.
- The euro area is more ambiguous, as if frontier AI production remains mainly concentrated in the U.S. and China, Europe may import productivity benefits through licensing and cheaper AI-enabled services, but with less domestic investment and potentially more capital flowing abroad. This is why Europe's R* story depends heavily on whether the Savings and Investments Union, capital-market integration and AI diffusion turn imported technology into domestic investment.
- The U.K. has strong AI research, finance and business services, but less compute scale than the U.S., so the productivity upside is real but the domestic capex impulse is smaller.
- Japan is the clearest case where AI may prevent R* from falling further, rather than pushing it sharply higher. Aging remains a structural drag, but robotics and AI can raise effective labour supply in sectors facing scarcity, especially health care, logistics, manufacturing and services.

While our next paper will examine the implications for market real rates in greater depth, our baseline is that AI is unlikely to produce a large or permanent rise in actual real rates. Productivity and investment may lift equilibrium R*, especially in AI-producing economies, but debt-sustainability pressures, aging, inequality and precautionary saving should provide offsetting downward forces, leaving real rates above the ultra-low pre-Covid regime, though not dramatically higher.

Conclusions

AI should be understood as a disruptive macro force whose impact will depend on how widely its benefits spread across economies, sectors, firms and workers. In our base case, AI is positive for long-term growth, but the path is gradual and uncertain, shaped by energy, capital, skills, governance, inequality and adoption bottlenecks. Our pillar framework shows why: producing frontier AI requires strategic autonomy across capital, electricity, minerals, data, chips and talent, while exploiting AI for structural growth requires diffusion, adaptable workers, trusted governance and environmental discipline. The U.S. has models, capital and cloud; China has scale and industrial depth; Europe has savings and science; Asian countries control robotics, minerals, memory and fabrication.

For investors, implications are constructive but selective: stronger GDP should support revenues, while structural forces keep rates above the ultra-low pre-Covid regime. Yet bottom-up dispersion will be large, implementation is early and some valuations already embed high expectations, but the AI value chain extends beyond the most visible technology leaders. Investors should therefore avoid treating AI as only a U.S. equity, chip or technology-sector story, and instead combine selective diversification, valuation discipline and forward-looking analysis of firms able to turn AI into competitive advantage and durable earnings growth.

This paper sets out a living framework rather than a fixed forecast, which we will continue to update as evidence improves. Our next paper will examine the capital-market implications in greater depth, including how AI may affect real rates, valuation discipline and, in particular, whether today's expected margin gains can prove durable once competition, regulation, taxation and implementation costs are taken into account.

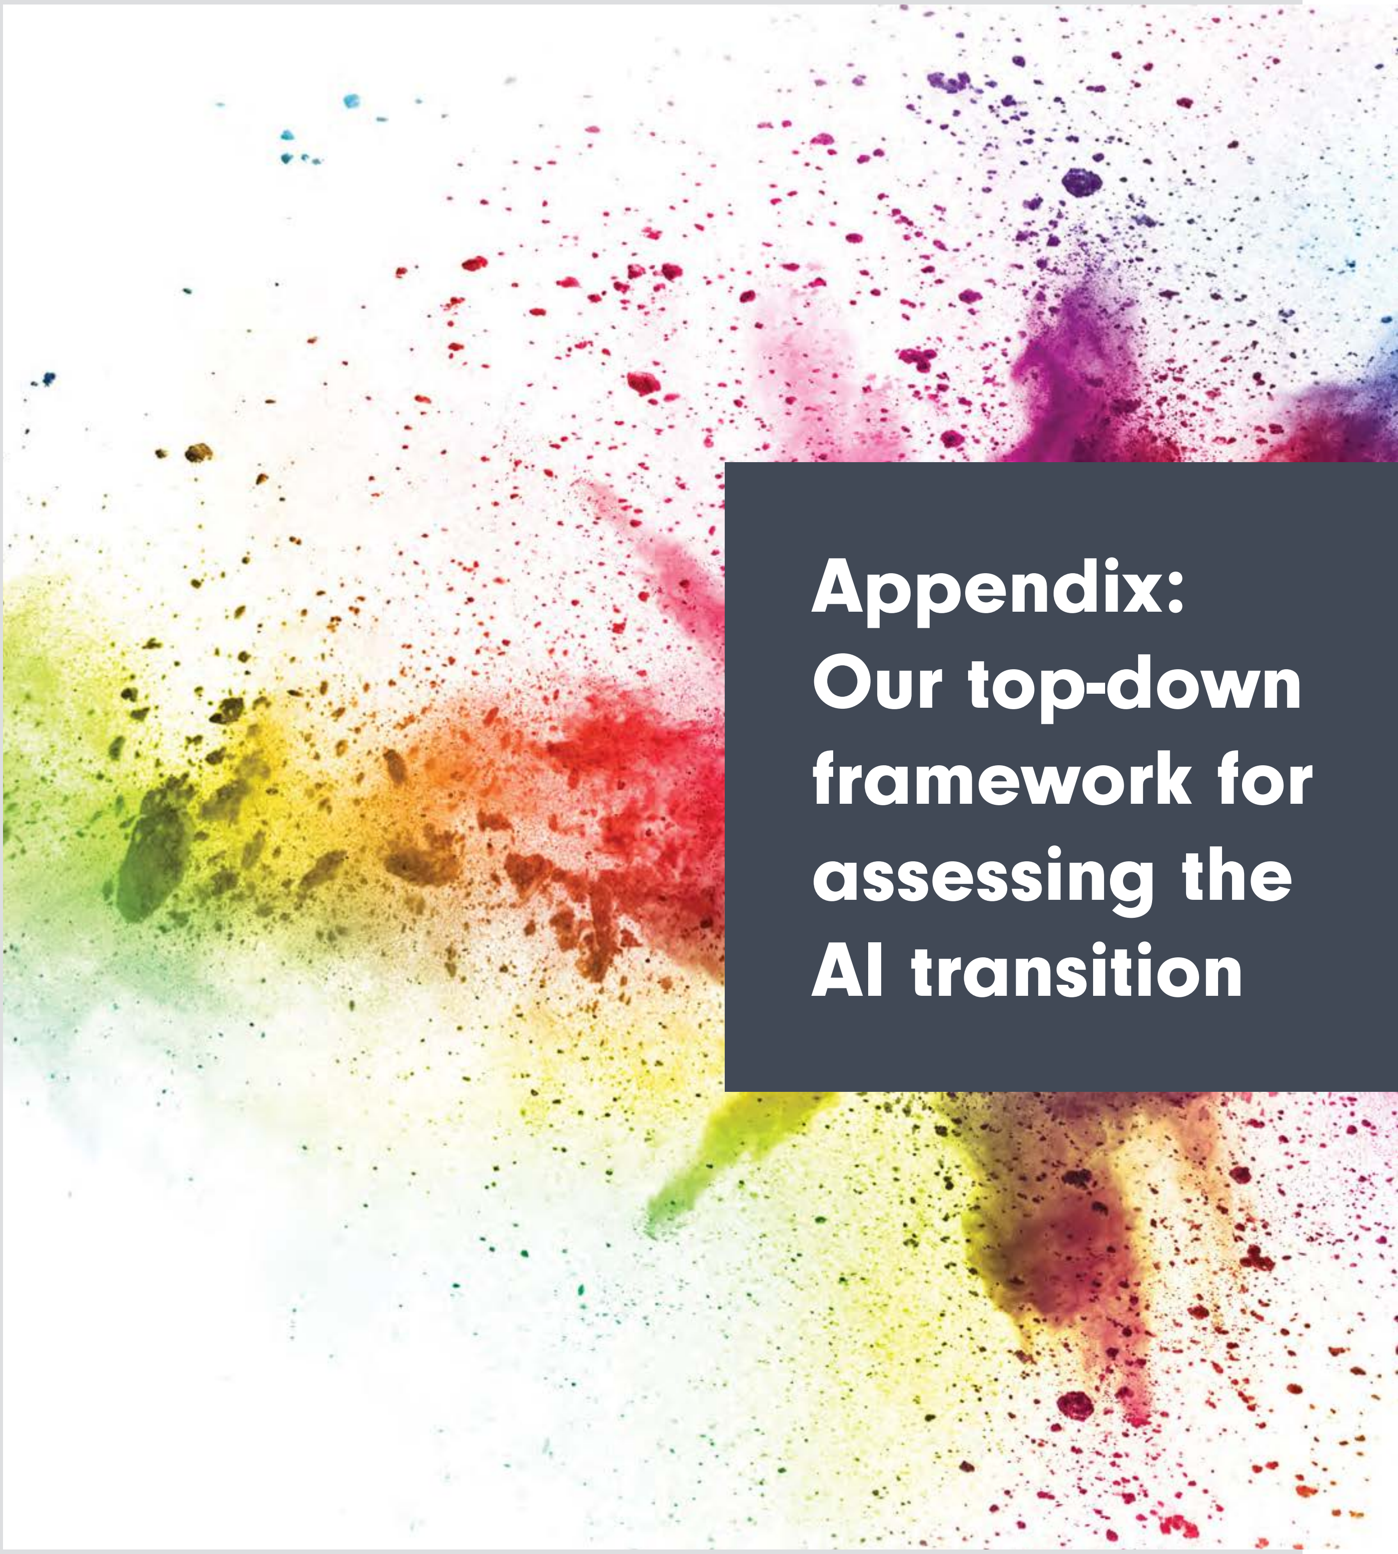

Appendix: Our top-down framework for assessing the AI transition

Pillar 1: Can a country build AI with resilient strategic autonomy?

This first pillar in our forward-looking framework for assessing AI's top-down impacts across real GDP growth, inflation and rates asks whether a country can build and sustain meaningful AI capability on its own terms. This requires access to AI inputs that is deep, reliable and secure enough to withstand a less benign external environment. The variables reflect the basic ingredients of AI production:

1. Public and private financing: Capital comes first, because public and private financing are needed to build costly systems, fund infrastructure, reshape production chains and absorb inevitable failures.
2. Energy and critical minerals: AI is physically intensive: without dependable electricity, suitable sites and key materials, compute cannot scale seriously.
3. Data and model inputs: Systems are only as useful as the information they are trained on and countries with stronger access to locally relevant, high-quality data are better placed to develop models suited to their economy, language and institutions.
4. Talent and entrepreneurial ecosystem: Talent and innovation ecosystem complete the picture, because none of these inputs translate into real capability without researchers, scientists, firms, universities and institutions able to develop the technology.

Geoeconomics runs through all these variables, as it shapes where inputs come from, the concentration of foreign dependencies, how easily they can be replaced, and how exposed a country is to supply-chain chokepoints. Over the past few years, chips and critical minerals have been the clearest examples; the U.S. has tightened controls on advanced chips, semiconductor equipment and high-bandwidth memory exports to China, while China has used restrictions on gallium, germanium, graphite and heavy rare earth elements as leverage in the wider technology contest. The recent U.S.-China summit showed that AI chips have moved from corporate supply chains into high diplomacy.

The conclusion is therefore clear: today, no country can build the entire AI production stack alone. However, the countries best placed to produce AI capabilities will be those able to secure critical inputs under stress, from capital, energy and strategic materials to data, talent, chips and cloud compute.

Chips: Geoeconomics and global supply chains

Chips are crucial as they are the entry ticket to the frontier AI race. Their value chain is global and complex, with several layers from design to manufacturing. Investor attention has recently focused on design, one of the most profitable chip stages, especially in the U.S. Yet, design cannot stand alone: advanced chips still need European machinery, Asian manufacturing, high-bandwidth memory and specialized packaging. The U.S. leads the compute ecosystem and its firms earn just over 50% of global chip revenues, but its manufacturing share fell to 10% by 2022, with Taiwan and Korea now indispensable for leading-edge foundry and high-bandwidth memory respectively.

Geoeconomic forces are also reshaping the AI semiconductor value chain. China still lags in certain critical layers of frontier hardware, yet this is changing quickly, as export controls have accelerated its push for self-reliance; significant progress is already visible in its domestic chips and model efficiency. Europe is adapting too: the European Commission's June 2026 Technological Sovereignty Package links the Chips Act 2.0 with the Cloud and AI Development Act to strengthen advanced semiconductor capacity, support investment, accelerate datacentre deployment and reduce strategic dependencies in cloud and AI infrastructure.

Cloud: The second bottleneck

The key question is who can secure chips or cloud compute under stress. Most firms, start-ups and public agencies will not own large GPU clusters, but access AI compute through cloud infrastructure, managed platforms, model APIs and enterprise software. Frontier producers, by contrast, seek dedicated chip access through ownership, long-term reservations, custom silicon or cloud partnerships.

Global Infrastructure-as-a-Service reached $171.8B in 2024, with Amazon, Microsoft and Google controlling around 70.6%, giving the U.S. unmatched power to convert chips into AI services. China has domestic scale through companies like Alibaba and Huawei, but less global reach and chip constraints.

Public and private financing

Capital, both private and public, is a strategic input in frontier AI. Private investment is mainly needed for commercially scalable AI, from model training and data centres to chips and cloud platforms. Public money also has critical roles to play, from supporting public education and research to cushioning households and workers through the transition.

The shape of private financing is evolving. Early AI investment was funded mainly from hyperscalers' cash flows, rapidly rising financing requirements mean that it is increasingly being paid for through issuance of bonds, private credit and structured vehicles. Hyperscaler leverage still looks healthy, but some deals sit partly off-balance-sheet, so AI obligations may be understated in headline ratios. A sharp test may come from credit markets if private-credit losses expose overleveraged AI infrastructure investments, capital could move abruptly from abundance to scarcity.

Across geographies, the U.S. boasts a financing privilege based partly on imported capital, which reduces the cost of capital for hyperscalers and frontier laboratories. The risk for the U.S. is that in a fragmenting world, higher geopolitical risk premia, tighter technology controls and reserve diversification could undermine this privilege. Europe could improve its position if the Draghi-Letta agenda delivers a savings and investments union that redirects household savings from deposits and foreign securities towards investments. China's relative position could strengthen if lower-cost open models shift value from expensive training towards cheaper inference and deployment. Recent usage data suggests Chinese models have gained meaningful share of token usage in recent months, highlighting the growing competitiveness of China's software and model layer despite continued hardware constraints.

Figure 8: U.S. hyperscalers capex likely to reach $1T next year (24M capex plans)

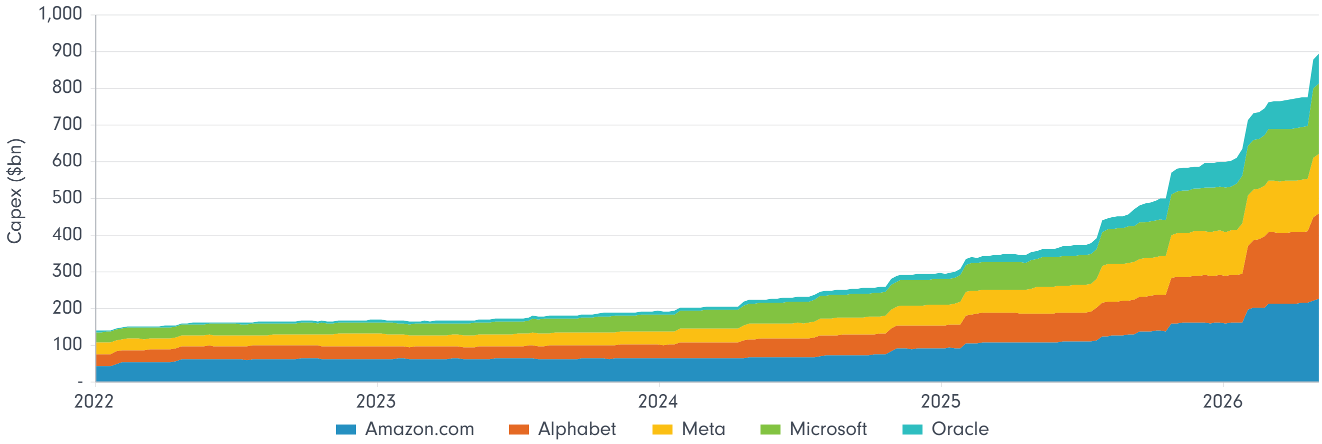

Source: Bloomberg, Fidelity International, 30 June 2026.

Venture capital

AI firms attracted $260B in 2025, equal to 61% of global venture capital, while U.S.-based firms received about $194B, around 75% of global AI venture deal value. The EU followed at $16B, while China and the U.K. each received $14B. The market is also increasingly shaped by concentration and mega-deals: Anthropic's $65B May 2026 round was more than four times the EU's entire 2025 AI venture inflow. Investor origin confirms the pattern: the U.S. supplied $124B, followed by the U.K. at $21B, China at $17B and the EU at $15B. However, China's published numbers probably understate capacity because a lot of its capital flows are through state banks and guided funds.

Public equity

The U.S. lead reflects scale, profitability and growth expectations, but also richer valuations and passive flows that allocate more money to the largest companies as index weights rise. Household balance sheets also matter: equities account for 55% of $72T of U.S. household liquid financial assets, while deposits account for only 14%. Europe has large savings, but they are not channelled efficiently enough into scale-up finance. EU household financial assets reached €40T in 2024, with 31% in currency and deposits, but direct listed-stock holdings were only 6% in the euro area, rising to 21% once funds, life insurance and pensions are included.

Private credit / structured infrastructure finance

AI datacentre financing requirements often exceed bank capacity and do not always fit unsecured bonds. Dedicated vehicles can own assets, borrow privately and repay through leases or capacity agreements with hyperscaler tenants, keeping upfront capex off technology companies' balance sheets while leaving exposure through leases, offtake commitments and guarantees. AI infrastructure capex over 2025 to 2028 could reach $3T, with $1.5T requiring external capital, including $800B from private credit. The Financial Stability Board estimates global private credit at $1.5T to $2.0T at end-2024, with the U.S. alone around $1T. Growth in private credit often leads to opaque valuations, leverage, liquidity mismatch and overbuilding risks, notable risks given that recent defaults have triggered redemption pressure in retail private-credit funds, exposing the mismatch between liquid promises and illiquid loans.

State funding

In the U.S., the Defence Advanced Research Projects Agency and National Science Foundation finance upstream research. The National AI Research Resource provides resources, while the current Administration has made removal of rules seen as innovation barriers central to the White House AI Action Plan. The EU response is more coordinated, with InvestAI aiming to offset weaker private platform scale through shared public infrastructure. China's state-capital system links AI research, semiconductor self-sufficiency, robotics, industrial automation and strategic procurement. Japan, the U.K., India and Canada matter, but remain smaller sovereign-capacity cases. Public finances also matter in financing diverse areas from education systems to social policies.

Energy and critical minerals

Energy and critical minerals are central to AI. Leading models need reliable, competitively priced electricity at scale, likely from a diversified mix of power sources, as well as the minerals and grid equipment behind it. No country fully controls this stack, so AI will deepen physical dependencies.

The International Energy Agency (IEA) estimates that data centres consumed around 415 TWh in 2024, about 1.5% of global electricity use, and projects 945 TWh by 2030, above Japan's entire current electricity consumption. AI is the main driver: demand from AI-optimized data centres is expected to more than quadruple, while U.S. data centres may account for almost half of electricity-demand growth to 2030. The constraint is generation plus delivery, because power must reach the right places reliably.

Existing grids were designed for older demand patterns, with the additional complexity that renewables like solar and wind are variable, often distant from demand centres, and require storage, flexibility and transmission to create economic value. The IEA estimates that around 20% of planned data-centre projects could face delays unless grid risks are addressed, while new transmission lines can take four to eight years in advanced economies.

A world of energy interdependence

The U.S. and China have the clearest scale advantages. The U.S. generated about 4,300 TWh in 2024 from a diversified, mostly domestic resource base, but data centres cluster in power markets where grid and planning frictions slow delivery. The U.S. looks increasingly energy-independent on a net basis, but still imports about 17% of domestic energy supply, reflecting needs of heavier-crude and Canadian gas because of pipeline geography and winter demand peaks. China too has vast domestic electricity production, although it also relies on significant energy imports. Renewables supplied about 35% of generation in 2024, rising 19% year on year, but coal still accounted for 62% of energy consumed, which matters if markets decide to distinguish between AI output and low-carbon AI output.

Europe's mix is cleaner, helped by France's nuclear fleet, Nordic hydro and wind, Iberian solar, but costlier and more import-dependent, with Germany and Italy still gas-exposed. In 2024, renewables supplied 47% of EU net electricity and nuclear more than 20%, but 2025 industrial electricity prices were more than twice U.S. levels and nearly 50% above China's. The U.K. generated 285 TWh in 2024, with low-carbon sources at about 65%, enough for clusters but not autonomy without faster grids, nuclear, offshore wind and storage.

Elsewhere, Japan, Korea and Taiwan are AI-hardware critical but energy-constrained. Japan's 2023 mix was 69% thermal, with 15% self-sufficiency. Korea generated about 60% from fossil fuels in 2024. Taiwan generated 83% from thermal power while importing more than 94% of its energy. Canada and Australia look attractive given hydro, land, solar and wind potential, although both need grids and storage. India has scale but heavy coal reliance, while the Gulf offers cheap energy and land but faces hydrocarbon and water-stress exposure.

Figure 9: Independence is vital in an increasingly fragmented world.

...characterised by the AI revolution, geopolitical risks, and the climate transition

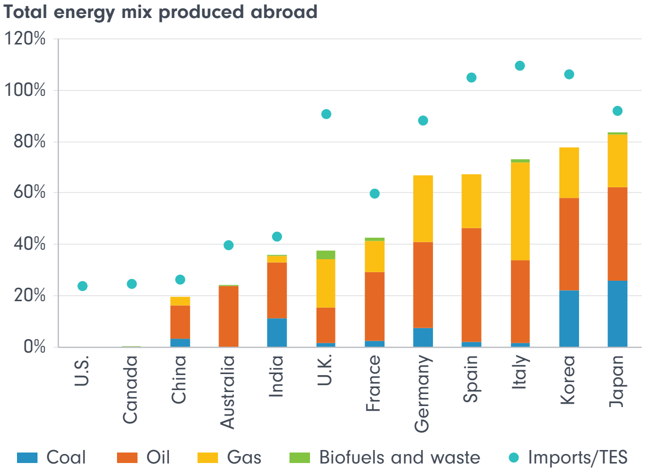

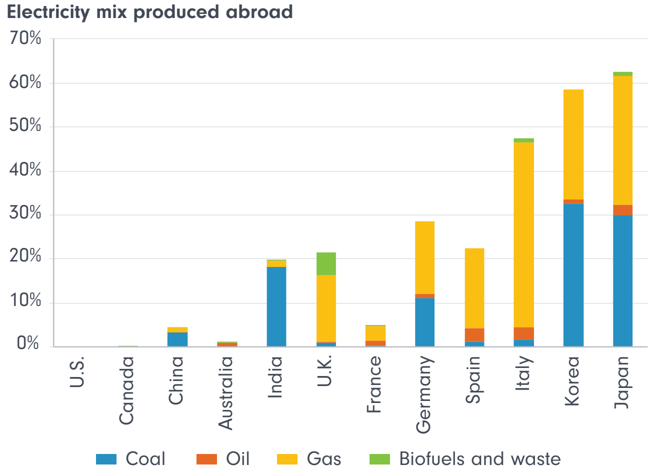

For illustrative purposes only. Assumptions are based on proprietary modelling, reflecting the views of investment professionals at Fidelity International.

Source: IEA, Fidelity International's calculations, June 2026.

Critical minerals: A strategic input

China is central in minerals production, as the dominant refiner of 19 of 20 energy-related strategic minerals (averaging about 70% by volume). China supplied 82% of world natural graphite in 2025, while 71% of U.S. rare-earth imports in 2021-24 came from China. Mining can look more diversified than refining, but this activity often sits in supply chains shaped by Chinese ownership, finance, joint ventures, offtake agreements and downstream refining (e.g. even nickel in Indonesia and cobalt in Congo).

Australia and Canada are key allied alternatives if they move from extraction into low-carbon processing. The U.S. remains import-dependent, especially on China, while the EU, U.K., Japan and Korea are resource-poor but can build resilience through recycling, high-value processing, advanced materials and allied supply chains. Critical minerals clearly introduce a geopolitical risk, as more than half of energy-related minerals face export controls; China has tightened restrictions on several minerals and processing technologies, so AI supply chains could be disrupted by rivalry as well as shortages.

Several trends could soften these constraints, although probably not remove them this decade:

- **Better model architecture, chip design and hardware-aware engineering** can deliver more capability per unit of input, but lower costs may expand total AI use, especially as reasoning models require more inference compute.
- **AI may also optimize energy production and distribution**, while accelerating breakthroughs in materials, batteries, storage and fusion.
- **Silicon photonics, glass archival storage and quantum computing** all offer significant long-term potential, but none materially changes near-term constraints.

Energy security is back

Investment in renewables and nuclear can reduce gas imports, but it cannot deliver full sovereignty. China controls 85% of solar supply-chain capacity, 95% of photovoltaic wafers, 80% of batteries and most strategic-mineral refining, while mined uranium needed for nuclear energy is concentrated in Kazakhstan, Canada and Namibia. Cost and speed favour renewables: IRENA puts 2024 installed costs at $0.69/W for solar, $1.04/W for onshore wind and $2.85/W for offshore wind, before grids, storage and backup; nuclear projects range from $2.7/W to $16/W. Also, reactors take seven years to build globally, over a decade in advanced economies, and 10 to 15 years in first-time nuclear countries. Small modular reactors may help sites but are unlikely to scale before 2030.

Energy independence is therefore a sequencing problem: renewables reduce fossil-fuel import exposure fastest before 2030, nuclear deepens low-carbon security in the 2030s and beyond, and the bottleneck may be the supply chain and grid rather than the plant.

Electricity grids: A hidden constraint

The IEA estimates that the world must add or refurbish more than 80 million kilometres of grids by 2040, roughly today's entire global grid. In the U.S., about 10,300 generation and storage projects were waiting to connect at end-2024. China leads in ultra-high-voltage transmission, storage and smart grids, but renewables have outpaced networks, creating curtailment and west-east mismatch.

Europe's grid is cleaner but older: 40% of EU distribution grids are over 40 years old, cross-border capacity must double by 2030, and investment needs about EUR600 billion this decade, rising to EUR1.2 trillion by 2040.

Figure 10: China dominates the refining and production phases of minerals needed for AI

While mining looks more diversified, at least geographically

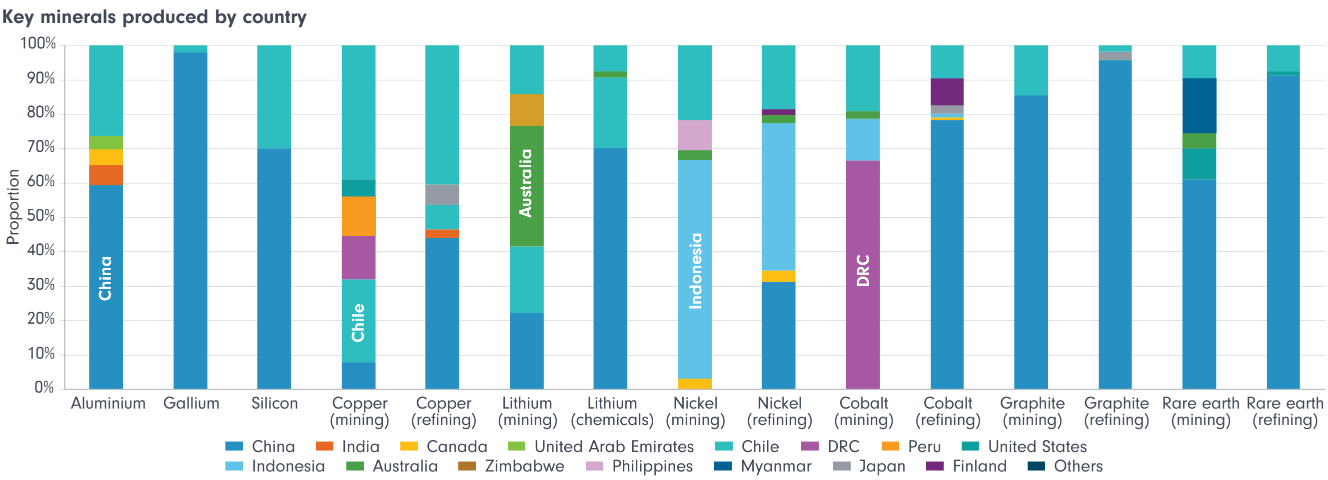

Source: U.S.GS, Fidelity International, 30 June 2026.

Data and model inputs: Quality and usage rights matter as much as quantity

Data is often described as the oil of AI, but unlike oil it can be reused, combined and repurposed many times, provided it is accurate and legally usable. Large populations may generate huge digital traces, but these may have limited value if they are unstructured, poorly labelled, legally unusable, under-represented linguistically or biased. Smaller countries can still hold valuable AI inputs through trusted health, financial, engineering, legal, geospatial, public-sector and enterprise data.

The U.S. has the strongest global flywheel, China the strongest sovereign domestic flywheel, and Europe the most deliberate trusted data-sharing regime. The next phase will reward sovereign, licensed, domain-specific and feedback-rich data: medical, industrial, engineering, financial, legal, tax, welfare, energy, transport, education and multilingual public-service datasets. This could benefit Europe, Japan, Korea and India. But data sovereignty is not AI sovereignty if domestic data mainly fine-tunes foreign models through foreign clouds and APIs, so the model layer remains strategically decisive.

Data scale

In relation to AI data, scale refers to the interaction between population, internet use, language, platform depth and target AI domains. China and India each have over 1.4 billion people. The EU has about 450 million if treated as one data space, the U.S. about 340 million and Japan about 124 million. Connectivity is high across these markets, though India is somewhat lower.

Anglophone countries, especially the U.S., U.K., Canada and Australia, benefit because English has dominated open AI training data. For example, English accounts for about 41% of common crawl primary-language HTML pages. China has vast domestic data in e-commerce, payments, logistics, mobility, manufacturing and public administration, supporting powerful Chinese-language models at home but reducing global portability. Europe faces language fragmentation, making interoperable, sector-specific and country-specific models vital, especially in health.

Domain relevance adds nuance. The U.S. has the strongest consumer, enterprise, developer and cloud data because its firms dominate search, software, advertising, social media, e-commerce, productivity tools and AI interfaces. China is strongest in payments, logistics, e-commerce, mobility, manufacturing and public administration. Europe, Japan and Korea are weaker in consumer platforms but strong in industrial and public-sector data, including manufacturing, robotics, transport, energy, health, finance and regulated services. The U.K. benefits from English, legal and financial sectors, universities and NHS depth, but lacks U.S. scale platforms.

Data usability

The U.S. historically enabled large-scale private data collection, creating a strong data-to-model flywheel, although copyright, privacy, publisher licensing and platform restrictions are now increasing uncertainty. Europe is stricter but aims to turn trust and interoperability into an advantage through the Data Act and Common European Data Spaces across health, agriculture, manufacturing, energy, mobility, finance, public administration, skills, science and the green transition. China's challenge is reciprocal access, as its 2024 cross-border data rules may reduce friction, but geopolitical frictions remain.

Building modern foundation models

- Pre-training teaches models statistical patterns from vast amounts of data.
- Post-training adapts them to follow instructions, answer questions, refuse unsafe requests and perform specific tasks using smaller but more curated datasets, such as instruction examples, expert demonstrations, preference data, evaluation data and, increasingly, synthetic data.

Talent and entrepreneurial ecosystem

Talent is the human core of the AI production stack, shaping both frontier AI production and AI diffusion. Frontier AI needs researchers, engineers, builders and founders moving between universities, labs and start-ups. Stanford's AI Index says nearly 90% of notable models in 2024 came from industry, so the key question is who converts STEM education, immigrant talent, venture capital, cloud and technology firms into one flywheel.

Future competition will be shaped by migration, retention, scale-up policy and education. The U.S. has been the best country at turning foreign students into AI researchers and founders, but its skilled-worker system is becoming harder to navigate, with recent changes affecting work permits and visa access, especially for start-ups and early-career researchers. Canada, the U.K. and China are responding with the 2026 International Talent Attraction Strategy, Global Talent and High Potential Individual visas, and K visas, respectively. The second trend is retention, as increasingly more top AI talent is staying at home, a shift that will help China first and India over time. The third trend is Europe's effort to fix the 'invent here, scale elsewhere' problem through the EU Startup and Scaleup Strategy and Scaleup Europe Fund.

Education and talent attraction

China and India dominate domestic technical supply. In 2020, China produced 3.57 million STEM graduates, 40% of the global total; India 2.55 million, 30%; and the U.S. 820,000, 9%. On intensity, in 2023, STEM represented over 40% of graduates in China, more than 30% in India and around 20% in the U.S. Germany is stronger than the U.S. at 36%, the EU average was 27%, Japan is lower, and Canada has 32% of new tertiary entrants in STEM.

Frontier AI also depends on specialized, mobile talent, where the U.S. lead is clearest. The OECD says the U.S. hosts 38% of all international doctoral students, while NFAP finds that around 70% of full-time graduate students in U.S. AI fields are international. That is the essence of the U.S. flywheel: it does not merely educate domestic talent; it converts global educational migration into research capacity, start-ups and frontier firms. The U.K., Canada and Australia share this openness advantage at smaller scale. Japan is much less magnetic, while India has huge domestic depth but still exports many top students and founders.

China is the strongest challenger because its talent base is increasingly elite. MacroPolo's AI Talent Tracker finds that 38% of elite AI researchers received undergraduate education in China, versus 24% in the U.S. and 10% in India, while U.S. institutions employ 59% of the world's elite researchers. The subtle point is that China produces a rising share of top talent, but the U.S. still captures much of it: only 11% of China-educated elite researchers work in China, while 72% work in the U.S.

Entrepreneurial conversion

Stanford records 40 notable U.S. AI models from 2024, versus 15 from China and three from Europe. One-thousand-nine-hundred-fifty-three AI companies were funded in the U.S. in 2025, 639 in Europe, 172 in the U.K., and 161 in China, which implies that roughly 56% of the global total falls to the U.S., 13% to Europe excluding the U.K., and about 5% each for the U.K. and China. The Global Entrepreneurship Monitor (GEM) reinforces the point: in 2024/25, 19% of U.S. adults aged 18 to 64 were starting or running a young business, a historically high level for the U.S. and nearly two-thirds above the average for high-income GEM economies. Immigrant founders remain central though, as they founded or co-founded 65% of top U.S.-AI companies.

Europe has Germany's technical depth, France's elite engineering base and AI clusters in France, Germany and the Netherlands, but not yet one flywheel. When ambitious AI firms need large customers, corporate partners and experienced scale-up leadership, the U.S. often becomes the more natural venue. The EU Commission claims that only 20% of European corporates actively engage with start-ups, versus 50% in the U.S. The U.K. and Canada are smaller hubs, Australia, Korea, Taiwan and Japan are more specialized, and the UAE and Saudi Arabia are building ecosystems through imported talent, capital and compute.

Pillar 2: Can a country exploit AI for structural economic growth?

This second pillar in our forward-looking framework for assessing AI's top-down impacts shifts the focus from producing frontier AI to exploiting AI across the economy. The decisive question is whether a country can turn models, software and automation into broad, lasting gains in productivity and social welfare, rather than allowing the benefits to accumulate mainly in a narrow group of frontier firms, financial centres and highly educated workers.

The assessment therefore rests on four variables:

1. Economic structure: AI diffuses faster across industries and sectors where activities are data-rich and organized at scale.
2. Workforce adaptability: Technology raises output only when workers and managers redesign tasks and acquire complementary skills.
3. The environmental impact of AI: AI growth will not be durable if it raises emissions, water stress, e-waste or local resistance.
4. Governance, safety and trust: Adoption depends on confidence that risks are managed.

AI and inequality

Inequality is the background force running through all four variables, and it should be analyzed in layers:

1. The first layer is global. Economies of the global north start with stronger connectivity, capital markets, education systems and institutions, while many economies of the global south still lack the basic foundations needed for AI diffusion. The UN estimates that about 30% of global population was still offline in 2024; internet use was 93% in high-income countries, but only 27% in low-income countries.
2. The second layer is sectoral and regional. AI may lift certain services and advanced manufacturing quickly, while agriculture, small retail, construction, care, local government and weaker regions adopt more slowly.
3. The third layer is individual. Even within the same country or sector, outcomes differ by education, income, age, gender, race, bargaining power and whether a person works in a firm able to invest in complementary technology. Gender illustrates this risk: a Danish study found women about 20% less likely than men to use ChatGPT, even within the same workplace and after controlling for detailed tasks. This matters because early users learn faster, shape workflows earlier and may capture more wage and promotion benefits.

This distinction matters because AI can fail economically even if it succeeds technologically. If the gains remain concentrated in a few firms, sectors and cities, they may not be large enough to move aggregate productivity, or they may be offset by weaker demand, job insecurity and political resistance. Eurostat reported that 20% of EU enterprises with at least ten employees used AI in 2025, up from 13% in 2024, but adoption remains uneven across countries and firm sizes. The U.S. may lead at the frontier, but it still faces a winner-takes-most risk because AI benefits are concentrated in large technology firms. China has scale and state capacity, but its urban-rural divide remains material. India has enormous diffusion potential, yet the relevant question is whether AI reaches informal workers, micro-enterprises and public services, rather than only large IT firms and digitally confident consumers.

Taxation implications

One of the most important trends to watch is taxation. Modern welfare states rely heavily on labour income, while effective corporate tax rates have decreased over the last 30 years (Figure 11): the OECD shows that, in 2023, social-security contributions were 26% of tax revenues and personal income taxes 25%, versus 12% for corporate taxes and 5% for property taxes. In 2025, the average tax wedge reached 35%, with Germany, France and Italy above 45%. If AI raises productivity while shifting income from labour towards capital gains and market power, companies may initially enjoy stronger margins and valuations, but political pressure to tax those gains will potentially rise if the benefits are not widely shared. Taxation is therefore both a redistributive instrument and a key variable for investors assessing the durability of post-tax corporate profits. The countries that benefit most from AI will need a new fiscal settlement, however global cooperation will be fundamental, although harder in a fragmented world, because profits, skilled workers, headquarters and data-centre investment can move across borders, making purely national tax strategies less effective.

Figure 11: Lower taxes and interest expenses have been tailwinds for companies profits

However interest rates have now increased and taxes can hardly continue to trend down

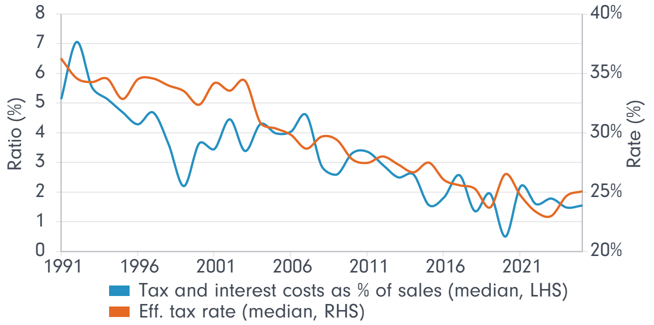

Source: Bloomberg, Fidelity International, 30 June 2026.

Economic structure

At its simplest, economic growth can come from higher labour productivity or more hours worked. The former ultimately depends on two tests: whether an economy has enough activities that AI can affect and whether adoption in those activities is broad enough to matter.

Going forward, two trends will matter most. The first is whether, and how quickly, companies can adapt their workflows. A lot still needs to happen before AI benefits show up at the economy-wide level. However, not all firms will benefit equally and there will be dispersion within sectors when it does, so disruption is also likely to intensify.

The second is whether industrial policy can re-shore strategic activities such as manufacturing, as is visible in parts of the U.S. This will be challenging because China's advantage is no longer just about cost, but also supplier depth, engineering capability and production know-how, which are hard to recreate quickly in the U.S. or Europe.

The timing of adoption across sectors and within economies will be important to monitor. While the first phase of AI benefits may favour economies like the U.S. and U.K., where LLM-led software and services adoption is strongest, progress in robotics, logistics and machine vision could shift later-cycle advantages towards more manufacturing exposed economies like China, Japan and Germany.

Service sector AI

AI can spread more quickly through services, in particular cognitive services where value is created from language, code, data, documents and decisions, from finance and software to health care administration and government.

For this reason, most high-income economies start with an advantage:

- In the U.S., private services-producing industries account for about 78% of GDP, similar to the U.K. and Singapore.
- Europe is large but less neat: services are about 66% of GDP. France and Spain are more service-facing, while Italy and Germany retain stronger industrial identities.
- Japan is also more service-based than often assumed, at 70%.
- China and India, at 56.7% and 49.9%, have narrower immediate cognitive-services exposure, but strong platform, IT, logistics and digital-public-infrastructure niches that could broaden over time.

But it is not just a question of how large the service sector is; the nature of the work also matters. Sector exposure creates the opportunity, but productivity gains arrive only when firms actually reorganise work around AI. Evidence suggests that adoption is indeed higher in the services most aligned with the cognitive channel: OECD data show that, in 2025, 57% of ICT firms and 37% of professional and scientific services firms used AI, the highest shares of any industry. Yet adoption is still nowhere near saturation. Across the OECD, only 20% of firms used AI in 2025, although this was up from 8.7% in 2023.

Figure 12: Different sectors have different exposure to AI (and robotics)

While current adoption is also uneven

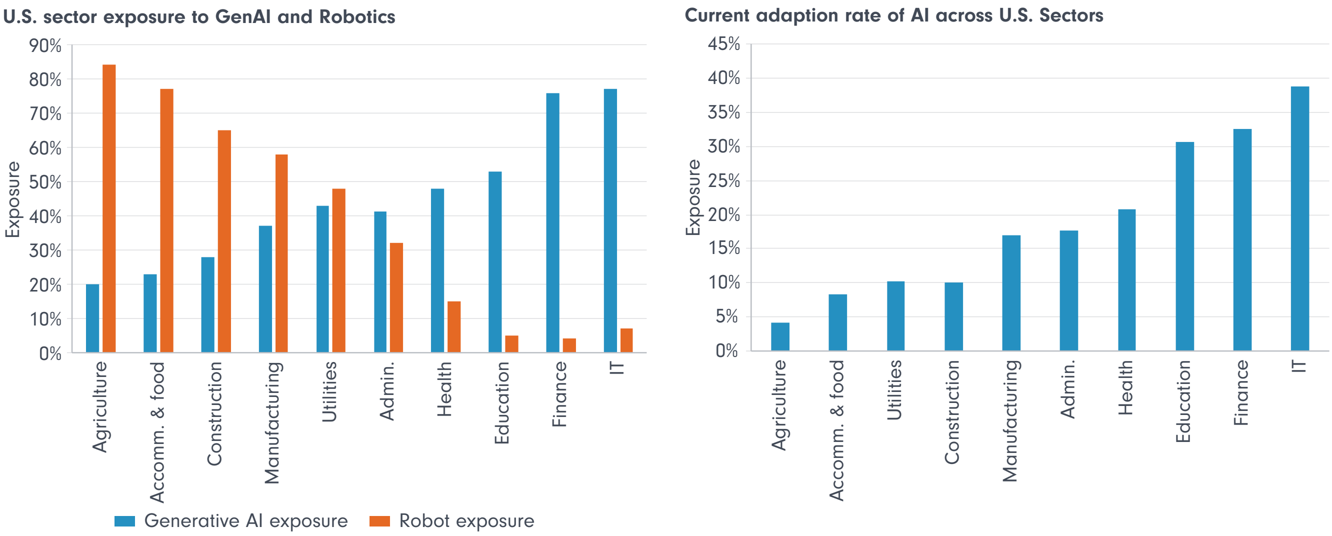

For illustrative purposes only.
Source: OECD, U.S. Census BTOS, Fidelity International, January 2026.

Physical AI

Stand-alone generative AI first affects white-collar services because the inputs are already digital. However, AI-augmented robotics extend the story to manufacturing, logistics and eventually agriculture, bringing blue-collar jobs into the AI transition.

Timing matters, however, as service AI can diffuse faster, while robotics requires machines, sensors, integration, safety systems and capital-stock replacement, so adoption is likely to be slower.

Manufacturing is becoming more machine-readable. Medium and high-technology industries represented 51% of global manufacturing value added in 2023; computers, electronics and optical products alone accounted for 12%, motor vehicles for 9%, chemicals and machinery for 8%. These are precisely the industries where AI-augmented robotics can matter most as it can reduce defects, accelerate design, improve maintenance, optimize inventories and make automation more flexible.

- China is central to physical AI and its policy framework is shifting from traditional automation towards high-end intelligent robotics integrated with AI. Its manufacturing share is 25% of GDP today and its robotics adoption is ahead of other countries with 295,000 industrial robots installed in 2024, 54% of global deployments. It has more than two million robots in operation and domestic robot suppliers serving 57% of its market, this is becoming an embodied-AI advantage.
- Korea, Japan and Germany are the other leading industrial-AI cases: Korea has manufacturing at 27% of GDP and the world's highest robot density, at 1,220 per 10,000 workers; Japan, at 21%, installed 44,500 robots in 2024; Germany, at 18%, installed 27,000, equal to 32% of Europe's total. Italy (9,000 installations), Spain and France (5,000) are smaller but matter too.
- The U.S. has manufacturing at 10% of GDP and 34,200 installations, but leads in advanced manufacturing, aerospace, defence and industrial software.
- India, at 13%, installed a record 9,100 robots and is already a top-ten manufacturing economy.

Workforce adaptability

Workforce adaptability, for both workers and managers, is the hinge between AI as a productivity revolution and AI as a labour-market shock. Innovation can destroy, augment and create new jobs, with the balance between these effects shaping economic and social outcomes. The IMF estimates that almost 40% of global employment is exposed to AI, rising to around 60% in advanced economies.

AI, robotics and aging will reshape labour scarcity, particularly for SMEs. The OECD has reported that nearly 40% of those experiencing a skills gap say that generative AI has already helped compensate for it, while a quarter say it has helped address worker shortages. Interestingly, the reported benefits are higher in Japan and Korea, where aging is already tightening labour supply. As reported in our Analyst Survey, China is especially relevant here, because AI-robotic innovation and adoption are accelerating, although current examples also show that full autonomy still requires better models, deeper hardware-software integration and richer 3D movement data. At the same time, however, AI agents may also erode entry-level service roles, weakening the ladder from graduate to experienced professional.

The plausible central case is therefore neither job collapse nor leisure utopia, but a material and uneven reorganization of work. Many jobs will change substantially and some workers may need to step outside their comfort zone, reinvent their role, or move into different occupations, rather than simply adapt at the margin. Bottom-up company evidence is supportive: in our 2026 Analyst Survey, 90% of our analysts claimed that at least some of companies they cover are benefiting from AI productivity gains, but 62% said that those benefits remain confined to a minority of firms. About 50% of analysts also reported that companies are doing more with the same number of workers, suggesting the first labour-market effect is often "same headcount, more output".

AI skills: A crucial adoption factor

The OECD reports that around 40% of employers (in particular, SMEs) in manufacturing and finance that had not adopted AI cited skills as the main barrier against adoption. It notes that few are likely to need advanced AI-specific skills, such as programming or model development, but more will need digital skills, data interpretation, problem-solving, creativity and managerial capability. On foundational skills, Japan remains the benchmark. The U.K. is also strong, ranking second among G7 economies in literacy and problem solving, while Canada looks similarly well placed. Germany is the strongest large EMU economy, but the euro area is uneven, with France, Spain and Italy weaker. The U.S. combines world-leading AI firms with only middling adult skills and a large low-skills tail, making diffusion powerful but socially uneven.

AI diffusion

Diffusion depends on firms. The U.S. is strongest in short-run adoption as it combines frontier firms, high management capability and rapid job reallocation. This accelerates adoption but can also transmit shocks quickly to clerical, administrative, customer-service and junior professional roles. The UK, Canada and Australia are also relatively flexible, but with less scale. Germany and Japan are more likely to adapt through large firms, internal training and process reorganization, while France, Italy and Spain face greater frictions because weaker adult skills meet more regulated labour markets. China's system can redeploy labour through industrial policy, but less through open, portable job-to-job mobility.

Adult learning

Adaptability depends increasingly on mid-career renewal. More than half of AI-using workers report employer-funded training and trained users are more likely to report better performance and working conditions. Europe has a formal advantage here, as it has stronger social institutions and policy language around lifelong learning, but the numbers are still not adequate for an AI shock: Eurostat reports that only 28% of EU adults aged 25 to 64 participated in education or training in the previous 12 months in 2024. China can mobilize scale, having pledged subsidized vocational training for more than 75 million people during the 14th Five-Year Plan. The U.S. relies more on employer-led and market-led re-skilling, which can be fast at the frontier but leaves many workers exposed. The UK, Canada and Australia sit between the two models: they have better average adult skills than the U.S. and more flexible labour markets than much of continental Europe, but their re-skilling systems still need far greater depth if AI disruption reaches white-collar services quickly.

The environmental impact of AI

Environmental outcomes will depend both on whether AI can be produced sustainably and if it can be used to solve environmental problems. It can support grid optimization, industrial efficiency, climate modelling, methane detection and materials discovery, with the IEA estimating 1.4 Gt of potential CO_2 reductions in 2035. However, energy mix transition, governments' policy consistency and international cooperation will decide whether those gains materialize.

At present, the most visible environmental pressure point is the data centre. The AI boom is turning abstract computing demand into concrete local questions about electricity grids, CO_2 emissions, water withdrawals, land use, noise, electronic waste adding a new layer of pressure to pre-existing global problems. At the macro level, some of these impacts may still appear modest relative to the wider economy, but they are landing on local systems that are already under strain. In the U.S., for example, rapid data-centre expansion has already triggered local pushback and congressional scrutiny in certain regions.

Looking at various regions, the EU is the clearest leader on environmental governance because it is trying to make AI's footprint measurable, comparable and eventually manageable. Its Energy Efficiency Directive requires large data centres to report energy performance, the European Commission is developing efficiency tools that include energy and water indicators, and the AI Act introduces energy transparency for general-purpose AI models.

Water constraints

Water is increasingly becoming a production constraint and could play a growing role in determining whether a country can produce frontier AI. It runs through the value chain, as data-centre cooling, power generation, semiconductor fabs, mining and mineral processing can all be water-intensive. Siting therefore matters. AI-training sites can be placed away from end-users, so developers often choose smaller towns or former industrial regions with cheap land, fast permits, tax incentives and grid access. The fiscal bargain can look attractive, because data centres raise property-tax revenues without adding many residents, but a 100 MW hyperscale site in the U.S. can consume around two million litres a day across its full footprint, equivalent to about 6,500 households.

The water risk is therefore local as much as national. A country may have enough water overall, yet a new data-centre cluster can provoke resistance if it sits in a hot, dry or sensitive watershed. Asia-Pacific is especially exposed, as it potentially accounts for around half of projected data-centre water consumption by 2030. Taiwan is the clearest warning signal: during its 2021 drought, reservoirs fell below 5% capacity and more than one million households and businesses faced rationing, industrial users were asked to reduce water use, and 74,000 hectares of rice fields lost irrigation while chip production was protected.

Emissions

At a time when global emissions need to fall, AI-driven data-centre growth is pushing in the opposite direction. The IEA estimates data centres produced around 180 Mt of indirect CO_2 in 2024, or 0.5% of global fuel-combustion emissions, already comparable to a country like Spain.

In the U.S. alone, data centres could reach 3-5% of combustion emissions by 2035. The decisive variable is therefore whether electricity systems decarbonise fast enough: China is exposed if coal remains important, while Europe benefits from cleaner power.

Input materials / E-waste

Wider environmental risk sits in hardware, minerals and e-waste. Servers, networks, cooling systems and AI devices will need rapid upgrading, creating obsolescence pressure. AI mineral demand is still small versus clean energy, but by 2030 data-centre expansion could require up to 2% of 2024 copper and silicon demand, over 3% of rare earths and 11% of gallium. This amplifies local risks in China, Indonesia, Chile and DR Congo, including child-labour concerns in cobalt and nickel-related pollution. E-waste is the other pressure point: 62 million tonnes was generated in 2022, only 22.3% was formally recycled, and generative AI could add 1.2-5 million tonnes by 2030.

Governance, safety and trust

AI regulation is moving beyond broad principles to practical implementation. The central challenge is to build an AI governance system that does not suffocate innovation, but channels it towards fairer, safer and more accountable outcomes, so that the benefits of AI are shared more widely across firms, workers and society as a whole.

Governments are already translating commitments to safe, accountable, human-centred AI into laws, standards, procurement rules and supervisory expectations. Generative AI has accelerated this shift because regulation now increasingly covers foundation models (general-purpose systems that can be reused across many applications, rather than only specific use cases such as credit scoring, recruitment or medical diagnosis).

Europe: Comprehensive

For EU countries, the central framework is the EU AI Act, which is comprehensive, horizontal and risk-based. It applies across the economy, covers many actors in the AI value chain and combines hard prohibitions on practices such as social scoring and harmful manipulation with stricter obligations where potential harm is higher. From August 2025, rules also began applying to general-purpose AI models, including foundation models, with obligations on documentation, transparency and additional risk-management duties for the most powerful models.

U.S. and UK: Flexible

The U.S. and U.K. share a more decentralized approach, although their institutional designs differ. The U.S. has no single federal AI law equivalent to the EU AI Act, relying instead on sector regulators, state-level initiatives, procurement rules and voluntary standards, especially the AI Risk Management Framework. The U.K. similarly asks existing regulators to apply common principles such as safety, transparency, fairness, accountability and contestability within their own sectors, while investing heavily in frontier-model evaluation through the AI Security Institute. The strength of the U.S. and U.K. approach is flexibility, speed and closeness to technical development, while the weakness is fragmentation and the risk that protections vary by sector, regulator, state or political cycle.

Four key governance tests

Cybersecurity: Frontier AI will increasingly be used both to defend and to attack critical systems. Anthropic's Project Glasswing and Claude Mythos Preview show the direction of travel: The same models that can help defenders find vulnerabilities can also raise fears about automated exploitation, particularly in banks and other critical-infrastructure sectors.

Deepfakes: Synthetic voice, video and identity fraud threaten the basic trust mechanisms of markets, elections and institutions.

User protections: Especially for children and teenagers, will be central as AI intersects with mental health, manipulation, privacy, education and child development. Australia's under-16 social media ban shows age-based controls can emerge quickly; similar pressure may reach AI tools through authentication, age assurance, access controls and safety testing.

Creation of new international frameworks and institutions: Some have already emerged, most notably the International Network of AI Safety Institutes, launched by AI safety bodies and government-mandated offices from countries and jurisdictions including the U.S., U.K., EU, Japan, Korea, Canada, Australia and Singapore. At this stage, these initiatives are better understood as the early plumbing of international AI governance rather than as global regulation: common language, common testing methods, comparable risk reports, shared incident information and diplomatic channels for future crises.

Important Information

This material is for Institutional Investors and Investment Professionals only and should not be distributed to the general public or be relied upon by private investors.

This material is provided for information purposes only and is intended only for the person or entity to which it is sent. It must not be reproduced or circulated to any other party without prior permission of Fidelity.

This material does not constitute a distribution, an offer or solicitation to engage the investment management services of Fidelity, or an offer to buy or sell or the solicitation of any offer to buy or sell any securities in any jurisdiction or country where such distribution or offer is not authorised or would be

contrary to local laws or regulations. Fidelity makes no representations that the contents are appropriate for use in all locations or that the transactions or services discussed are available or appropriate for sale or use in all jurisdictions or countries or by all investors or counterparties.

This communication is not directed at and must not be acted on by persons inside the United States. All persons and entities accessing the information do so on their own initiative and are responsible for compliance with applicable local laws and regulations and should consult their professional advisers.

This material may contain materials from third parties which are supplied by companies that are not affiliated with any Fidelity entity (Third-Party Content). Fidelity has not been involved in the preparation, adoption or editing of such third-party materials and does not explicitly or implicitly endorse or approve such content. Fidelity International is not responsible for any errors or omissions relating to specific information provided by third parties.

Fidelity International refers to the group of companies which form the global investment management organization that provides products and services in designated jurisdictions outside of North America. Fidelity, Fidelity International, the Fidelity International logo and F symbol are trademarks of FIL Limited. Fidelity only offers information on products and services and does not provide investment advice based on individual circumstances, other than when specifically stipulated by an appropriately authorised firm, in a formal communication with the client.

Europe: Issued by FIL Pensions Management (authorised and regulated by the Financial Conduct Authority in UK), FIL Investment Management (Luxembourg) S.à r.l. (authorised and supervised by the CSSF, Commission de Surveillance du Secteur Financier), FIL Gestion (authorised and supervised by the AMF (Autorité des Marchés Financiers) N°GP03-004, 21 Avenue Kléber, 75016 Paris) and FIL Investment Switzerland AG.

UAE: The DIFC branch of FIL Distributors International Limited is regulated by the DFSA for the provision of Arranging Deals in Investments only. All communications and services are directed at Professional Clients and Market Counterparties only. Persons other than Professional Clients and Market Counterparties, such as Retail Clients, are NOT the intended recipients of our communications or services. The branch is established pursuant to the DIFC Companies Law, with registration number CL2923, as a branch of FIL Distributors International Limited, registered in Bermuda. FIL Distributors International Limited is licensed to conduct investment business by the Bermuda Monetary Authority.

In Hong Kong, this material is issued by FIL Investment Management (Hong Kong) Limited and it has not been reviewed by the Securities and Future Commission.

FIL Investment Management (Singapore) Limited (Co. Reg. No: 199006300E) is the legal representative of Fidelity International in Singapore. This document / advertisement has not been reviewed by the Monetary Authority of Singapore.

In Taiwan, independently operated by Fidelity Securities Investment Trust Co. (Taiwan) Limited 11F, No.68, Zhongxiao East Road, Section 5, Taipei 110, Taiwan, R.O.C. Customer Service Number: 0800-00-9911.

In Korea, this material is issued by FIL Asset Management (Korea) Limited. This material has not been reviewed by the Financial Supervisory Service and is intended for the general information of institutional and professional investors only to which it is sent.

Issued in Japan, this material is prepared by FIL Investments (Japan) Limited (hereafter called "FIJ") based on reliable data, but FIJ is not held liable for its accuracy or completeness. Information in this material is good for the date and time of preparation and is subject to change without prior notice depending on the market environments and other conditions. All rights concerning this material except quotations are held by FIJ and should by no means be used or copied partially or wholly for any purpose without permission. This material aims at providing information for your reference only but does not aim to recommend or solicit funds /securities.

For information purposes only. Neither FIL Limited nor any member within the Fidelity Group is licensed to carry out fund management activities in Brunei, Indonesia, Malaysia, Thailand and Philippines.

FIL Investment Management (Australia) Limited (ABN 34 006 773 575, AFSL No. 237865) and FIL Responsible Entity (Australia) Limited (ABN 33 148 059 009, AFSL No. 409340) are the legal representatives of Fidelity International in Australia. This is intended as general information for a person who is a "wholesale client" under section 761G of the Corporations Act 2001 (Cth) and has been prepared without taking into account any person's objectives, financial situation or needs. Not for distribution to or use by retail clients.

GCT260704GLO

FIDELITY CANADA INSTITUTIONAL™

FCI-4554050 07/26