Investment Manager's Report

This report forms part of the Strategic Report section

Ben Rogoff
Partner, Technology

Alastair Unwin
Partner, Deputy Fund Manager

Market review

Equity markets delivered a third consecutive fiscal year of gains in the 12 months to the end of April 2026, with global and US indices both returning 28.8%, in sterling terms, including dividends. This capped an extraordinary three-year run in which the S&P 500 doubled from its October 2022 lows. Most major markets posted strong fiscal year 2026 (FY26) returns, flattered by April 2025 volatility around Liberation Day tariff uncertainty and the subsequent relief rally as reciprocal tariffs were paused.

Having entered the year with near-record country (US), sector (IT) and stock ('Magnificent Seven' (Mag7)) concentration, FY26 finally rewarded diversification – Europe (+21.6%), Japan (+28.1%) and South Korea (+148%) all delivered strong returns. Much of the ex-US performance reflected multiple expansion as the long-standing US premium, even adjusted for growth, narrowed. In the US, company valuations remained elevated versus history and returns were almost entirely driven by earnings. The dollar weakened by 1.7% versus sterling during the fiscal year, while the trade-weighted US dollar spot index (DXY) declined by 1.4%.

While US stock-level returns were narrow, sector returns showed greater breadth. Technology still delivered strongly in the US, while financials and miners staged a powerful recovery in Europe, partly as AI capital expenditure (capex) spilled into infrastructure-related areas. Rising single-stock volatility also gave active managers greater scope to add alpha.

FY26 was defined by robust global growth, central bank interest rate cuts, upside to AI capex and growing optimism over the broader tailwinds from AI adoption. Markets climbed the proverbial 'wall of worry' as tariff concerns and erratic policy kept uncertainty levels elevated.

The One Big Beautiful Bill Act (OBBBA), passed in early July 2025, delivered substantial stimulus via corporate tax cuts and investment incentives. Financial conditions stabilised as Middle East tensions eased, inflation moderated and labour markets cooled, prompting the Federal Reserve (Fed) to resume cutting interest rates in September. Consumer spending proved resilient, led by higher-end consumers, and the labour market held up despite tariff and policy uncertainty. Inflation ran slightly above the Fed's 2% target, but did not re-accelerate as feared, allowing three further 25 basis point (bp) rate cuts in the second half of 2025.

Late summer and autumn brought fluctuating rate expectations and a prolonged US government shutdown, triggered when Congress missed the 30 September appropriations deadline. At 43 days, it was the longest in US history, ending on 12 November. Sentiment was tested by pockets of credit stress and renewed tariff threats, particularly around Chinese rare earth export controls. Fed Chair Jerome Powell publicly clashed with President Trump over the pace of cuts and Fed independence. Momentum stalled in Q4, though equities held up, supported by resilient company earnings, easing inflation, supportive monetary policy and continued strength in AI-related capex.

Early 2026 saw a sharp rotation in equity markets from growth to value, secular to cyclical, asset-light to asset-heavy, US to ex-US. AI disruption risk came into acute focus in late January as weaker-than-expected results from Microsoft, SAP and ServiceNow were poorly received against an inflection in coding capabilities from Anthropic's Claude Code. Terminal value concerns caused technology stock valuations to fall sharply, driving the software sector's forward price to earnings (P/E) ratio from 35x to below 20x (i.e. investors previously willing to pay £35 for every £1 of earnings are now paying less than £20).

AI concerns spread to data-intensive industries and intermediaries – media; information services; brokers; business services – and Goldman Sachs' AI-at-risk basket was down 22% year to date by mid-February. Even traditionally resilient, high-quality businesses weren't spared, posting their steepest 12-month losses since the pandemic. Conversely, cyclicals and consumer industries

rallied as the ISM Manufacturing Index turned positive and the premium investors would pay for perceived 'AI defensibility' expanded. This resulted in an increase in market breadth and dispersion. The consumer staples sector – seen as sheltered from AI – reached the same valuations as technology, a sector that delivered one of its weakest starts to a year in five decades.

Geopolitical uncertainty rose sharply in the final months of the fiscal year. A US military buildup in the Middle East culminated in coordinated US and Israeli strikes on Iran in late February, raising concerns about wider regional conflict and disruption to traffic through the Strait of Hormuz. The conflict injected significant uncertainty around inflation, growth and the path of monetary policy, driving demand for safe-haven assets as the oil price rose to a multi-month high. The final weeks delivered a sharp reversal. On 8 April, President Trump announced a two-week ceasefire which held in some form through to the end of the month, and investors rushed back into AI stocks despite tighter financial conditions and a less favourable rate outlook.

Technology review

After a volatile start to the calendar year navigating DeepSeek and Liberation Day tariff shocks, the technology sector (as per the Company's benchmark, the Dow Jones Global Technology Net Total Return Index, sterling adjusted) led the market's recovery during the Company's financial year, returning +55% for the 12 months to 30 April 2026. Large and small-cap technology companies both delivered strong returns – the Russell 1000 Technology Index +51.7% and Russell 2000 Technology Index +63.8% – while the Bloomberg Magnificent 7 Total Return Index and NASDAQ 100 lagged at +46.8% and +38.8% respectively. Concerns over the impact of higher AI investment on margins and cashflows, plus uncertain returns on investment (RoI), drove a more two-way debate on many Mag7 constituents.

AI again dominated investor attention, market returns and portfolio outcomes, with the continued improvement in model capabilities the most important driver of the investment narrative. This culminated in a step-change in enterprise adoption and AI revenue growth into calendar year 2026 that surpassed even the most bullish expectations. This was all against a backdrop of the highest US tariff regime in nearly a century and a sharp valuation reset in incumbent application software.

The clearest immediate financial implication of AI model improvement was accelerating hyperscaler (the largest cloud companies) investment in AI infrastructure and the rapid consumer and enterprise adoption it enabled. ChatGPT grew from 500 million weekly active users at the start of the fiscal year to 900 million by year-end. Consensus 2026 hyperscaler capex rose from $314bn at the start of 2025 to $751bn by April 2026 – up 140% in 16 months.

OpenAI and later Anthropic struck a series of large multi-year compute commitments with Oracle, Microsoft, Amazon and Google, alongside neo-clouds including CoreWeave and Nebius. Anthropic completed a $30bn round of institutional funding, its seventh successful round, in February 2026 at a $380bn post-money valuation; OpenAI closed a $122bn round in March 2026 at $852bn. Cross-investments involving AI labs, hyperscalers and semiconductor chip providers raised concerns about 'circular financing', though the arrangements were small in scale, transparent and strategically justifiable. Concerns eased as AI token demand and revenue came through well ahead of expectations.

Extraordinary capex numbers, compute commitments and equity investments reflected their confidence in continued scaling laws and the size of the AI opportunity. Late 2025 and early 2026 saw a significant inflection in model capabilities: Google Gemini 3 (November 2025), Anthropic Claude Opus 4.5/4.6/4.7 (November 2025-April 2026) and OpenAI GPT-5.4/5.5 (January-March 2026). The step-change in model performance, alongside major improvements in agentic harnesses (orchestration layers around foundation models that provide planning loops, memory, execution environments and guardrails), plus the ability to call deterministic software tools where required, enabled models to pursue multi-step tasks autonomously rather than simply respond to individual prompts and expanded the length of time a model could work autonomously by orders of magnitude. OpenClaw – an autonomous open-source personal AI agent – went viral in January 2026 and became the fastest-growing open-source project in history.

These developments marked the true arrival of agentic AI and a critical inflection point: the shift from AI assistance to execution, as AI models began doing economically valuable work autonomously. Enterprise adoption promptly accelerated: according to Ramp, token consumption among its customers grew 14x between January and April 2026. The most striking proof point was Anthropic's annualised revenue, which rose from $9bn at the start of 2026 to $30bn by early April.

Repeated positive revisions to hyperscaler capex flowed through to AI infrastructure suppliers. Demand accelerated beyond the industry's ability to supply it and multiple bottlenecks emerged across the new AI computing architecture. The Philadelphia Semiconductor Index (SOX) returned 145% as compute demand inflected, particularly as agentic workloads grew exponentially. Semiconductor equipment manufacturers saw record orders and backlog, TSMC reported industry-leading utilisation and raised its long-term AI revenue growth targets, and sentiment

Investment Manager's Report continued

improved on a reinvigorated Intel under new leadership supported by fresh capital and a mandate for domestic chip production. Ongoing China export uncertainty remained a headwind, but demand proved more resilient than feared as Beijing's threat to restrict rare earth exports appeared to temper US efforts to tighten controls.

Networking stocks were pulled into the AI trade as new compute architectures required new networking topologies to deliver performance and address power and other bottlenecks. Suppliers of connectors, cables, componentry, power systems and cooling to the data centre buildout also delivered strong returns. Many hardware-centric subsectors that benefitted from the AI inflection had consolidated and were reluctant or unable to add capacity, driving higher pricing as demand exploded. Memory and storage were the best examples – clean room shortages and HBM (high bandwidth memory) conversion squeezed commodity supply, prompting customers to seek longer-term agreements. Hard disk drives also benefited as the most cost-effective solution for storing and retaining training and inference data.

Software lagged materially with the iShares Expanded Tech-Software ETF (IGV) returning -14.3%. Q3 2025 brought a negative inflection in sentiment as even mildly disappointing earnings were punished heavily. Concerns about the terminal value of SaaS (Software as a Service) businesses gained traction as model makers moved up the technology stack, coding tools proliferated and seat-based models came under pressure. Subsequent model releases redoubled fears over AI's role in replacing traditional human-facing software workflows as code became abundant, pressuring shares further in early 2026.

Internet companies also struggled as although advertising and e-commerce growth remained strong, significant capex weighed on free cashflow and AI narrative swings drove sentiment. Alphabet was a notable exception – initially weak as traffic share loss and decelerating paid click growth highlighted the ChatGPT competitive threat. It subsequently rallied when a judge in an antitrust case against Google ruled for more lenient remedies than the Department of Justice (DoJ) had sought, removing a major overhang on its Search business. Alphabet's AI narrative improved alongside an accelerating product cadence: deeper Gemini integration in Chrome and AI Search, Nano Banana image/video generation, then Gemini 3 and Project Genie 3. The market also reappraised Google's TPU (tensor processing unit) initiative as a credible GPU (graphics processing unit) competitor and Search held up far better than feared.

Meta Platforms (Meta) delivered solid AI-driven monetisation gains and continued to surprise with the scale of its capex commitments, but the failure to ship a competitive frontier model weighed on the stock. Apple posted solid headline numbers and reached a $4trn market capitalisation (cap), though Apple Intelligence remained underwhelming as a consumer-facing AI proposition. The iPhone refresh cycle proved more durable than feared, partly offsetting unresolved regulatory threats to the services business.

Portfolio performance

The Company significantly outperformed its benchmark during the period, with net asset value (NAV) per share increasing by +102.2% during the financial year versus +55.0% for its benchmark, the sterling-adjusted Dow Jones Global Technology Net Total Return Index.

The Company's share price rose +109.0% over the period, reflecting both strong growth in the value of its underlying assets and active share buybacks that helped close the gap between the share price and asset value. That gap – the discount – narrowed from -11.3% to -8.3%. The Company bought back 55.8 million shares during the year, more than the 38.9 million repurchased the year before, leaving 1,114.2 million shares in issue at 30 April 2026.

While naturally reluctant to celebrate periods of strong performance, we should acknowledge a remarkable year, even if absolute returns were flattered by a low starting point given the weak close to the prior financial year. The Company's relative performance was among the strongest in its history and its outperformance versus peers – already strong across most timeframes – has extended meaningfully, evidenced by top decile performance over one, three, five and 10 years versus its Lipper peer group.

The dominant driver of relative performance was our 'AI maximalist' positioning, which aligned tightly with accelerating AI adoption, surging infrastructure investment and a sharp bifurcation between AI 'haves' and 'have-nots' across the technology sector.

While significant underweight positions in Microsoft (+417bps) and Apple (+360bps) were among the largest stock-level contributors, the breadth of our AI-related exposure across multiple subsectors delivered considerably greater alpha. Stock selection contributed positively across every market-cap sector and major geography, with outstanding performance in small and mid-cap stocks.

The semiconductor sector and its supply chain remained central to the AI story. AI capex drove a surge in HBM demand, which spilled over into firmer commodity DRAM and NAND pricing. The result was extraordinary returns across the memory complex over the reporting period, most notably SanDisk (+374bps from a +3,257% stock return), SK Hynix (+140bps) and Micron Technology (+95bps). Higher data storage requirements and strong

hyperscaler capex meeting tighter supply also drove strong returns in hard disk drives, an area the Company had not held meaningfully for years, but, applying our 'AI lens', revisited during the year. This produced significant positive contributions from Seagate Technology (+199bps) and Western Digital (+107bps).

While the large (average 10.6%) but underweight (-4.1%) position in NVIDIA detracted -58bps from relative performance as we found opportunities elsewhere and tried to balance single-stock risk with enthusiasm for the AI theme, this was more than offset by broader AI-related exposure. The Company benefited from overweight positions in Advanced Micro Devices (AMD; +167bps), TSMC (+65bps) and Intel (+45bps), alongside key supply chain enablers including Elite Material (+314bps), TTM Technologies (+139bps) and Ibiden (+77bps). Semiconductor equipment and manufacturing companies also performed well on AI- and memory-related demand: LAM Research (+165bps) and KLA (+91bps), plus smaller holdings such as Disco (+32bps) and FormFactor (+50bps). Tokyo Electron (-33bps) provided a modest offset.

The Company's exposure to networking provided more than 1,000bps of relative performance. As previously discussed, the sector – encompassing chips, cables, fibre, switches and optical components – appears well positioned to benefit from the shift toward denser, higher-performance computing and the growing need to interconnect AI training clusters across multiple data centres. The theme enjoyed a sharp recovery from the depressed Liberation Day-induced lows and produced a plethora of large contributors: Lumentum Holdings (+379bps from a +1,402% stock return), Ciena (+306bps), Corning (+204bps), Fujikura (+201bps), Celestica (+178bps), Coherent (+129bps), Asia Vital Components (+127bps), Credo Technology Group Holding (+94bps) and Fabrinet (+74bps).

Closely related, the off-benchmark exposure to the data centre power and cooling theme also contributed strongly, benefiting from capex strength and the increasingly power-intensive nature of AI-optimised servers and data centres. Strong order books and earnings momentum drove positive contributions from Delta Electronics (+146bps), GE Vernova (+93bps), Vertiv Holdings (+89bps), Siemens Energy (+80bps) and Caterpillar (+36bps), more than offsetting modest detractions from other power-related holdings such as Belimo Holdings (-27bps). We leaned on our AI team's experience and domain expertise as we scaled exposure to these technology-adjacent areas.

Another major driver of relative performance was our significant underweight in the software sector, which materially underperformed during the year. This positioning reflected our out-of-consensus view that incumbent software is unlikely to prove a good conduit for AI – a view the market embraced as enterprise software multiples compressed and several leading SaaS franchises came under sustained pressure. Microsoft (+417bps) underperformed the benchmark by 53% over the period as expectations adjusted to AI-driven competition and the defensive nature of elevated capex. Underweight or zero positions in Salesforce.com (+116bps), Intuit (+85bps), SAP (+87bps), ServiceNow (+80bps), Adobe Systems (Adobe)(+72bps), Oracle (+67bps) and IBM (+61bps) all contributed positively, partially offset by smaller positions in MongoDB (-75bps) and Snowflake (-33bps).

The internet sector proved more challenging. The largest single stock detractor was an underweight in Alphabet (-235bps), reflecting a strong post-April recovery supported by a favourable DoJ outcome, successful model launches (Nano Banana; Gemini 3; Genie) and a reappraisal of the strategic value of Google's first-party chip and networking capabilities. This was partially ameliorated by the Google equity call options, which added back 122bps. The arrival of agentic AI and collapsing cost of code raised terminal value questions for Spotify Technology (-54bps), Carvana (-44bps), Netflix (-41bps), AppLovin (-37bps) and Shopify (-34bps). Robinhood Markets offered a modest offset (+44bps).

Unsurprisingly in a strong market, the largest negative contributions came from cash and the NASDAQ put options. An average cash position of around 4.8% was the single greatest drag (-277bps), while we paid away 130bps on NASDAQ put options designed to ameliorate the impact of sharp technology sector drawdowns.

As mentioned in prior reports, we view both cash and the NASDAQ put strategy within the context of an overall portfolio whose beta – owing to our growth and AI focus – remains well above one. In the prior year, the NASDAQ put options worked effectively during the January-April selloff. They also allowed us to maintain the pro-AI shape of the portfolio during early April volatility, which undoubtedly contributed to the Company's strong rebound. Finally, the use of equity call options – implemented to protect against upside risk in select stocks where we hold large underweight positions – contributed 125bps to relative performance during the financial year. Foreign exchange (FX) was a modest headwind to Company returns as sterling strengthened slightly versus the US dollar, given the majority of the portfolio's holdings are denominated in non-sterling currencies, primarily the dollar and dollar-linked currencies.

Market outlook

We consider the market outlook primarily through the lens of whether any potential macro or market shock is likely to derail the AI story. The answer so far has been 'no' – or 'not yet' – as the S&P 500's rebound from October 2022 lows has been among the strongest since 1928.

Investment Manager's Report continued

Assuming Middle East events do not cause sustained oil price increases or petrochemical shortages, US growth should remain firm as the drag from tariffs gives way to tax cuts and OBBBA (One Big Beautiful Bill Act) stimulus. Political risk is likely to remain elevated, however, especially around the midterms where a change in control of the House and/or Senate could limit the administration's ability to pass legislation and bring further geopolitical volatility to markets. Real wage growth and loose financial conditions should support consumer spending, though lower income families may continue to feel the squeeze. Cyclical indicators like the ISM Manufacturing and Purchasing Managers' indices (PMI) have picked up and tax incentives and deregulation should drive business investment.

The most important market questions for the Company's outlook are (1) whether breadth returns and (2) whether AI will prompt a market regime shift.

(1) Breadth

Technology – and associated US and growth exposure – outperformed so dramatically for so long because profit growth and returns on equity (RoE) rose substantially relative to other sectors and geographies, driven by the scaling up of cloud computing, the growth of the digital economy and increasing returns to scale. Prodigious cash generation was reinvested in high-RoIC (return on invested capital) businesses, rivals were acquired or outcompeted and the share of earnings increased.

After years of mega-cap US technology dominance, investors face a broadening set of opportunities and have thus far been rewarded for diversification. Elevated valuations, tight credit spreads and strong 'risk-on' leadership suggest there is little macroeconomic risk left to price out of broader markets so earnings growth should be a key contributor to future returns. The durability of any broadening trade ultimately depends on whether earnings breadth improves. We are optimistic conditions are in place for a recovery in market breadth as almost every industry is being reimagined in the wake of AI, though an oil shock pushing inflation higher and setting back the case for interest rate cuts makes this more challenging near-term.

NTM Earnings And Price - S&P 600 To S&P 500

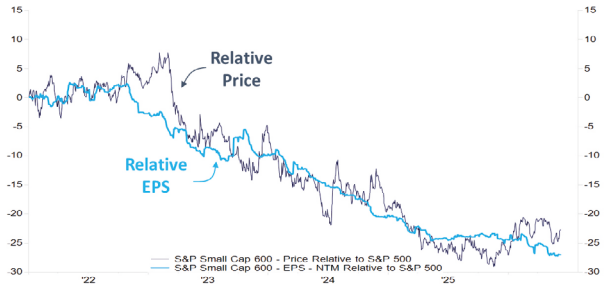

(2) Market regime shift

We have long been interested in the potential for a regime shift after many years of US 'Big Tech' dominance, though our conclusions have generally been cautious. A combination of AI and geopolitical developments is changing what the market wants. Following the global financial crisis (GFC), capital systematically rotated towards high-quality, asset-light compounders generating elevated returns on invested capital with above-trend growth. AI disrupts this framework at its foundation. By compressing the time and cost required to replicate a company's competitive advantage, it accelerates the erosion of the ability for the best companies to earn returns well above their costs of doing business. The theoretical endpoint is that exceptional profitability becomes harder to sustain, with more companies earning only ordinary returns. This could favour a longer-term shift towards asset-heavy, short-cycle, commodity-adjacent businesses already earning close to their cost of capital. Those fortunate businesses with growing competitive strengths, sitting on increasingly scarce complementary assets in an AI-first world, could see their shares meaningfully rerated upwards by the market.

AI also redefines scarcity, and asset prices will respond. Prior eras were defined by what was scarce: in the Malthusian Age, land determined the population it could support; in the Industrial Age, capital became reproducible through machines, making labour the bottleneck and enormously valuable. That is the world we still inhabit – wages are high because human intelligence cannot be replicated cheaply. Transformative AI challenges that logic. If cognitive and eventually physical labour become reproducible at marginal cost, what AI cannot reproduce appreciates in relative value. Relative prices shift. Energy, land, water, manufacturing expertise with genuine barriers to entry, authentic human experiences, proprietary data, brands and regulatory franchises all become relatively more valuable precisely because they cannot be created in a data centre. While the potential distributional and governance questions around this shift are profound and unresolved, the investment implications to us are clear: be willing to pay up for what remains scarce in a world of abundant intelligence.

Barclays Long-Short Quality Factor, 1-year & Rate of Change

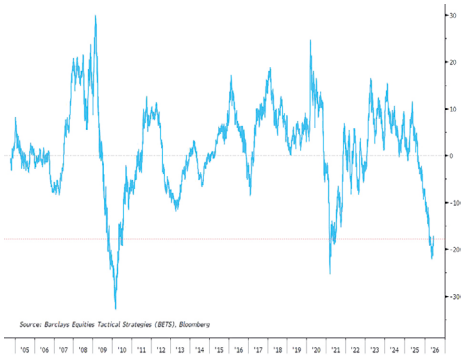

Labour market risk

The bear case to this thesis runs through the job market. If AI displaces existing work faster than new work can be created, the disinflationary effect may come via a collapse in overall demand rather than a productivity boost. A jobless expansion would exacerbate wealth inequality and likely trigger political responses such as universal basic income, AI taxes, labour protections and compute caps. The standard reassurance – that diffusion rates rather than capabilities determine labour market impact and that Goldman Sachs estimates fewer than 3% of workers are at high risk of full displacement by current AI – we find unconvincing. The major difference from prior GPT (general purpose technology) cycles is that many of the complementary changes will be provided by AI itself: electricity did not create new physical plants or reorganise work but future AI could create entirely new ways of doing knowledge work that displace whole jobs rather than automating parts of them. The labour market experience is likely to be a 'jagged frontier' where tasks, and potentially jobs, fall in scope of AI and never return. There is scant evidence of material concern in credit spreads, bond yields or inflation breakevens which is precisely why we think it is underpriced.

% of global AI adoption index vs youth unemployment change (pp) since December 2022

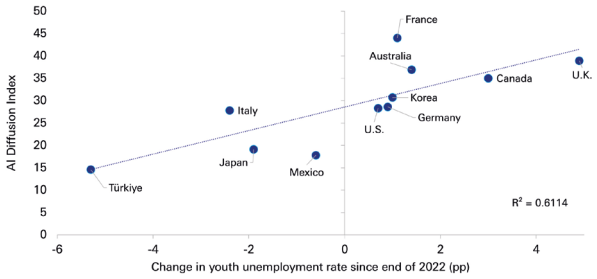

Source: Deutsche Bank

Despite this, our base case remains constructive – global growth stays resilient, corporate earnings expand and central banks ease over time. However, geopolitics has muddied the waters, and tail risks from a much more significant conflict to persistent inflation, to AI monetisation uncertainty, are more numerous and consequential than at any point since the pandemic. Our core conviction is that the AI story is just as exciting, but the investment regime is shifting.

The forces that supported 'just own US large-cap tech and Treasuries' – globalisation; falling rates; digitisation; favourable demographics – are giving way to a regime that rewards diversification, exposure to physical assets and underappreciated regions. This does not mean abandoning technology exposure but the risk/return profile of portfolios robust to a wider range of outcomes has become more attractive: more global; more diversified; more resilient.

Our bull case rests on the AI infrastructure cycle continuing and delivering greater productivity gains and economic growth than disruption in the near term, just as it did in the mid-to-late 1990s. Growth accelerates, inflation comes down and rate cuts continue, providing conditions for valuations to break higher rather than hold stable. Per the Census Bureau, AI adoption among US establishments now stands at 18.9%, with adoption among firms of >250 staff already at 34.8%. Dispersion is significant: in Q4 earnings calls, 70% of S&P 500 companies referenced AI (54% in productivity terms) but only 10% offered detailed use cases and just 1% quantified a measurable earnings impact.

US real GDP growth

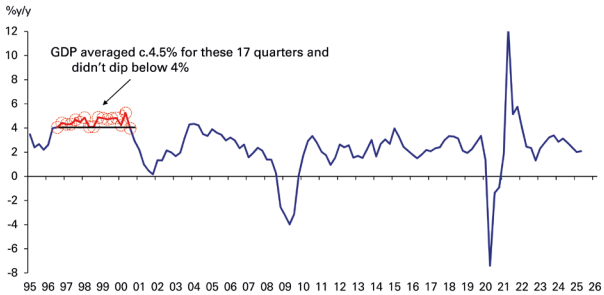

Source: Deutsche Bank

Market/macro risks

Valuations appear extended relative to history, although this has been true for several years amid improving margins and structurally high returns on equity. Equity risk premia have fallen back toward pre-GFC levels, leaving the market historically expensive and increasingly concentrated at a time when parts of its profit engine may be slowing.

US household equity allocations have also reached record levels, while elevated retail participation and growing use of zero-dated call options leave markets more vulnerable to shifts in sentiment. Such a shift could be triggered by higher risk-free rates, potentially driven by rising energy and other input costs, expansionary fiscal policies, including the OBBBA and higher defence spending, and already elevated fiscal deficits.

With just seven AI-exposed stocks accounting for roughly half of S&P 500 earnings per share growth in 2025, with a similar contribution expected in 2026, broader market performance is becoming increasingly tied to continued AI progress, particularly at a time when the largest companies may need to accelerate investment ahead of corresponding revenues.

Investment Manager's Report continued

Technology outlook

(1) AI driving accelerating IT budgets

Table 1. Worldwide IT Spending Forecast (Millions of U.S. Dollars)

	2025 Spending	2025 Growth (%)	2026 Spending	2026 Growth (%)
Data Center Systems	505,634	51.6	787,990	55.8
Devices	791,663	9.7	856,189	8.2
Software	1,254,449	12.8	1,443,621	15.1
IT Services	1,715,650	6.2	1,870,197	9.0
Communications Services	1,296,409	3.3	1,358,553	4.8
Overall IT	**5,563,805**	**10.5**	**6,316,550**	**13.5**

Source: Gartner (April 2026)
https://www.gartner.com/en/newsroom/press-releases/2026-04-22-gartner-forecasts-worldwide-it-spending-to-grow-13-point-5-percent-in-2026-totaling-6-point-31-trillion-dollars

Calendar year 2025 delivered one of the best years for IT spending in recent history at 10.5% growth, exceeding earlier expectations of 9.8% and well ahead of 7.7% in 2024. Together, these two years represent the strongest back-to-back growth since 1995-96. All the upside to Gartner's January 2025 forecast came from data centre systems, which grew 52% year on year (y/y) versus expectations of 23%. Other categories came in line or light of this.

For 2026, worldwide IT spending is forecast to remain strong at 10.5% y/y, with growth more clearly skewed towards AI as model and agentic progress drive enthusiastic enterprise adoption following disappointing early gains from copilots. Data centre systems are now expected to exceed $780bn in 2026, up from earlier expectations of $650bn. Software is forecast to accelerate to 15.1% y/y, though this includes GenAI (generative AI) model spending growing 80%. Device spending is expected to decelerate to 8.2%, partly due to rising memory prices and extending replacement cycles.

CIO surveys consistently rank AI as the top IT priority for 2026. Citi's Q4 2025 survey placed AI first, ahead of cybersecurity, digital transformation and robotics/automation. Jefferies estimates 12% of IT budgets are now allocated to AI, up from 6.5% in its previous survey, with 24% of CIOs now delivering production use-cases per Citi. Rising confidence in productivity gains likely explains why 63% of CIOs expect AI-related spending to affect hiring plans.

(2) Earnings growth

For calendar year 2026, the S&P information technology sector is forecast to deliver revenue and earnings growth of 28% and 48% respectively, well ahead of the wider market's 10.8% and 22.9% (although earnings growth is flattered by some unrealised investment gains, e.g. in SpaceX and Anthropic stakes). Outperformance is expected to extend into 2027 with revenue and earnings growth of 17.7% and 24.7% compared to 6.9% and 13.7% for the broader market. While these forecasts might appear at odds with elevated geopolitical uncertainty, the AI imperative and corporate earnings have thus far proved more resilient than feared.

S&P 500 consensus earnings growth (Y/Y): CY 2026

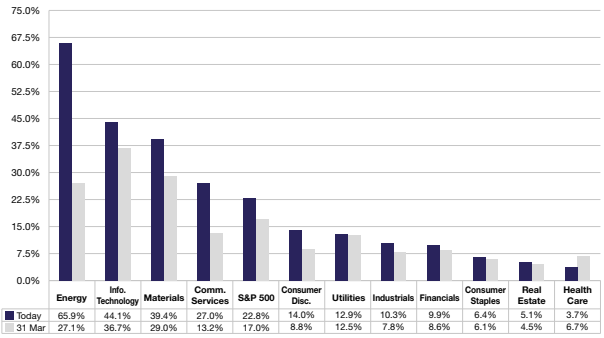

	Energy	Info. Technology	Materials	Comm. Services	S&P 500	Consumer Disc.	Utilities	Industrials	Financials	Consumer Staples	Real Estate	Health Care
Today	65.9%	44.1%	39.4%	27.0%	22.8%	14.0%	12.9%	10.3%	9.9%	6.4%	5.1%	3.7%
31 Mar	27.1%	36.7%	29.0%	13.2%	17.0%	8.8%	12.5%	7.8%	8.6%	6.1%	4.5%	6.7%

Source: FactSet

First-quarter reporting season has been supportive with blended earnings growth of 30.7% y/y, though headline numbers mask significant subsector divergence – semiconductors (+49%) doing much of the heavy lifting while others delivered more modest growth.

The most significant risks to the sector's earnings profile are geopolitics, AI disappointment and concentration risk (the Mag7's outsized share of sector earnings). Large FX moves would also have a disproportionate impact – technology has the highest international revenue exposure of any S&P 500 sector at 55% versus the S&P 500 average of 39%.

(3) Valuation

Technology's valuations expanded modestly through most of the past year, but recent pronounced weakness in software and internet stocks – groups perceived as AI losers – has driven a contraction in headline valuations. The sector now trades on a 25.1x forward earnings, down from 26.3x last year and sits between five (25.5x) and 10-year (24.3x) averages. The S&P 500's current forward P/E of (c.21x) remains elevated by historical standards, at a level seen only a handful of times in recent history (2021 and 1998-2000).

S&P 500 information technology sector forward P/E

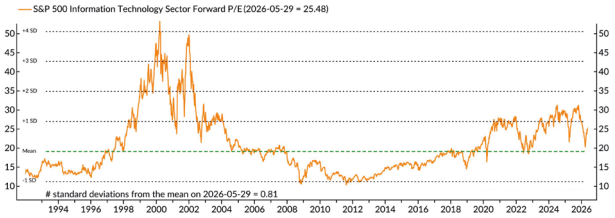

Source: Ned Davies, S&P DJ Indices and MSCI (GICS).

On a relative basis, technology's P/E trended lower for much of the past year. More recently, AI disruption fears, coupled with a robust Q4 earnings season, triggered a sharper derating. However, subsequent sector strength has seen the sector's relative P/E recover to 1.2x, its 20-year average.

(4) Repricing AI risk

Headline valuations belie record subsector divergence, as AI disruption has created clear winners and losers. Those perceived as benefitting from AI – particularly chip makers – have reached valuations not seen since early 2021 while software recently traded at its lowest valuations since 2016. The software plunge saw almost 80% of stocks suffer a >30% drop, almost unheard of outside cyclical bear markets. To us, this feels less like a sudden dislocation rather than an overdue reaction to fundamentals that have been diverging from semiconductors since at least 2020. We remain significantly underweight software and the internet on long-held concerns about their relevance in an AI-first world.

US software vs US semis: Relative price & relative revenue revisions

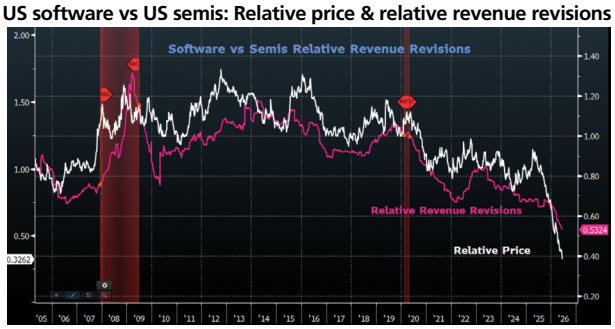

Source: Baird/Bloomberg.

(5) Mag7 challenged

After years of sharp outperformance, Mag7 returns have been moderating. In 2025, only NVIDIA and Alphabet outperformed the S&P 500, reflecting diverging AI fortunes. The group now accounts for 31% and 25% of the S&P 500 market cap and forward earnings respectively, compared to 29%/22% a year ago.

Mag 7: Market-Cap vs earnings

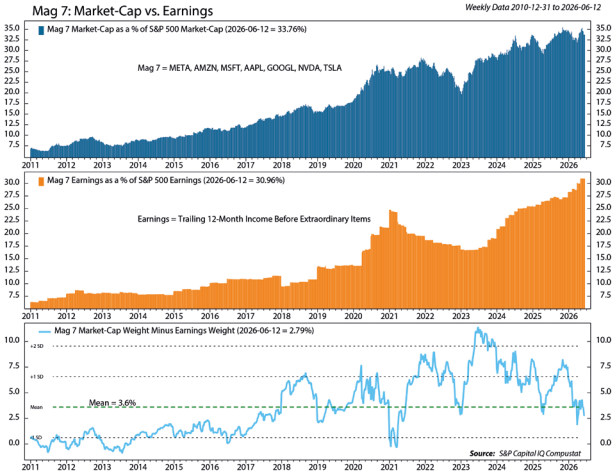

Source: Ned Davis, S&P Capital IQ Compustat.

The premium valuation enjoyed by this remarkable group of natural monopolies is now being challenged by a combination of prosaic and potentially existential factors – converging growth rates, elevated AI investment and rising capital intensity. While we believe current hyperscaler capex is rational and in long-term shareholder interests, interest alignment will be challenged near-term as hyperscalers – particularly those without frontier models – become asset-heavy 'quasi-utilities'. Against this backdrop, the absence of tangible AI benefit is likely to result in capex being increasingly perceived as defensive.

Risk to Hyperscalers: CapEx Growers Have Historically Underperformed

Source: NDR

We see agentic AI as a threat to existing cloud platforms, especially those without a frontier LLM (large language model) and/or their own chips. The LLM is the 'brain' of any agentic system. Without a frontier model, cloud platforms may cede the most lucrative part of the AI value chain – the orchestration layer – and with it application programming interface (API) revenue and control over the 'scaffolding' of core tools and management infrastructure. AI also threatens other massive profit pools – advertising (a $1trn market by 2028 dominated by Google, Meta and Amazon), e-commerce (Amazon accounts for 40% of US online retail) and software (Microsoft's 72% PC of the operating system market and 450 million productivity software users).

Investment Manager's Report continued

Competing visions

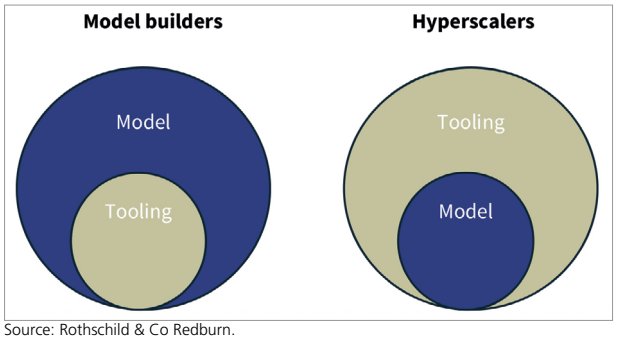

Source: Rothschild & Co Redburn.

The absence of a frontier model is a key reason we have become considerably more cautious on Meta and Microsoft, and why we meaningfully increased Alphabet exposure (which also has its own competitive chip) ahead of and after Gemini 3. There are still many ways for each Mag7 company to 'win' in the AI era, but the range of outcomes has shifted. Apple remains an outlier, lacking a frontier model but spending modestly on AI capex, and we wonder if the company might change tack once Tim Cook steps down as CEO in September.

In The Magnificent Seven (1960), only three of the seven gunfighters survive. We suspect the Mag7 trade may be approaching a similar moment and for those competing directly, they obviously cannot all win, even if the overall return on investment for AI is very positive.

Mag7 now accounts for 29% of the portfolio versus >50% of the benchmark. The burden of proof has shifted onto incumbents to demonstrate their relevance in an AI-first world. We expect to exit anything we believe to be impaired, holding significantly smaller equity positions in incumbents augmented with out of the money (OTM) call options to protect against significant rallies.

(6) Implications of a post-Mag7 world

Aside from 2022, the Mag7 has outperformed the S&P 500 in nine of the past 10 calendar years. Should recent underperformance extend, the investment ramifications would be significant – some are already unfolding.

After another exceptionally narrow market in 2025 (only 30% of S&P 500 stocks beat the index), we are hopeful the backdrop becomes more active-manager friendly. This may take the form of improved breadth – the recent low in the Russell 2000 vs S&P 500, back to levels last seen around 2000, feels like a significant portent and a potential moment for small-cap companies to reassert themselves. We are also intrigued by the potential for AI tool use to materially increase the audience for smaller businesses by helping discovery and analysis.

(7) Change of leadership?

S&P 500 Information Technology sector high-low logic index

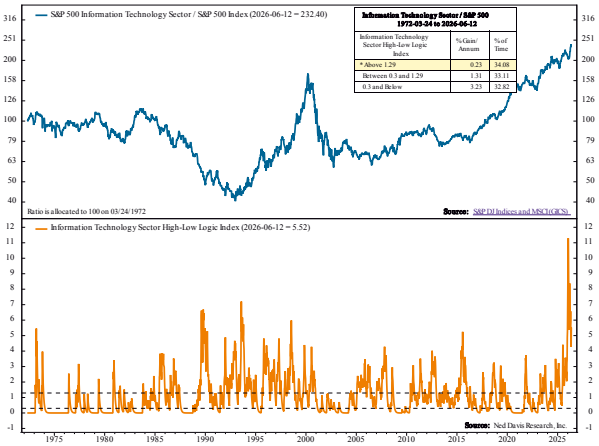

Source: Ned Davis, 20 February 2026.

We are even more excited about a potential change of technology leadership as the AI cycle extends. This long-held view was recently supported by a fascinating spike in the S&P Information Technology High-Low Logic Index, which measures the number of technology stocks simultaneously making new highs and lows. The index recently surged to 11.3, by far its highest reading since 1972, with prior peaks between 5.0 and 7.0 during the early and late 1990s. What it measures is breadth divergence: the IT sector at peak relative highs versus the S&P 500 with internal breadth is the most fractured ever. We believe this represents AI disruption, with semiconductors and infrastructure surging to new highs while software, IT services and internet break to new lows. At the headline level things look fine; beneath that, we see a sector in turmoil.

While others continue to push the dot.com parallel, the high/low reading looks more akin to the technology leadership rotations of the early 1990s PC cycle.

The PC cycle parallel

Like the AI cycle, the early 1990s PC cycle was hardware-driven and enabled by scaling laws. IBM's 1981 PC quickly became the standard business microcomputer. The fateful decision to use an Intel microprocessor was reframed by Moore's Law and Intel's merchant volumes. Intel's fifth-generation Pentium (1993) unlocked a processing throughput bottleneck much as HBM has enabled training of massive AI models, while Windows 3.1 (1992) sold three million copies in its first two months and turned the PC into a viable platform.

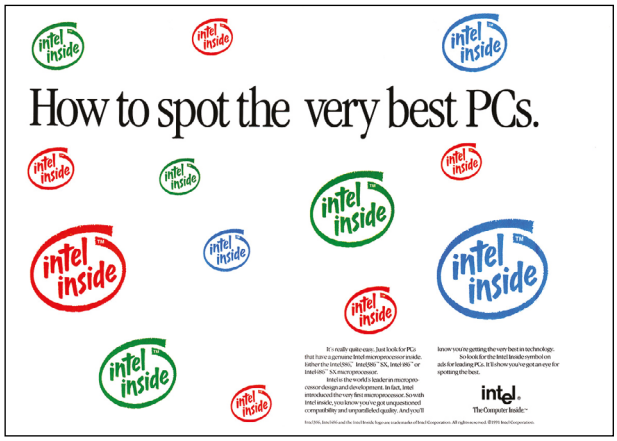

Global personal computer unit sales, 1975–2011

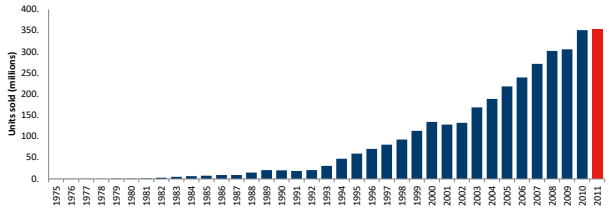

Sources: Jeremy Reimer, "Total Share" (1975–1993); Gartner / IDC worldwide PC shipments (1994–2011).

PC units grew from 24 million to 150 million between 1990 and 2000. Microsoft's revenues increased 30x between 1986 and 1995 at 25% net margins, and its market cap grew roughly 70x in nine years from IPO, overtaking IBM in early 1993 despite generating a fraction of its revenue – highly reminiscent of recent growth inflections at OpenAI and Anthropic. The incumbents, meanwhile, were in freefall: IBM posted an $8bn loss in 1993, the largest in US corporate history; DEC, the world's second-largest computer company in 1988, lost money in almost every year between 1991 and 1996; Wang Laboratories, which had controlled up to 80% of office word processing, filed for bankruptcy in 1992.

What lessons might we learn from the PC cycle?

1. Rapid leadership change: In 1993, IBM's share price halved and its market cap fell below Microsoft's for the first time, despite IBM generating $62bn in revenue to Microsoft's $3.8bn. PC and client-server winners rapidly usurped incumbents within indices, years before fundamentals 'justified' the rotation.

2. New (invisible) market opportunities: Early PC winners were later augmented by software applications (Adobe; Corel), utility vendors (Symantec; Norton), networking (3Com; Novell) and an entirely new gaming hardware layer. In 1993 – the same year IBM posted the largest loss in US corporate history – Jensen Huang founded NVIDIA and ATI (now AMD) listed in Toronto.

3. Moving up the stack: Having established Windows as the platform, Microsoft systematically moved into applications. Lotus 1-2-3 was displaced by Excel; Novell NetWare's networking dominance evaporated when Windows NT (1993) added built-in networking; WordPerfect, Borland and Netscape each dominated their category until Microsoft leveraged its platform to enter, bundle and win. Anthropic's recent moves into tool use, computer use and model context protocol (MCP) connectors look less like feature additions and more like the early stages of a Microsoft-style platform expansion from model provider toward the application layer.

4. Architectural shift: The PC cycle is usually framed as the natural successor to previous hardware cycles – mainframe to minicomputer to PC – but it represented something more fundamental: vertically integrated computing was replaced by a horizontally layered alternative where Intel owned the processor, Microsoft owned the operating system and commodity assemblers owned the box. Value migrated from integration to specialisation, hardware to software, proprietary to open. The PC democratised computing but destroyed every incumbent built on a vertical moat.

AI is following a similar pattern of misidentification. It is widely articulated as the natural successor to the cloud because both are scaled forms of computing delivered remotely. However, the underlying stack is profoundly different: parallel rather than serial, power-hungry rather than power-efficient, with model builders adopting heterogeneous approaches to architecture, training, inference and power. Before AI, the unit of compute was the server; today it is the rack and tomorrow it may be the data centre itself. AI may democratise cognition, but the infrastructure required is anything but commoditised. The AI cycle is therefore likely to look less like on-premise to cloud – which destroyed branded hardware but preserved much of the existing software stack – and far more like mainframe to PC, which destroyed the incumbents entirely.

5. Greater competition: Investors should brace for significantly greater competition. The new cycle has created a window for entrants to displace incumbents. Microsoft may have felt it covered its bases when it invested (brilliantly) in OpenAI back in 2019, just as IBM may have done when it launched its PC in 1981. However, the PC cycle shows that revenues, installed bases, brands and customer loyalty count for little when an architectural shift renders technical foundations obsolete. Lotus had the dominant spreadsheet. WordPerfect had the dominant word processor. Novell had the dominant network. None of it mattered. Current debates about near-term software model challenges look remarkably small relative to AI's likely disruptive impact on existing clearing prices and digital profit pools over time.

Investment Manager's Report continued

(8) End of US exceptionalism?

This might partly explain why US market leadership is waning. The Company is at its largest US underweight in many years, principally from reducing software/internet exposure and reinvesting in semiconductors/hardware assets in Japan, South Korea and Taiwan, as well as in the US.

However, all is not lost for US technology. America enjoys strong or dominant positions in AI chips, storage, power, optical networking and other components, and is winning the frontier model race led by Anthropic, OpenAI and Google. Beyond AI, the US holds strong positions in aerospace/defence, space, autonomous vehicles and humanoid robots.

This is likely to become more apparent as the IPO (initial public offering) market recovers following another subdued year. The outlook for issuance is genuinely exciting should OpenAI and Anthropic become public as reported. SpaceX could also ignite the IPO market while refreshing the quality and narrative of the US-listed technology universe. In our view – from a distance, given we have not seen their financials – all three are world-class, pure-play assets on critical investment themes. We are also keeping a close eye on Databricks and Stripe, both scaled mission-critical platforms.

(9) Later, but still not a bubble

After a remarkable year, some of our 'bubble' indicators have moved up a notch. However, we still feel AI stocks are not in bubble territory and the current backdrop more closely resembles the mid- rather than late-1990s. Valuations look expensive relative to history but undemanding relative to the market – and far from 2000, when technology comprised 32% of market cap but only 12% of earnings, compared to 43% of the market cap (including communication services) and 38% of earnings today. The IPO market remains far from bubble conditions – just 31 venture-capital-backed technology companies went public last year (14 in 2024), versus 370 and 261 in 1999 and 2000 respectively. The average age at listing was 14 and 12 years for 2024 and 2025 deals, rather than five years for the 1,400 companies that went public between 1995 and 2000.

Loss-making companies making IPOs shot up in the 1990s

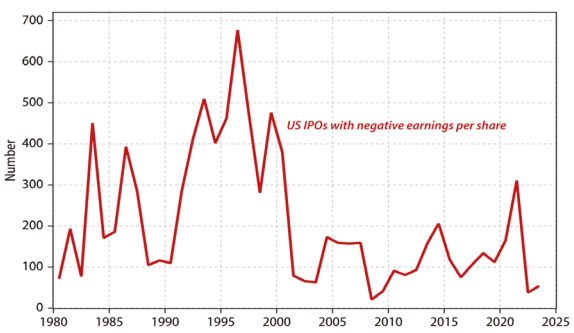

Source: Gavekal Research/Macrobond.

M&A (mergers and acquisitions) are another useful bubble indicator. Even ignoring software deals we consider defensive or private equity-related, 2025 was much busier for AI-related M&A – disclosed deal value of $120-125bn including software and security, c$50-55bn for pure hardware, infrastructure and talent acquisitions compared to $10-12bn in 2024. The 5-6x like-for-like increase looks more reminiscent of the mid-1990s when telecom and internet M&A increased 4-5x between 1995 and 1996. As a share of index market cap (up almost nine-fold since 1996), last year's activity looks even less significant and most of the largest deals were cash, not equity financed.

There has been a greater use of debt financing relative to earlier AI capex. According to the *Financial Times*, Oracle, Meta, xAI and CoreWeave have structured over $120bn of AI data centre financing through off-balance-sheet SPVs (special purpose vehicles), while UBS estimates $125bn flowed into AI-related project finance during 2025. We are not unduly concerned – most resemble sensible financing structures for long-life infrastructure assets, not dissimilar to aircraft leasing. The best borrowers continue to enjoy easy access to capital – Google's $32bn bond sale in February was 10x oversubscribed and completed within 24 hours. Nevertheless, the industry's need for external capital speaks to the scale and maturity of current capex plans. Pockets of exuberance are to be expected: in August 2025, OpenAI CEO Sam Altman called it "insane" that some tiny AI startups – "three people and an idea" – were being funded at billion-dollar valuations. While data points like this are unsettling, excitable private markets also reflect an acute shortage of AI talent.

(10) Volatility likely

We still consider current conditions more analogous to the mid- rather than late-1990s and, as previously warned, the volatility that accompanied that period should be anticipated. The NASDAQ experienced seven corrections of more than 15% between 1995 and 1998, despite delivering substantial returns. Volatility should be considered an inherent feature of transformative technological transitions. This was certainly the case in 2025, a remarkable year punctuated by DeepSeek (which saw an AI leaders basket drop 18%) and Liberation Day tariffs (which triggered a 9% S&P 500 decline, pushed the VIX – a measure of volatility in the S&P 500 – to 60 (anything above 50 is cause for concern) and wiped roughly $5trn in market cap). Realised volatility ran at 19% – the 83rd historical percentile – and the largest peak-to-trough S&P 500 drawdown was 19%, nearly double the annual median of 10%. We expect heightened volatility to persist – elevated P/E ratios, tight corporate spreads, rising government debt, unprecedented AI capital spending, broadening AI disruption and record retail participation all leave markets with less margin for error should narratives shift.

AI cycle update

(11) Rapid adoption: users and tokens

AI adoption continues to significantly outpace historic trends seen with the PC, internet and smartphone. OpenAI recently announced it had reached 900 million weekly active users (WAU) with ambitions to grow to 2.6 billion WAU by 2030 while Google's Gemini AI surpassed 750 million monthly active users (MAU). User growth has helped drive AI usage, as measured by token growth which had been compounding at 4-5x annualised before appearing to accelerate to 9x annualised growth in January. In Altman's words, reasoning and agentic have taken AI into a "new phase where frontier AI moves from research into daily use at global scale".

(12) Enterprise adoption begins

This is evident from Anthropic's revenue trajectory. Having ended 2024 at $1bn, Anthropic's annualised recurring revenue (ARR) reached a remarkable $14bn in February 2026, making it, at the time, the fastest-scaling business to business (B2B) company in the history of software. Since then, ARR exploded to $19bn in March, $30bn in April and – according to reports – $47bn in May. This is truly unprecedented; no company in history, to our knowledge, has ever reached $30bn of ARR within four years, let alone $47bn.

This revenue breakout reflects an inflection in enterprise AI adoption which had previously trailed consumer adoption due to early scepticism and inconsistent model performance. This is clearly no longer the case, with measurable productivity gains driving demonstrable RoI. Anthropic's analysis of Claude usage shows a median 84% time saving per conversation, concentrated in the 80-90% range for multi-step cognitive work. AI is collapsing the coordination cost of knowledge work, creating a self-reinforcing adoption loop: freed capacity justifies further investment in AI infrastructure and tooling.

As discussed in last year's Annual Report, agentic AI represents the architectural breakthrough that enables enterprise adoption because enterprises are built on repetitive, multi-step workflows that are too complex for a simple prompt but do not require human judgement at every step. Unlike chat, agentic AI can integrate with existing systems and complete work. By 2028, technology research firm Gartner expects 15% of day-to-day work decisions to be made autonomously by AI agents.

(13) Rapid model progress

The emergence of agentic AI reflects the extraordinary pace of underlying model progress during 2025, a year that began with the DeepSeek R1 shock and ended with an unprecedented wave of frontier releases. AI systems evolved rapidly from conversational chatbots into reasoning-driven, multimodal agents able to operate across text, image, audio and video with expanding autonomy. Frontier closed-source models delivered major advances in reasoning, hallucination reduction, multimodal capability and million-token context windows, with leading systems achieving gold medal standard mathematical performance alongside near-human or superhuman coding ability.

Open-source models also advanced rapidly, led by a wave of Chinese releases, including DeepSeek R1, that brought frontier-level performance to dramatically lower price points. Momentum accelerated further in early 2026 as systems such as Claude Code and open-source autonomous agents demonstrated how quickly AI is evolving from a passive assistant into a practical digital worker, while models such as Claude Mythos Preview hinted at levels of reasoning, autonomy and technical sophistication previously associated only with highly specialised human experts.

Investment Manager's Report continued

(14) Task-completion time horizon

Time horizon of software tasks different LLMs can complete 50% of the time

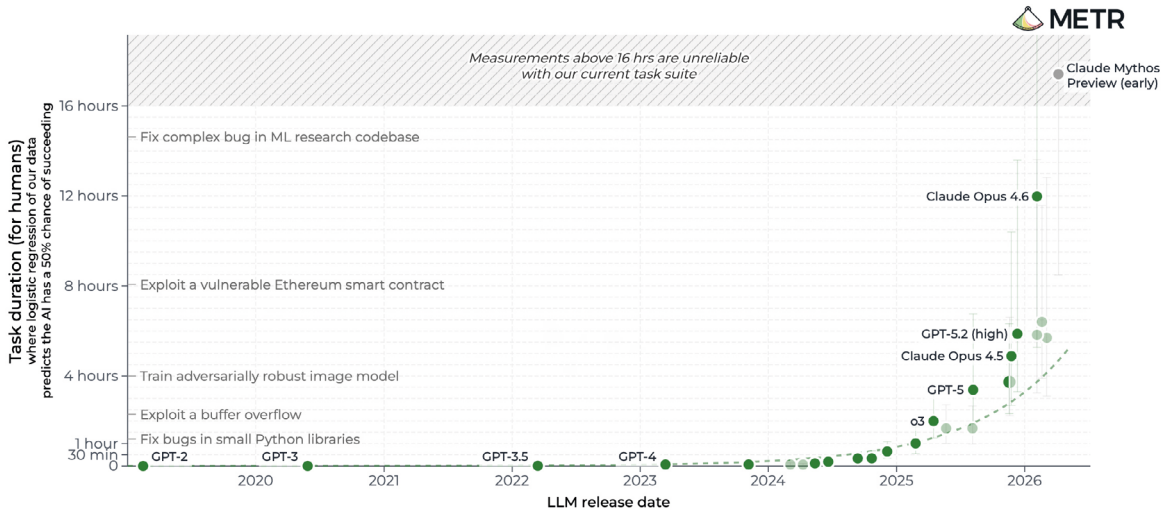

As previously discussed, agentic AI introduces the prospect of non-human scaling in cognitive labour. Accordingly, one of the most important frontier AI benchmarks today is the duration and complexity of tasks a system can reliably complete autonomously. The so-called task-completion time horizon measures the length of a task – based on expert human completion time – that an AI agent can successfully complete at a given level of reliability. A 50% time horizon represents the task duration at which an agent succeeds half the time. Until recently, the duration of tasks AI systems could complete was doubling roughly every seven months – already an extraordinary rate of progress – but since 2024 this pace has accelerated to closer to every four months. At this threshold, ChatGPT (2022) could autonomously perform tasks requiring less than a minute of human expert effort; GPT-5.2 (2025) could manage tasks exceeding six hours; Anthropic's Opus 4.6 (2026) – the model underpinning Claude Code – has reportedly extended this horizon to 14.5 hours.

(15) The coding breakthrough

Claude code: Daily GitHub Commits

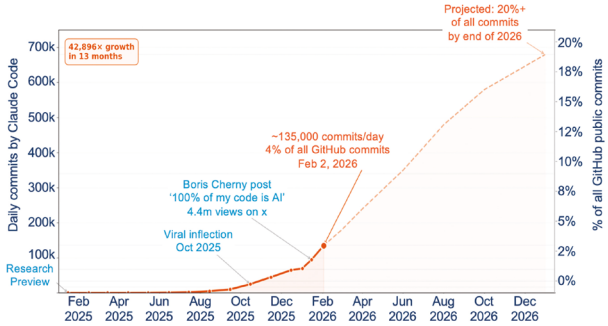

Source: Claude generated from Semianalysis.

This remarkable progress has appeared first in coding. In early 2025, Anthropic launched Claude Code, a command-line coding agent able to autonomously navigate codebases, write and edit code, run tests and commit changes with minimal human oversight. By September, Claude Code had reached $1bn ARR and, in December, Boris Cherny – its creator – stated that 100% of his commits (unit of code) were written by Claude + Opus 4.5. Daily GitHub commits attributed to Claude Code rose to 4% of all public GitHub commits, an extraordinary 43,000× increase in roughly 13 months. A commit is not code itself, but rather a snapshot of changes to a repository and one of the few observable, standardised time-series signals for software production.

While this may appear to be 'just' faster coding, it is better understood as non-human scaling and helps explain why AI has entered a more disruptive phase. The release of Cowork in January extended Claude Code to non-technical users. More importantly, as AI shifts from tool to autonomous worker, RoI follows. Although just one example, Anthropic's revenue trajectory also undermines earlier claims that AI revenues were insufficient to justify current levels of capex. Nor should it surprise us that systems capable of autonomously completing multi-hour tasks possess far greater economic value than those limited to minutes. Just as skyscrapers transformed the value of land by conquering vertical space (and thus it became economic to replace early skyscrapers with taller structures on the same site very quickly), so conquering cognitive depth is transforming the value of AI.

Gillender Building (1897) replaced by (2x taller) Bankers Trust Building (1910) after just 13 years

(16) Code as a reasoning substrate

The coding breakthrough also matters because code has become part of the reasoning process itself, rather than merely an output. Frontier models increasingly use code to externalise state, memory, iteration and precision into deterministic systems better suited to those tasks. A useful analogy is using AI to calculate a tax return. Previously, the LLM had to perform the entire calculation within token space: interpreting the tax code, applying thresholds, carrying figures forward and holding all dependencies simultaneously in memory, with small reasoning errors compounding along the way. Today, the model is not doing the tax return directly – it writes the system that does the tax return. The reasoning and synthesis still come from the model, but calculation, state management and iteration are offloaded into Python. By externalising or offloading cognition, a probabilistic LLM produces deterministic, repeatable output.

(17) Capex and the AI race

In his 2020 shareholder letter, CEO Jeff Bezos quantified the value of Amazon Prime in terms of time saved. He estimated Prime saved members more than 75 hours annually versus physical store trips, valuing that at roughly $630 per member against a subscription price of $119. Time is also a useful lens through which to view AI progress, given that many of the key benchmarks are denominated in time, including autonomous task duration and the acceleration of human work. Applying this framework to the developed-world knowledge wage bill ($23trn), assuming half is addressable by current AI and using Anthropic's estimate of 80% time savings, generates $9-10trn in equivalent labour value. Applying an Amazon-style 16% capture rate suggests a potential AI revenue opportunity more than $1trn today – and this captures only the current visible market opportunity. In 1998, the internet addressable market was framed as a share of the existing $500bn global advertising market; today, e-commerce, digital advertising and video streaming together exceed $7trn. AI may unlock similarly large invisible markets that are difficult to frame today.

(18) Looks like convergence…

From a distance, AI models appear to be converging. On benchmarks such as MMLU (massive multitask language understanding), leading systems now cluster above 85%, while open-source models often launch at 90% of frontier performance and close much of the remaining gap within months. Chinese models are believed to trail US frontier labs by less than a year, while architectures themselves are converging around mixture-of-experts (MoE) designs. DeepSeek's release of R1 in January 2025 crystallised the commoditisation narrative: it matched OpenAI's o1 on many benchmarks while charging just $2.19 per million output tokens versus o1's $60. Today, open-weight models such as Qwen3 and Meta AI's Llama 4 offer near-frontier capability at close to zero marginal cost.

This convergence, however, is largely a mirage. Static benchmarks obscure widening gaps in economically valuable capabilities. On SWE-bench Pro – a demanding real-world coding benchmark – frontier models such as Claude Opus and GPT-5 solve more than 20% of tasks, while open-weight Qwen3 32B achieves just 3.4%. Frontier systems also reason differently, exhibit distinct strengths and produce materially different outputs. Unlike compute or storage, cognition is not fungible. That helps explain why closed-source models still account for roughly 80% of token usage and 96% of industry revenue despite costing, on average, eight times more per token than open alternatives.

Investment Manager's Report continued

(19) Front-runners breaking away from the pack

Today, the frontier appears increasingly concentrated around three companies: Anthropic, OpenAI and Google DeepMind. Everyone else is either a tier below or reliant on one of these for their intelligence layer. That these three companies' models are clustered or trading positions in benchmarks is not evidence of commoditisation. It means that only they are today able to compete at this level. Meta is spending enormous sums, but Llama is today a strong open-weight model, not a frontier one. Amazon Nova competes on cost, not capability. Smaller and open-weight models can match frontier performance on isolated tasks but degrade sharply when each step depends on the previous one.

This heightened level of intensity is evident in soaring capex budgets, high-profile talent wars and faster model iteration. In less than a month during late 2025, four frontier labs launched their most powerful models: Grok 4.1, Gemini 3, Claude Opus 4.5 and GPT-5.2. Since then, the pace has only accelerated. The leading labs are locked in a pace that is self-reinforcing: every capability jump resets the field and compresses the interval to the next one, with perceived leadership shifting with each new model breakthrough. Huge differences in token prices for the most capable models indicate the market reality of the returns to frontier intelligence.

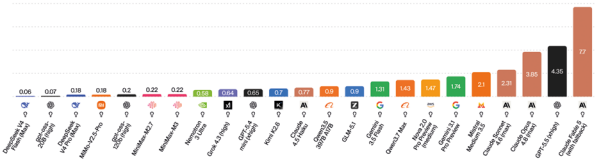

The AI race is one the smartest companies know they cannot afford to lose. Now that agents are viable, it should be clear that the LLM increasingly resembles the operating system of the agentic world – the layer that interprets intent and orchestrates execution. Any company without a frontier model is building on someone else's operating system, exposed to margin compression, feature absorption and – in an age of abundant code – rapid obsolescence.

Figure 5....now $790bn

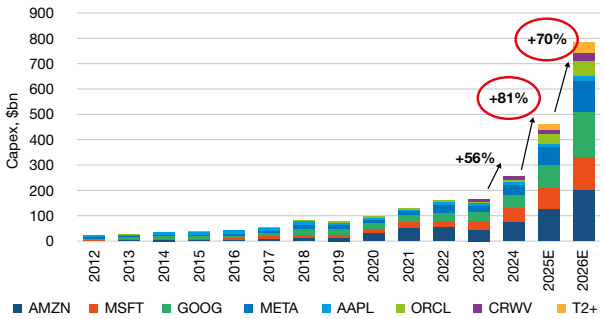

Capital expenditure program as share of GDP
Annual average over program duration or year if specified

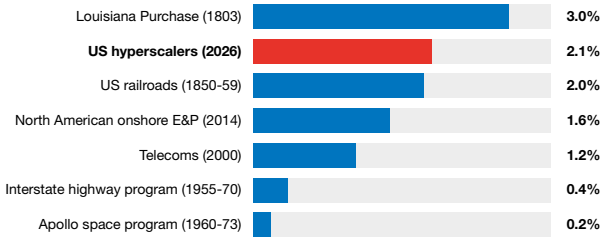

Source: Companies World Bank, Wall Street Journal and S&P Global.

That is why capex budgets can be both historic and entirely rational. Following an extraordinary 2025, AI capex is now expected to approach $800bn (+70% y/y) this year. Amazon, Alphabet, Meta, Microsoft and Oracle alone plan to spend as much as $690bn (+60% y/y) from already historic 2025 levels, while China is also scaling rapidly with $125bn spent on AI capex last year. Between the two blocs, global AI infrastructure spending is approaching $1trn annually. The conviction behind it is simple: the industry is supply constrained, AI workloads will consume every available unit of compute capacity and the penalty for underbuilding is existential. The cost of staying in the race is also compounding, with each generation of frontier model requiring roughly an order of magnitude more compute alongside a rapidly growing inference burden – inference workloads already consume over 55% of AI-optimised infrastructure spending in 2026, up from a third in 2023.

(20) Who could leave the table?

While each participant in the AI race has their own motivations, history suggests that as the ante rises the number of players thins out. The xAI/SpaceX merger – while framed around the vision of orbital data centres – may have been a strategy to ensure that xAI had access to the capital required to stay in the AI race. Oracle plans to raise $45-50bn in debt and equity this year having spent like a hyperscaler without the balance sheet of one. At the same time, others could yet join the game: Microsoft could enter the frontier race having relied on OpenAI for model access, with Microsoft AI CEO Mustafa Suleyman explicit in wanting Microsoft models "at the absolute frontier". Apple also remains a wildcard, with $130bn in cash and its CEO due to step down in September. As for the AI labs, both continue to raise capital relatively easily: in February, Anthropic closed a $30bn funding round while OpenAI raised $110bn, the largest private financing in history.

(21) Disruption, accelerated

In today's digital economy, attention scarcity is monetised at every layer. Instead of the so-called 'long tail', we created vast new digital gatekeepers that helped humans make informed choices from an almost infinite range of alternatives. Data exploded 32-fold in a decade, helping to defend, optimise and scale the digital winners into the natural monopolies that many are today. However, AI represents three principal risks to these profit pools: abundant code, non-human actors, and a natural language interface.

Abundant code: The modern digital economy rests on roughly 2.8 trillion lines of code written over the past two decades, of which perhaps one trillion lines remain actively maintained. In 2010, the world produced around 33 billion lines of new or modified code per year; by 2024, AI coding assistants were producing 256 billion lines, with 41% of all new code now AI-generated. Today, tools like Claude Code and OpenAI's Codex can independently write software, improve existing code and even submit completed updates with minimal human involvement. The cost per line of code has followed: from $7-15 pre-AI, to $2-4 with assistants, to sub-$2 with agentic systems – and falling, as token prices decline at roughly 90% per year. The digital platforms were built when code was scarce and expensive. Today, the marginal cost of code is converging on zero.

Non-human actors: Pre-AI digital profit pools were earned by solving a real problem: helping humans navigate infinite choice and monetising the time/value trade-off at the heart of every digital interaction. 'Time is money' has underpinned two centuries of economic logic, and the entire margin structure of the digital platforms, but agents have infinite time. A human will pay for a shortcut; an agent will simply do the work. That distinction matters because the moats that took a decade to build – network effects; data flywheels; switching costs – were durable in a world where iteration was slow, learning was expensive and time was scarce. Agents face none of these constraints. They operate on a different efficient frontier, optimising for compute costs, not time.

Natural language interface: Every previous interface revolution made software easier to use, e.g. command line to graphical user interface to touchscreen. Each also diminished the value of incumbency, breaking habits and rendering previous workflows obsolete. However, the LLM is significantly more disruptive. It responds to intent, not instruction, unbundling the human operator from the application itself and making software redundant as an interface layer.

Abundant code lowers the barriers to entry. Non-human actors erode the willingness to pay. The natural language interface unbundles the user from the application. Together, they represent a profound challenge to many digital incumbents.

Investment Manager's Report continued

In information services, AI can approximate datasets that were previously valuable because they were hard to assemble rather than fundamentally secret. The moat was never the data, but the human labour required to collect, synthesise and curate it. AI collapses that labour by 80-90%. In software, three moats look challenged – data trapped inside the application, workflow lock-in from learning the user interface (UI) and integration complexity between systems. In an AI-first world, the application becomes a backend service that the LLM orchestrates, not a product that a human operates.

Advertising – worth $1trn by 2028 – also appears AI-exposed. Performance advertising is today built on attention scarcity with platforms monetising the space between intent and action. Agentic commerce may also reshape e-commerce – Morgan Stanley estimates agentic gross merchandise value could reach $190bn in a base case and $385bn in a bull case by 2030, implying 10-20% of US e-commerce. Content is likely to be highly disrupted too, as high-quality AI-assisted video content continues to improve and become significantly cheaper. Last year, UK creative agency headcount fell 14% year-on-year, the steepest decline since records began. Within five years, McKinsey estimates $60bn of content revenue could be redistributed.

These areas potentially at risk from disruption are not exhaustive. Any business whose competitive advantage exists purely as executable code sits on a different risk curve. Disruption risk extends well beyond code; agentic AI systematically attacks confusion, inertia and opacity – three of the most profitable features of the consumer economy today.

Technology risks

As we have previously outlined, the principal technology risk to the AI cycle and our portfolio is that the pace of model progress, adoption or monetisation fails to justify the historic level of infrastructure investment currently underway. Much depends on continued confidence in scaling laws; however, these are empirically observed rather than theoretically derived and there is no fundamental reason they must continue indefinitely. Data scarcity, diminishing returns from pre-training or an architectural breakthrough that invalidates the current paradigm could each flatten the current trajectory.

Regulation adds a further layer of uncertainty. The EU AI Act is now in force, US state-level legislation is proliferating and unresolved copyright litigation could yet establish precedents that materially alter the economics of frontier model development. More broadly, sufficiently rapid advances in frontier capability could provoke direct state intervention, including licensing regimes, restrictions on model releases or greater government oversight of frontier labs, any of which could alter how AI value accrues across the ecosystem.

There are also risks associated with the financing of this investment cycle. While we are broadly comfortable with the increasingly interconnected capital flows between industry participants – which we view as a form of vertical integration in a stack no single company fully controls – they nevertheless create system-level fragility. The structure depends on AI revenues materialising on a timeline compatible with capital deployment. JP Morgan estimates cumulative AI capex requirements of $5.3trn through 2030, with a $1.4trn financing shortfall emerging even after hyperscaler cashflows, investment-grade bond issuance, securitisation and leveraged finance are fully utilised. Private credit is expected to bridge much of that gap, although recent indicators have been mixed.

Hyperscaler balance sheets are also becoming materially more capital intensive: the roughly $2trn of AI assets planned by the five largest hyperscalers could imply annual depreciation expense of $400bn by 2030, greater than their combined 2025 profits. Except for Alphabet, none of the hyperscalers currently controls a leading frontier model, meaning current spending increasingly resembles a Red Queen dynamic in which participants must keep running simply to maintain competitive positioning.

The AI cycle is also exposed to a range of infrastructure and supply-chain risks. AI demand remains constrained by the availability of power, networking, advanced semiconductors, memory, packaging and data centre capacity, leaving the industry vulnerable to bottlenecks, delays and cost inflation. Power may prove the most underappreciated constraint. Training and serving frontier models increasingly requires gigawatt-scale infrastructure. US grid interconnection queues are measured in years, and new gas turbine delivery times now extend to 5-7 years.

Geopolitical tensions represent a further material risk, particularly around Taiwan and US/China relations, given Taiwan's dominant role in advanced semiconductor manufacturing and China's importance across both supply and demand. Export controls, tariff escalation, restrictions on rare earth minerals or disruption to semiconductor supply chains could materially affect AI infrastructure deployment, sector valuations and, by extension, our 'AI maximalist' portfolio positioning.

Our AI bull case: cognitive abundance

The industry's most prominent figures have made predictions that would once have sounded absurd. Anthropic CEO Dario Amodei has argued that AI-enabled biology and medicine may condense 50-100 years of human progress into 5-10. In his recent essay *The Adolescence of Technology*, he defines "powerful AI" as a system smarter than a Nobel Prize winner across most relevant fields, able to carry out tasks autonomously, with enough compute to run millions of instances of itself – a "country of geniuses in a data centre" – and in 2024 suggested this could be as little as 1-2 years away. DeepMind CEO Demis Hassabis has spoken of drug discovery becoming a thousand times more efficient while Sam Altman has suggested AI could compress a decade of scientific progress into a single year. While these claims echo Ray Kurzweil's long-standing view that technological progress compounds far faster than human intuition allows, they also assume **cognitive abundance:** AI cheap and continuous enough to put billions of minds to work in parallel.

There is a precedent for abundance turning a known resource from input to substrate. By 1901, oil - refined into kerosene - lit the world's lamps. While it had already made Standard Oil one of the most valuable companies on earth, everything that moved the economy - the trains, the ships, the factories - ran on coal. For the same heat, oil cost four times as much.

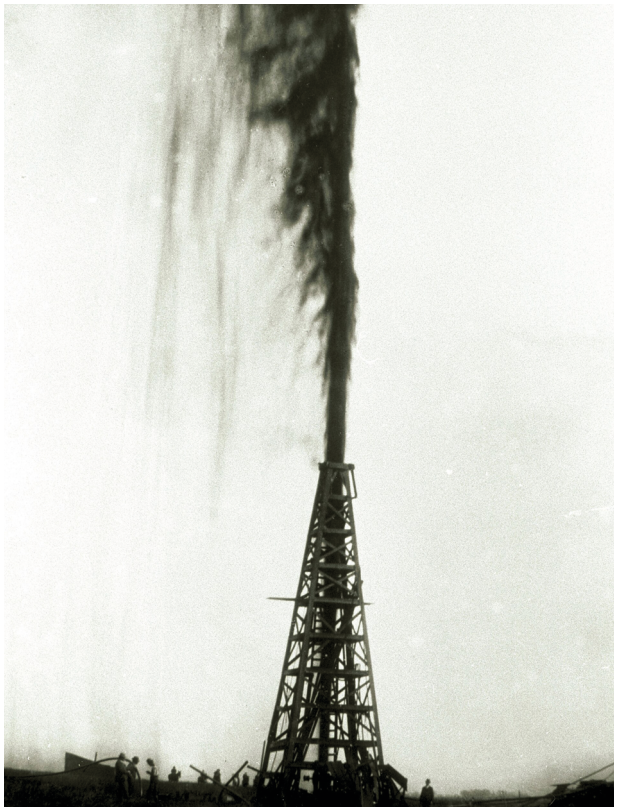

Spindletop changed that. On 10 January 1901, the Lucas gusher on Spindletop - a salt-dome hill in Texas - blew oil 150 feet into the air at 100,000 barrels a day, more than every other US well combined at the time. The price collapsed to three cents a barrel – for a time, cheaper than water - and oil became something to burn rather than merely to light. Within the year, railroads and steamships began converting from coal, and within a decade a lamp fuel had become the substrate of the industrial world.

Today, intelligence is the kerosene of the modern economy. The breakthrough in task duration may be AI's Spindletop. If so, the claims of Altman and his peers are not so fantastical — and the question becomes not whether to dismiss them, but what would need to be true for them to be even directionally right.

The easiest way to underestimate AI is to mistake the current interface for the underlying capability. ChatGPT is to AI what the telegraph was to electricity – a narrow early application of a far more significant general-purpose technology. We ask AI to slot into workflows designed around scarce, expensive human attention, just as early factories bolted electric motors onto steam-era layouts. The deeper opportunity is to redesign around abundant cognition from the outset. The binding constraint today is not cost – at current token prices, even an Edison-scale invention is computationally cheap – but continuity: reliable, persistent, reality-linked cognition over long durations. That bottleneck does not look permanent. Context windows are expanding, memory architectures improving, agents are learning to coordinate tools and sub-tasks over extended horizons.

The deeper implication is institutional. No single human mind could hold the Apollo space programme end-to-end; tens of thousands of people each understood a fragment – heat shields, guidance software, orbital mechanics etc made compatible through shared standards, hierarchies and records. Bureaucracy, in Weber's sense, was itself a form of externalised rationality: the historical solution to cognitive scarcity was to break problems into human-sized pieces and build institutions to coordinate them. AI does not merely substitute for the specialists working on those fragments. As context, memory and agency scale, it begins to encroach on the institution's deeper function – the coordinating intelligence that keeps the whole problem in view.

Investment Manager's Report continued

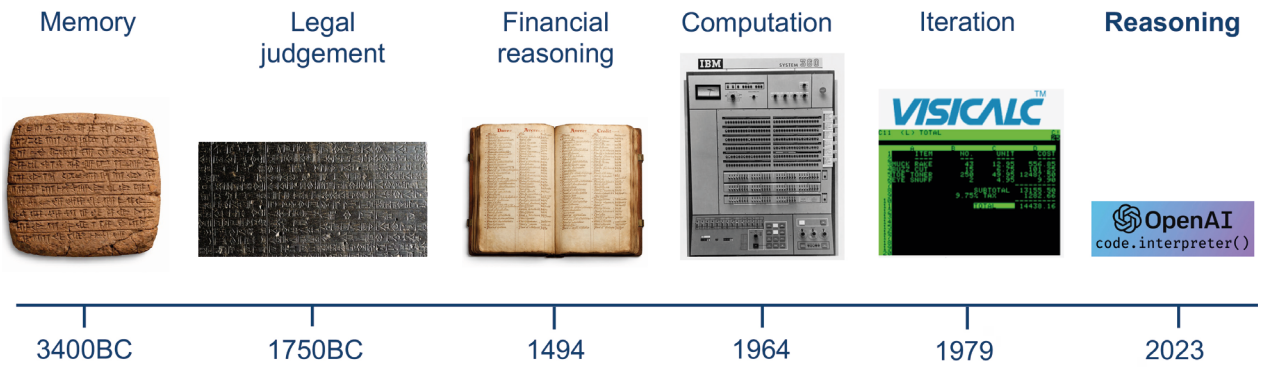

The richest problems – protein folding; materials discovery; climate modelling; drug design – sit at the intersection of domains no individual human can hold simultaneously in their mind. Specialisation was the workaround; the polymath was the rare exception. Abundant cognition collapses that constraint, allowing combinatorial search across domains that humans could only ever explore sequentially. The prize is problem-solving systems whose logic we can observe in operation without ever fully containing it in thought. Asked why LLMs work at all, Noam Shazeer, one of the architects of the transformer, replied: "My best guess is divine benevolence. Nobody really understands what's going on."

Oil transformed existing markets – ships and railways – but its real consequence lay elsewhere. Oil was a liquid; the change of state opened markets a solid could never reach: the automobile, and behind it the aeroplane, the highway, the petrochemical economy – none of which was visible from the boiler room. So it is likely to prove with cognition. The cheaper, faster analysis available today is remarkable but still only substitution. The transformation lies in the markets that could not previously exist at all: the space of things worth trying that we have never been able to afford to try.

This is why the 'low-hanging fruit has been picked' objection collapses. The sequence space for a modestly sized protein is 10^{260}; nature has produced perhaps 10^{13} across all species. The space of drug-like molecules is estimated at 10^{60}; chemistry has synthesised roughly 10^{8}. DeepMind's GNoME work expanded the catalogue of stable inorganic crystals from 20,000 to 421,000 in a single study – centuries of experimental progress in one pass and still a negligible fraction of the wider design space. The fruit was never too high. We have barely entered the orchard.

Ben Rogoff & Alastair Unwin
Polar Capital Technology Trust plc

9 July 2026