

For professional clients and qualified investors only.
Capital at risk.



AI: A new division of labour?

By: Matthew Rodger



ChatGPT's launch back in late-2022 seems inauspicious in hindsight. The tool, which was powered by OpenAI's* GPT-3.5 model, couldn't tell the time properly, or follow a game of tic-tac-toe. And yet its release, not even four years ago, birthed an artificial intelligence (AI) rush that's among global macro's most important trends. Yet underlying it lies the solution to a simple problem computer scientists have wrestled with for decades.

Computers are known to solve well-understood processing problems (say, multiplication) quickly. But while they were very good at calculation, they always needed a well-specified algorithm, which many human-level skills (recognising a friend on the street, translating a language) just don't have. Although computers have improved substantially over the last 80 years and are now central to modern life, getting them to perform these higher-order tasks proved elusive. That clear separation of

tasks that computers could and could not do is now broken. Today's large language models (LLMs) are now capable not just of replicating human outputs in language and images, but operating computers, conducting scientific research and interacting with one another. Regardless of your view on AI, the pace of improvement in recent years has been extraordinary.

Best model intelligence over time




Figure 1: Artificial Analysis as at July 2026. This chart tracks intelligence improvements across OpenAI*, Anthropic*, and Google* models from early 2023 through mid-2026, revealing steady advancement. The index (derived from standardized benchmarks) shows models tightly clustered at the frontier. The scoring reflects relative performance gains, not a 60x absolute improvement. **Past performance is not a guide to the future. Assumptions, opinions, and estimates are provided for illustrative purposes only. There is no guarantee that any forecasts made will come to pass. *For illustrative purposes only. Reference to a particular security is on a historic basis and does not mean that the security is currently held or will be held within an L&G portfolio. The above information does not constitute a recommendation to buy or sell any security.**

Stress-testing the bull case

Although the progress in the technology has been remarkable to watch, there remain questions (both novel and prosaic) that persist in AI's wake around scaling, labour markets and profitability. These questions determine whether AI remains a technological story or becomes a durable, long-term macro and market driver.

An infrastructure boom with no ceiling in sight


Firstly, on scaling. It's well-known that AI relies on colossal volumes of computation and memory chips to both train new models and to operate those in service. Glutinous demand among builders of this infrastructure for ever-more chips to power AI models has propelled Nvidia* to the world's most valuable company, and Korea's memory manufacturers (Samsung* and SK Hynix*) to trillion-dollar valuations. The rush to deploy these chips in datacentres has also driven a surge in demand for power, cooling and communications infrastructure worldwide, as physical constraints to deployment tighten.

Yet some have wondered if this buildout is based on a misconception. Can this trend continue indefinitely? It's true infrastructure has sustained further model improvements, powering reasoning capabilities and agentic tool use, but it's possible some new constraint could come to derail the returns on all this spending.

It's true infrastructure has sustained further model improvements, powering reasoning capabilities and agentic tool use, but it's possible some new constraint might come to derail the returns on all this spending.

Sceptics have various anxieties. Some worry about diminishing returns from this spending, with the effect on model quality from each new chip or dataset dragging down the healthy returns that had been expected. This is essentially what happened with railways in the 1800s, which saw repeated booms and busts with each new wave of investment. In more recent history, the dot.com bubble saw miles of fibre cable and server farms built without the uptake people expected. Others worry if the capabilities of LLMs are as strong as proponents claim. Sceptics say we gain much of our knowledge of the world through experiences and experimentation, much of which mere text (on which LLMs are based) cannot capture. If true, this implies a hard ceiling which LLMs will never breach and a limit on the possibilities (and returns) on this brand of AI.

Infrastructure Capex as Share of US GDP



Sources: NBER, FCC/US Telecom, company filings, as of July 2026.

Figure 2: Capital expenditure on selected infrastructure throughout US history. Earlier booms in railroad spending, alongside servers during the dot.com bubble, clearly visible. AI investment (based on company filings) recently outpaced telecoms infrastructure at its peak in the early-2000s, though remains some way behind railroad investment, which reached a peak of a little below 5% of GDP.

Assumptions, opinions, and estimates are provided for illustrative purposes only. There is no guarantee that any forecasts made will come to pass. *For illustrative purposes only. Reference to a particular security is on a historic basis and does not mean that the security is currently held or will be held within an L&G portfolio. The above information does not constitute a recommendation to buy or sell any security.

However, so far, proponents claim, these remain distant risks to the AI theme, with little sign of slack emerging in infrastructure or in model performance. Labs continue to pay exorbitant funds for each new chip, making for a healthy return on existing AI infrastructure and encouraging further investment. Breakthroughs at the model level continue to proceed at pace, quietening fears of diminishing returns so far. OpenAI’s* ChatGPT solved previously unanswered problems in mathematics and AI lab Anthropic* launched Claude Mythos, which proved capable at cracking cyber risks in systems previously deemed the gold standard. Despite the concerns of sceptics, diminishing returns so far remain a distant prospect.


The flesh is weak?

Secondly, on the labour market. Technology’s interaction with human labour is one of the central questions of economics, one which goes back to the discipline’s founding. Adam Smith, the founder of political economy, highlighted 250 years ago how machines “abridge and facilitate” human handlers, making workers more productive and encouraging profitable usage of resources elsewhere in the economy.

Many jobs that we have today, from airline pilots to social media managers, would have been inconceivable to those in Smith’s day, but follow from this relationship between labour and technology. Moreover, in the fields that became more automated (agriculture, manufacturing) even the poorest in our society today is far wealthier than the richest in 1776. The wealth of access we have to information, medical care, and travel, is again a consequence of continual technological change that’s benefitted human labour. Does this mean we should relax about AI’s job impact, confident this natural process will repeat itself?

Sceptics argue not, for two reasons. First, they highlight that technological transitions can be tough for affected workers to navigate. Artisanal weavers and telephone operators had to reckon with technological change that automated away their skills, forcing them to settle for less in a new labour market. Many expect that much the same will happen to call centre workers and translators, and perhaps lawyers and coders, with AI automation steadily eroding labour’s importance. Taken together, it’s probably true that the pace of AI improvement has been faster than in previous technology shocks, like railways or electricity. Second, the more alarming fear is that AI itself is different to railways and smartphones, as it can learn and act independently outside of human control, spelling trouble for any occupation observed long enough. If AI capabilities prove as powerful as boosters claim, a much wider impact on human labour could follow.

AI occupational exposure



Source: Felton, Raj & Seamans (2023), ‘How will Large Language Models like ChatGPT Affect Occupations and Industries?’

Figure 3: A chart adapting research from Felton, Raj & Seamans (2023) showing the exposure of each occupation to LLM capabilities. Scores for over 770 occupations are accumulated and group into sectors. White-collar work, particularly law, education and business operations are highly exposed to AI in this framing, with physical labour in construction and agriculture more insulated.

Despite how intense this debate has been, the one thing we are confident in so far, is that AI has yet to produce the large labour market displacement the jobs bears expected. Youth unemployment among 20-24-year-olds in the US remains broadly unchanged since early 2022 and has been falling over the past six months, contrasting with softer trends in the UK and suggesting labour markets have so far absorbed AI's arrival with more resilience than feared. Although job openings in affected sectors (like coding) have receded, it's tricky to attribute this directly to AI with various other shocks (rate hikes, energy crises) hitting the global economy.

Early signs do point to a productivity uplift in AI-exposed roles, alongside hiring slowdowns in white-collar professions, helpful for growth at a tremulous time.


Early signs do point to a productivity uplift in AI-exposed roles, alongside hiring slowdowns in white-collar professions, helpful for growth at a tremulous time. However, meaningful upward wage pressure has yet to materialise, suggesting firms are absorbing

efficiency gains rather than changing labour conditions (in wages or inputs) outright. We remain open minded on the jobs question and will monitor incoming data to watch for a painful AI-induced shift in labour's influence on the global economy.

The value chain is clear; the winners are not

Lastly, there remains a question about AI's lasting profitability. Although equity markets remain abuzz with excitement, debate still rages over where the ultimate source of value will eventually lie, which depends on the nature of the technology itself. Three plausible sources of value lie in infrastructure, datacentres and AI labs. Infrastructure providers, in particular chipmakers (like Nvidia* or TSMC*) and energy companies (like Siemens* or Mitsubishi*), benefit from an ever-intensifying pursuit of AI's scaling laws, with a clear 'winner' still in doubt. Datacentre providers, from established players like Oracle* but also newer entrants like CoreWeave*, benefit from a broad diffusion of AI, to maximise usage, and a wide selection of capable models to run, to encourage retention. Labs, like OpenAI* and Anthropic*, continue to spend vast sums in pursuit of more capable models, in the hope of capturing a valuable niche (like coding or legal work) or force other players to withdraw. Four years into the buildout, it remains unclear which business model will win out.

AI value chain: where does profitability sit?



Sources: Bloomberg (NVDA, TSM, ASML, MSFT, GOOG, AMZN - T12M as of June 2026; The Information; FT.

Figure 4: Performance metrics from across selected AI-oriented companies, grouped by their place on the supply chain, as either infrastructure, datacentre provider or research lab. The chart displays strong margins for those upstream (infrastructure, datacentre) but negative margins for the labs, despite strong revenue growth. There remains considerable uncertainty about which layer in this chain will be the most profitable in future. ***For illustrative purposes only. Reference to a particular security is on a historic basis and does not mean that the security is currently held or will be held within an L&G portfolio. The above information does not constitute a recommendation to buy or sell any security.**

The value of an investment and any income taken from it is not guaranteed and can go down as well as up, and the investor may get back less than the original amount invested.

All three frameworks could fall over. Infrastructure providers could see their rosy earnings projections fade away if scaling laws reach a limit or their IP edge erodes. Datacentres could become commoditised, like off-patent drugs, eroding their margins considerably. And labs might struggle to sustain their rate of spending if models don't bring in the revenue they promise. Models made open source, with weights freely available to users, are (by at least some measures) closing the performance gap against peers produced by OpenAI or Anthropic, potentially eroding lab revenue growth as diffusion proceeds. Yet we believe claims that the technology as a whole is worthless vapourware is well wide of the mark. Consequently, we believe a balanced approach between these three AI 'sectors' could pay dividends in the years ahead.

Could productivity optimism meet policy friction?


AI's immediate effects on the macro economy have been mixed. Aggregate productivity has been slow to materialise, despite strong anecdotal evidence. Thus far, the boom has probably been inflationary, boosting prices in electronic components (particularly memory chips) and electricity (suddenly in demand to power data centres). Indeed, many of the delays in constructing new datacentres have come from procedural difficulties with grid connections and electricity access, with regulatory change in infrastructure permitting expected to have positive

impact. Investment in data centres has supported growth, but with many of the components (chips, gas turbines) imported, leaving the net effect on US GDP has likely been muted (even if Taiwan and South Korea are growing rapidly).

Longer term, we believe there's every reason to feel optimistic about future growth. AI's performance in services dominant in rich countries like law and healthcare, implies a substantial growth pick-up, even if there's some labour market disruption and if diffusion takes time. Firms may not realise these gains immediately, but as they experiment with new workflows and AI tools, they should discover new methods of working more productive than those which came before. Newer models have crossed key capability thresholds, such as in computer programming, that will encourage a much wider 'computerisation' of the economy, as the costs of writing software decline. We judge that even if model performance halts, the application of AI still has much further to run, giving us a positive bias on global growth.



Avg Cost of Intelligence - All Models per Threshold



Source: artificial analysis.ai, as of June 2026.

Figure 5: The average cost of intelligence, measured by the average price of 1 million tokens among models of a certain score on the Artificial Analysis intelligence index. This shows consistent declines in the average price of models in a given intelligence class since the sample begins in 2025. **The value of an investment and any income taken from it is not guaranteed and can go down as well as up, and the investor may get back less than the original amount invested.**

AI will also raise subtle policy challenges, particularly in implementation. For now, AI systems lack key attributes of 'personhood', such as the capacity to bear responsibility or culpability. Liability can, in principle, be assigned through institutional design (as with corporations or other legal fictions), but this may take time to design well and for economic agents to get used to.

Effects on the tax system are also going to be substantial. As more production is delegated to capital (in AI programmes, datacentres) how will the tax system change to accommodate? Labour income in rich countries is currently taxed much more than capital income. While this gap is designed to encourage productive investment in the economy, if AI's effects produce a bonanza for capital owners along with disruption for workers, it'll be tricky to politically defend capital's special treatment.

A new Great Game


The rollout also takes place in a parlous geopolitical situation. Amid this backdrop, AI has already become a foreign policy tool, as global powers race to define and spread global standards and benchmarks (through China's WAICO, the EU's AI Act and the US Pax Silica initiative) that favour domestic AI producers. The rivalry between the US and China threatens Taiwan, where most of the world's advanced AI chips are made. Continued advances made by American AI labs, if deemed a security threat by Beijing, could encourage greater threats to the world's AI chip supply. Further development of autonomous drone or cyber-based weapons, powered by AI, could spiral into a dangerous arms race between the US and China, with grave consequences for global security. Moreover, the diffusion of AI technology empowers more marginal

groups with dangerous capabilities, in cyberwarfare, bioweapons or chemical weapons, that heighten tail risks to wider society.

The Revolution's Reckoning

We believe the AI theme poses material opportunities and risks to markets over any investable time frame. In our view, the lesson, for now is one of cautious optimism. The technology has already exceeded many expectations since GPT-3.5's launch nearly four years ago. The fact that such systems have gone from not telling the time to building codebases and legal briefs speaks volumes about our ability to be positively surprised by trends that seem inconspicuous at first. The potential of the technology to radically transform the world, as the railway and the internet did, is evident, but so too are the risks: infrastructure investment may prove excessive, and value could prove hard to extract.

Beyond markets, the technology raises difficult questions about labour's role in the economy, distribution of AI's gains, and geopolitical risk, that will persist for years into the future. History suggests that transformative technologies never follow a straight path; periods of enthusiasm are often accompanied by setbacks, misallocated capital and unforeseen consequences. Nevertheless, the breadth of AI's potential applications, combined with the pace of recent progress, gives us confidence that its long-term impact will be substantial. In such an environment, humility is warranted. Our approach will be to follow incoming evidence on AI's development, impact and consequences, remaining alert to sudden changes. After all, history suggests the biggest effects of technology are often those that are hardest to see.



AI glossary

Agentic AI: AI systems capable of independently planning and executing multi-step tasks (browsing the web, writing and running code) with minimal human instruction.

Data centre: A facility housing the servers and computing hardware used to train and run AI models. AI's energy and cooling demands have driven a global surge in data centre construction.

Large language model (LLM): The class of AI model underpinning tools like ChatGPT. Trained on vast quantities of text, LLMs learn statistical patterns in language, enabling them to generate fluent, contextually relevant responses.

Reasoning models: A newer generation of LLMs that pause to 'think through' problems in structured steps before responding, improving performance on complex tasks like mathematics or coding.

Scaling laws: Empirical relationships showing that model performance improves predictably as more compute, data, and model parameters are added. These laws have justified enormous infrastructure investment, though debate persists over whether they hold indefinitely.

Tokens: The basic units AI models process, roughly three-quarters of a word on average. Model pricing and capability benchmarks are typically quoted per million tokens, making tokens the standard unit of AI's economics.

Training vs. inference: Training is the computationally intensive process of building a model; inference is running it to generate responses. Both require significant hardware, though their demands differ.



From the investment floor: Portfolio implications

Based on current expectations, in the Q2 earnings season, the US information technology sector is expected to deliver year-on-year earnings growth of 60%. Capital spending related to the AI build-out is supporting turbocharged profits, and reshaping both equity and credit indices. Concentrated equity markets may increase market sensitivity and dependency on the performance of a handful of stocks in Europe, North America and, especially, Asia. Corporate updates from the tech giants at times rival economic data releases or the pronouncements of central banks. That has upended the world of asset allocators and demands a new focus on the sustainability and impact of the unfolding technological revolution.

Chris Jeffery, Head of Macro,
Co-Head of Asset Allocation

Our clients will judge firms less by how much information they possess, and more by how effectively they turn intelligence into better decisions, faster adaptation, and more resilient investment outcomes.

Will Riley, Head of Solutions

Capital spending related to the AI build-out is supporting turbocharged profits, and reshaping both equity and credit indices

Theme	Thesis	Portfolio Guidance
Scaling and AI infrastructure	The buildout shows no ceiling yet; compute, power and memory remain the binding constraints, sustaining returns on infrastructure spend.	Comfortable keeping an AI skew in listed equities, in light of expected gains from diffusion and continued scaling.
Labour and productivity	AI has yet to displace labour at scale; early productivity gains are being absorbed by firms rather than paid out in wages.	Remain constructive on growth assets overall, owing to expected future gains from AI adoption.
Value chain and profitability	The value chain is clear but the winners are not; margins sit upstream while labs still burn cash and open-weight models pressure pricing.	Balanced exposure across infrastructure, datacentres and labs rather than a single 'winner'; avoiding concentration in any one layer.
Inflation, energy and scarcity	Scarcity in chips and (increasingly) power makes the boom inflationary; energy and permitting are the real rate-limiters in the West.	Infrastructure, energy and grid bottleneck plays; real-asset-linked and inflation-aware exposures; avoid longer-dated bonds.
Geopolitics and standards	US–China rivalry over Taiwan, chips and AI standards raises tail risk; AI is now an instrument of foreign policy.	Diversify chip-supply concentration risk; hedge geopolitical tails; selected non-US and market-neutral alpha.

Key Risks

The value of an investment and any income taken from it is not guaranteed and can go down as well as up, and the investor may get back less than the original amount invested.

Important Information

The views expressed in this document are those of Legal & General Investment Management Limited and/or its affiliates ("LGIM"- or "L&G", "we" or "us") as at the date of publication. This document is for information purposes only and we are not soliciting any action based on it. The information above discusses general economic, market or political issues and/or industry or sector trends. It does not constitute research or investment, legal or tax advice. It is not an offer or recommendation or advertisement to buy or sell securities or pursue a particular investment strategy. Past performance should not be taken as an indication or guarantee of future performance and no representation, express or implied, is made regarding future performance.

No party shall have any right of action against L&G in relation to the accuracy or completeness of the information contained in this document. The information is believed to be correct as at the date of publication, but no assurance can be given that this document is complete or accurate in the light of information that may become available after its publication. We are under no obligation to update or amend the information in this document. Where this document contains third-party information, the accuracy and completeness of such information cannot be guaranteed and we accept no responsibility or liability in respect of such information.

This document may not be reproduced in whole or in part or distributed to third parties without our prior written permission. Not for distribution to any person resident in any jurisdiction where such distribution would be contrary to local law or regulation.

© 2026 Legal & General Investment Management Limited, authorised and regulated by the Financial Conduct Authority, No. 119272. Registered in England and Wales No. 0209 1894 with registered office at One Coleman Street, London, EC2R 5AA.

Unless otherwise stated, references herein to "L&G", "we" and "us" are meant to capture the global conglomerate that includes:

- **European Economic Area:** LGIM Managers (Europe) Limited, authorised and regulated by the Central Bank of Ireland as a UCITS management company (pursuant to European Communities (Undertakings for Collective Investment in Transferable Securities) Regulations, 2011 (as amended) and as an alternative investment fund manager (pursuant to the European Union (Alternative Investment Fund Managers) Regulations 2013 (as amended)).
- **Japan:** Legal & General Investment Management Japan KK (a Japan FSA registered investment management company).
- **Hong Kong:** issued by Legal & General Investment Management Asia Limited which is licensed by the Securities and Futures Commission.
- **Singapore:** issued by LGIM Singapore Pte. Ltd. (Company Registration No. 202231876W) which is regulated by the Monetary Authority of Singapore.

The L&G Stewardship Team acts on behalf of all such locally authorised entities.

D012848