Artificial Intelligence, Semiconductors, and the New Politics of Technology: Governance, Industrial Policy and Global Markets in an Era of Strategic Competition

Micaela Lopez

Supervised by Dr William Vlcek
School of International Relations
University of St Andrews

1 September 2026

I. Introduction

Artificial intelligence (AI) and the semiconductor technologies that support it have undergone a rapid change in political and economic significance. Technologies recently discussed primarily in terms of commercial innovation, productivity and competition are increasingly treated as matters of national security, industrial resilience and geopolitical power. Frontier AI models—the most advanced general-purpose AI systems currently available—can be used in intelligence, cybersecurity and military applications, while the advanced semiconductors required to train and operate them depend on highly concentrated supply chains. Governments are therefore becoming more directly involved in determining who may develop, supply and access technologies largely shaped by private firms and global markets.

This shift is visible in the changing relationship between governments and technology companies. In the United States, disputes over the military use of private AI systems have become questions of procurement and supply-chain security, while frontier models have faced pre-release scrutiny and, in some cases, export restrictions. At the hardware level, export-licensing controls on advanced processors and chipmaking equipment, alongside efforts to relocate semiconductor production, demonstrate the strategic importance attached to computing capacity.

The consequences extend beyond American policy. The European Union has pursued greater technological sovereignty, seeking to reduce dependence on foreign technology and infrastructure. South Korea and Taiwan are managing the opportunities and vulnerabilities created by their strong positions in semiconductor manufacturing, while China is accelerating domestic substitution and applying AI more extensively within manufacturing. These responses differ according to countries' existing technological capabilities and positions within the global economy.

The transformation also remains commercial. Firms are experimenting with AI's effects on labour and productivity; frontier developers are attracting unprecedented private capital; and major technology companies are committing hundreds of billions of dollars to computing infrastructure. Existing legal institutions must meanwhile address questions involving copyright, patent inventorship and competition in markets already being reshaped by generative AI.

The importance of semiconductors to this analysis is therefore not merely technical. Advanced AI depends on physical infrastructure: processors, memory, manufacturing equipment, data centres and energy. Decisions about AI governance increasingly intersect with decisions about who controls that infrastructure, where it is produced and which users may access it. This connection between software capabilities and material supply chains provides the basis for examining AI as both an economic technology and an object of governance.

This paper traces these developments across governance, semiconductor and industrial policy, regional responses, markets and labour, and intellectual-property and antitrust law. It examines how the growing strategic importance of AI and its supporting infrastructure appears differently across institutional and economic settings, and what these developments reveal about the interaction between technology, the state, markets and law.

II. Governance: AI as a National-Security Object

The rapid development of frontier AI models has begun to change the terms on which governments interact with technology companies. AI models are no longer treated solely as commercial products whose risks can be addressed through conventional technology regulation. Their potential applications in intelligence, military operations and cybersecurity increasingly bring them within the institutional framework of national security.

The Anthropic-Pentagon Dispute

The dispute between Anthropic and the Pentagon provides a striking example. Anthropic had become the first frontier AI company to deploy models on US government classified networks. Its disagreement with the Pentagon concerned the limits of military use: Anthropic maintained that Claude could support essentially all lawful military applications while refusing mass domestic surveillance and fully autonomous weapons. It argued, in part, that current frontier models were insufficiently reliable for autonomous weapons systems.[1]

In March 2026, the Pentagon formally designated Anthropic a “supply chain risk”, a national-security designation allowing the government to restrict use of a supplier’s technology. Contractors were prevented from using Anthropic technology in work for the US military, although their non-Pentagon use of Claude remained unaffected.[2] Such powers had previously been associated mainly with foreign companies perceived as security threats, including Huawei; Anthropic was the first American company publicly subjected to the designation.[3]

Anthropic challenged the action in court, alleging violations of its First Amendment and due-process rights. A federal judge granted a preliminary injunction, finding that Anthropic was likely to succeed on its claim that the designation constituted retaliation for publicly challenging the government’s contracting position.[4] The dispute exposed a broader governance problem: governments increasingly depend on privately developed AI, yet firms may impose technical or ethical restrictions on systems embedded in state functions.

Frontier-Model Vetting

---

[1] Anthropic, "Statement on the Comments from Secretary of War Pete Hegseth," 27 February 2026.
[2] Reuters. "Pentagon Designates Anthropic a Supply Chain Risk." 5 March 2026.
[3] Ibid.
[4] Politico. "Judge Pauses Pentagon’s Punishment for Anthropic." 26 March 2026; Anthropic PBC v. U.S. Department of War. Complaint for Declaratory and Injunctive Relief. U.S. District Court for the Northern District of California, 9 March 2026.

This security framing also expanded beyond military procurement. Executive Order 14409, “Promoting Advanced Artificial Intelligence Innovation and Security,” issued in June 2026, created a framework through which developers could voluntarily provide the federal government access to advanced models for up to thirty days before public release. Agencies could then assess cybersecurity and other national-security risks.[5]

The order represented an adjustment to an administration that had previously favoured a comparatively hands-off approach to AI, but it stopped short of conventional licensing. Participation remained voluntary, while pre-release code audits, licensing and permitting requirements were excluded.[6] The policy therefore sought to balance industry competitiveness with government capacity to assess emerging risks.

The testing criteria remained largely classified, creating uncertainty about how models would be assessed and what would follow from identifying serious risks. The framework was also expected to focus on state-of-the-art proprietary systems while exempting open-source or open-weight models.[7] Oversight therefore increased through cooperation with a small number of frontier developers rather than comprehensive regulation of the wider AI market.

Fable, Mythos and Export Controls

The limits of cooperation became visible on 12 June 2026, when the Department of Commerce used national-security authorities to restrict foreign-national access to Anthropic’s Fable 5 and Mythos 5 models after concerns that Fable’s safeguards could be bypassed to expose cybersecurity capabilities. Because Anthropic could not reliably restrict access by nationality, it suspended both models for all users.[8] Anthropic disputed the government’s assessment.

The choice of regulatory instrument was significant. Export controls traditionally regulate the international movement of strategically sensitive goods and technologies; applying them to access to an AI model treated model capabilities themselves as a potentially controlled strategic resource. The restrictions were lifted after Anthropic implemented an additional safeguard approved by the government’s Center for AI Standards and Innovation.[9]

The Anthropic dispute is particularly useful because it shows that the relationship between the state and frontier developers cannot be described simply as regulation of private industry. The government is also a customer of AI systems, while developers retain practical control over how their models are designed and deplyoyed. Procurement decisions can therefore

[5] Executive Order 14409. "Promoting Advanced Artificial Intelligence Innovation and Security." 2 June 2026; New York Times. "Trump Signs Executive Order Seeking Oversight of AI Models." 2 June 2026.
[6] Atlantic Council. "Reading Between the Lines of Trump’s New Executive Order on AI." 3 June 2026.
[7] Politico. "White House AI Vetting Plan to Exempt Lower-Cost ‘Open’ Models." 4 August 2026; The Guardian. "The White House’s Plan to Vet Potentially Dangerous AI Is Cloaked in Secrecy." 7 August 2026.
[8] Anthropic. "Statement on the US Government Directive to Suspend Access to Fable 5 and Mythos 5." 12 June 2026; Financial Times. "Anthropic Scrambles after Trump Administration Freezes Its Top AI Models." 14 June 2026.
[9] Financial Times. "White House Lifts Ban on Anthropic Models." 30 June 2026.

become instruments of technology policy, and disagreements over model safeguards can acquire consequences extending beyond an ordinary commercial contract.

The framework also illustrates a broader institutional difficulty. Pre-release review may give government specialists earlier access to potentially important capabilities, but voluntary participation leaves substantial responsibility with developers themselves. Classified assessment criteria may be justified by security concerns, yet they also make it harder for firms and outside researchers to evaluate how consistently risks are being judged. The result is a form of governance that is more interventionist than ordinary market oversight but less comprehensive than a licensing system.

Taken together, these developments do not amount to a settled American model of AI governance. Their mixture of procurement powers, voluntary review, classified benchmarks and export controls instead shows a change in the object being governed. As frontier models become useful for intelligence, cyber operations and military activity, their development and distribution increasingly resemble questions associated with strategic infrastructure. This security logic is still more visible in the semiconductor hardware on which advanced AI depends.

III. Semiconductors and the Export-Control Regime

If frontier AI models are increasingly treated as strategic capabilities, the semiconductor supply chain provides the physical counterpart. Training and operating advanced AI depends on specialised chips, manufacturing equipment and production capacity concentrated among relatively few firms and countries. US policy has consequently combined export licensing, restrictions on processors and manufacturing equipment, and efforts to diversify production geographically.

BIS Licensing and Advanced AI Chips

At the centre of the US export-control system is the Department of Commerce's Bureau of Industry and Security (BIS), which restricts transfers of sensitive dual-use technologies—technologies with both civilian and potential military or security applications. Export licensing allows the government to distinguish between permitted and prohibited transactions rather than imposing a complete embargo.

By 2026, the licensing process itself had become economically significant. A CSIS survey of major US technology exporters found that 56 per cent of respondents reported average BIS review periods exceeding 180 days, despite a 90-day statutory benchmark, while more than half reported losing business or contracts because of delays.[10] The small, semiconductor-heavy sample limits how widely the findings can be generalised, but it illustrates how security screening can impose commercial costs even where exports are ultimately permitted.

[10] Center for Strategic and International Studies. "Delays and Uncertainty in the Export Licensing Process: Findings from a CSIS Survey of U.S. Technology Exporters." 21 April 2026.

Nvidia's H200, an advanced processor used for computationally intensive AI workloads, shows how licensing can calibrate access rather than simply prohibit it. In January 2026, BIS allowed applications for exports of the H200, AMD's MI325X and comparable processors to China to be assessed case by case, subject to conditions including customer screening and independent testing in the United States.[11] More advanced generations remained restricted. The policy therefore involves a continuing judgement over how much technological access can be permitted without undermining US strategic advantages or accelerating demand for Chinese substitutes.

Reshaping Production: TSMC in Arizona

Export controls address who can acquire advanced technology, but US policy also addresses where it is produced. Advanced semiconductor manufacturing remains concentrated in East Asia, particularly Taiwan, making disruption in the Taiwan Strait a potential economic and security risk.

TSMC's Arizona expansion represents an attempt to reduce this exposure. In 2026, the company announced another $100 billion in US investment, bringing its total planned commitment to American manufacturing to $265 billion. The expansion includes semiconductor fabrication facilities, or "fabs", where chips at two nanometres and below will be manufactured, alongside advanced packaging—the processes used to combine and connect sophisticated semiconductor components.[12]

Geographic diversification, however, does not immediately reproduce the semiconductor ecosystem developed in Taiwan over decades. Production in Arizona has faced higher costs, workforce requirements and an incomplete supplier network. Advanced manufacturing also depends on manufacturing yield—the proportion of chips produced that function correctly—as well as specialist engineers, materials and equipment.[13] The project is therefore better understood as gradual diversification than straightforward self-sufficiency.

ASML and Allied Controls

The international structure of semiconductor production limits unilateral controls. ASML, a Dutch company, dominates advanced lithography equipment: machinery used to project extremely small circuit patterns onto semiconductor wafers. Its most advanced extreme-ultraviolet (EUV) systems have been prohibited from export to China since 2019, while

[11] U.S. Department of Commerce, Bureau of Industry and Security. "Department of Commerce Revises License Review Policy for Semiconductors Exported to China." 13 January 2026; U.S. Department of Commerce, Bureau of Industry and Security. "Revision to License Review Policy for Advanced Computing Commodities." Federal Register 91, no. 10 (15 January 2026): 1684.
[12] Financial Times. "Chipmaker TSMC to Invest Another $100bn in US Production." 16 July 2026; National Institute of Standards and Technology. "TSMC Arizona." Accessed 17 June 2026.
[13] Nikkei Asia. "TSMC's $165bn US Expansion Push: 'Inevitable' but Risky." 26 March 2025; Nikkei Asia. "How Arizona Is Cultivating Talent for a US Chip Revival." 31 December 2025.

later Dutch restrictions extended to certain deep-ultraviolet (DUV) systems.[14] US strategy therefore depends partly on coordination with the Netherlands, Japan and other semiconductor-producing economies.

Controls on existing capacity are harder to enforce than restrictions on new sales. Chinese manufacturers have reportedly upgraded older ASML DUV machines using secondary-market components and third-party engineering support. Techniques such as multiple patterning, in which lithography steps are repeated to create smaller and more complex chip features, can extend older equipment's capabilities, although at higher cost and lower efficiency.[15]

Enforcement and Circumvention

Restricted Nvidia chips and complete AI servers have also continued to reach Chinese buyers through intermediary jurisdictions and informal supply chains. By June 2026, tighter enforcement had raised black-market prices, while investigations in Taiwan and Singapore examined alleged diversions through false destinations or end-user declarations.[16]

The H200 decision also demonstrates why export controls are economically difficult to calibrate. A restriction may preserve a technological advantage by limiting access to advanced computing power, but it can also reduce sales by American firms and strengthen incentives for customers to seek alternatives. The policy problem is therefore dynamic: the effectiveness of a control depends not only on the technology restricted at the time a rule is written, but also on how suppliers, purchasers and rival producers respond afterwards.

TSMC's investment complements this strategy by addressing supply resilience rather than access alone. Export controls attempt to limit the technologies available to strategic competitors, whereas domestic industrial investment seeks to reduce American dependence on geographically concentrated production. These policies share a security objective but operate through different economic mechanisms: one constrains transactions and the other attempts to redirect capital and productive capacity.

Allied coordination creates a further trade-off. The semiconductor supply chain is international precisely because different firms specialise in different stages of production. This specialisation creates technological chokepoints that can make controls effective, but it also means that governments cannot easily act independently. Restrictions imposed by one country may be weakened if equivalent equipment or services remain available elsewhere, while broader coordination can impose costs on firms in allied economies as well as on the intended target.

Circumvention does not necessarily mean that controls are ineffective: higher prices, shortages and more complex diversion routes may indicate that access has become harder. It does, however, reveal the difficulty of enforcing technological restrictions across global supply chains. Semiconductor controls therefore operate through repeated rulemaking, adaptation and

---

[14] Financial Times. "Netherlands to Restrict Chip Exports after US Pressure over China Threat." 8 March 2023.
[15] Financial Times. "China Boosts AI Chip Output by Upgrading Older ASML Machines." 18 December 2025.
[16] Financial Times. "Nvidia's Banned AI Chips Double in Price on China's Black Market." 23 June 2026; Nikkei Asia. "Singapore Seizes $42m Home in Nvidia Chip Smuggling Case." 1 July 2026.

enforcement. Chips, fabrication capacity and manufacturing equipment remain commercial assets, but they are increasingly governed as strategic resources.

IV. Regional Responses

The strategic importance of AI and semiconductors has produced different responses outside the United States because economies occupy different positions within the technology supply chain. Europe is seeking to reduce dependence on foreign digital infrastructure; South Korea and Taiwan are managing the opportunities and vulnerabilities created by semiconductor strength; and China is accelerating substitution and automation in response to external restrictions and domestic pressures.

European Digital Sovereignty

The European Union's response is increasingly framed through technological sovereignty: reducing dependence on foreign providers for strategically important technologies and digital infrastructure. In June 2026, the European Commission proposed a European Technological Sovereignty Package covering semiconductors, cloud computing, AI and open-source technology. Chips Act 2.0 seeks to expand European semiconductor capabilities, while the proposed Cloud and AI Development Act aims to triple European data-centre capacity over five to seven years.[17] The objective extends beyond chip production to the infrastructure through which AI is developed and deployed.

This concern predates the recent AI boom. Gaia-X, launched in 2020, sought an interoperable European data infrastructure based on common standards rather than a single centralised cloud provider. By 2025, more than 180 sector-specific data spaces were moving towards implementation.[18] The project nevertheless illustrates the difficulty of coordinating firms and governments while competing with established US providers.

Although Europe hosts many data centres, it accounts for roughly 18 per cent of global IT capacity compared with approximately half in the United States. A small group of US hyperscalers—very large cloud-computing providers operating data centres at global scale—controls a particularly large share of infrastructure used for computationally intensive workloads.[19] Greater European autonomy therefore requires physical investment as well as regulation.

Draft revisions to the Chips Act would also give the Commission stronger crisis powers during shortages. The proposals emerged partly from the Nexperia disruption, which showed how dependence on geographically fragmented semiconductor supply could affect European

---

[17] European Commission. "Commission Proposes Tech Sovereignty Package to Strengthen Europe's Digital Autonomy and Resilience." 3 June 2026.
[18] Institut Polytechnique de Paris. "Gaia-X: The Bid for a Sovereign European Cloud." 18 June 2025.
[19] Center for Economic Policy Research. "Powering the Digital Economy: The Global Expansion of Data Centers and Its Energy Implications." 28 March 2026.

automotive production.[20] Europe's approach is therefore not technological isolation, but an attempt to combine access to trusted foreign providers with greater domestic resilience.

East Asia: Production, Concentration and Substitution

South Korea's response builds on an existing industrial advantage. Samsung Electronics and SK Hynix dominate high-bandwidth memory (HBM), a specialised form of high-speed memory used alongside advanced AI processors, and together account for nearly 80 per cent of the global HBM market. In June 2026, South Korea announced an approximately Won800 trillion public–private expansion involving four new fabrication plants and investment in advanced packaging and next-generation chips.[21]

Strong semiconductor exports contributed materially to South Korean growth in 2026, but large investments also create exposure to the semiconductor cycle. New fabrication capacity takes years to become operational, so present investment partly assumes continued hyperscaler demand into the next decade.[22] South Korea is therefore expanding an established advantage while accepting the risk of future overcapacity.

Taiwan faces a different form of concentration. TSMC's centrality to advanced manufacturing has supported strong exports and investment, but by May 2026 the company represented approximately 43.8 per cent of the market capitalisation of Taiwan's listed companies.[23] The government has sought to spread AI-related growth into robotics, healthcare, energy and traditional manufacturing.[24] TSMC's overseas expansion also creates a tension: diversifying global production may reduce geopolitical supply risk while gradually dispersing some of the capacity that gives Taiwan strategic importance.

China's response is shaped more directly by restrictions on foreign technology. Chinese authorities have added processors from companies including Huawei and Cambricon to government procurement programmes and encouraged state-owned institutions to replace foreign hardware.[25] Huawei's Ascend processors have become a significant alternative to Nvidia, although software adaptation, performance and production constraints remain. Domestic AI chips have gained market share, but estimates still place China's overall AI computing capacity substantially below that of the United States.[26] Export restrictions have therefore encouraged substitution without eliminating the technological gap.

China is also promoting industrial and humanoid robotics as a response to demographic decline. Its working-age population is projected to contract sharply, while the country already

[20] Financial Times. "EU Wants Crisis Powers to Seize Control of Chip Supplies." 28 May 2026; Euronews. "European Carmakers Face EU Pressure to Diversify Chip Suppliers." 21 May 2026.
[21] Financial Times. "Samsung and SK Hynix Plan $600bn Chipmaking Expansion." 29 June 2026.
[22] Ibid; Financial Times. "Chips, Ships and Guns: South Korea Booms on AI Race and Global Conflict." 7 June 2026.
[23] Taipei Times. "Analysis: How AI Has Rewired Taiwan's Stock Market." 31 July 2026.
[24] Nikkei Asia. "Taiwan GDP Grows Nearly 13% in Q2, Supercharged by AI and US Ties." 31 July 2026.
[25] Financial Times. "China Adds Domestic AI Chips to Official Procurement List for First Time." 10 December 2025.
[26] The Wall Street Journal. "Inside China's All-Out Push to Catch Up with American AI Chips." 23 July 2026.

installs more industrial robots annually than the rest of the world combined.[27] Automation may alleviate future labour shortages while displacing workers in the nearer term, and policymakers have acknowledged employment risks alongside proposals for retraining and social-security adjustment.[28]

The European strategy consequently combines industrial policy with infrastructure policy. Increasing chip capacity without expanding cloud and data-centre resources would address only one part of the AI supply chain. Conversely, expanding computing capacity while remaining dependent on foreign semiconductor production would leave other strategic dependencies intact. The emphasis on sovereignty reflects this interconnectedness, although reducing dependence can conflict with the efficiencies created by an internationally specialised technology market.

South Korea and Taiwan show that technological concentration can be economically advantageous and strategically uncomfortable at the same time. Strong demand for AI hardware raises exports, investment and corporate valuations, but it also increases exposure to changes in global AI spending. In Taiwan's case, concentration has an additional geopolitical dimension because the global importance of domestic semiconductor production is itself part of the island's strategic significance.

China provides the clearest example of policy adaptation to external controls. Restrictions have raised the cost of accessing leading foreign technologies while increasing incentives to invest in domestic processors, software ecosystems and manufacturing techniques. This does not mean that substitution is immediate or complete. Performance gaps, manufacturing constraints and the cost of switching away from established technologies can persist even as domestic alternatives improve. The response therefore demonstrates how controls can reshape the direction of investment without fully determining its eventual outcome.

These regional responses do not form a single policy trajectory. Europe is pursuing greater autonomy, South Korea is expanding industrial capacity, Taiwan is managing concentration, and China is combining technological substitution with automation. In each case, however, the strategic turn in technology policy is producing consequences for investment, firms and labour.

V. Markets, Firms and Labour

The economic effects of AI are developing on several levels simultaneously. Firms are experimenting with AI as a substitute for and complement to labour; frontier developers are attracting extraordinary valuations; and technology companies are committing unprecedented capital to AI infrastructure. These developments do not yet point towards a single economic conclusion.

---

[27] Financial Times. "Robot Nation: China's Bid to Beat Its Demographic Decline." 24 June 2026.
[28] Ibid.

Productivity, Employment and Adjustment

The most immediate question is whether generative AI complements workers or replaces them. Early evidence is mixed. Some studies identify vulnerable groups: employment among workers aged 22 to 25 reportedly fell relative to older workers in occupations highly exposed to AI—that is, jobs containing a relatively high proportion of tasks AI could perform or assist with—while firms recruiting for generative-AI positions reduced junior hiring more sharply than comparable firms. Broader studies, however, found much weaker effects, including little evidence of major changes in the overall occupational composition of the US labour market.[29]

Part of this disagreement is methodological. Researchers use different measures of AI exposure, company adoption and occupational change, while government statistics are poorly suited to tracking a technology spreading unevenly across industries.[30] Private datasets may provide greater detail but are less representative of the economy as a whole.

Productivity gains and labour displacement can also occur simultaneously. AI may allow an employee to complete more work while reducing the number of additional workers a firm needs to hire. Conversely, firms adopting AI may expand if the technology lowers costs or creates new products. Individual layoffs therefore provide limited evidence of economy-wide effects, particularly because recent technology-sector reductions also reflect pandemic-era over-hiring and broader cyclical conditions.[31]

Policy responses are developing before the eventual labour-market effect is known. Raise US, a $500 million initiative involving technology companies, state governments and foundations, proposes community-college partnerships, retraining, wage insurance and support for moving employees into new roles.[32] Nearly 200 economists and technology leaders also warned in 2026 that AI could produce both large gains in living standards and significant displacement.[33] The central uncertainty is therefore not simply whether AI creates or destroys jobs, but how quickly adjustment occurs and how its costs and benefits are distributed.

Corporate Strategy, Pricing and Valuations

Capital markets are assigning extraordinary value to frontier developers. In March 2026, OpenAI completed a $122 billion financing round valuing the company at approximately $852

[29] Brynjolfsson, E., B. Chandar, and R. Chen. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence." Working paper, Stanford Digital Economy Lab. Revised 12 August 2026; Lichtinger, G., and S.H. Maasoum. "Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data." 2025; The Budget Lab at Yale. "Tracking the Impact of AI on the Labor Market." 15 June 2026.

[30] New York Times. "AI Is Reshaping the Economy. Good Luck Measuring How." 2 July 2026.

[31] The Economist. "Don't Blame AI for Your Job Woes." 6 November 2025; New York Times. "Is AI Replacing Tech Workers or Providing an Excuse for Job Cuts?" 1 June 2026.

[32] New York Times. "Big Companies Aim to Ease AI Transition for American Workers." 25 June 2026.

[33] "We Must Act Now: A Statement on AI's Transformation of the Economy." 2026; New York Times. "Nearly 200 Economists and Tech Leaders Warn of AI Threats." 13 July 2026.

billion. Two months later, Anthropic raised $65 billion at a $965 billion post-money valuation.[34] These financings provide capital for chips, data centres and model development while illustrating investor expectations about future AI revenues.

Both firms are increasingly focused on enterprise customers and have considered public listings, although OpenAI later considered delaying its offering until 2027.[35] Their pricing strategies also show how the business model is evolving. AI agents can consume substantially more computing resources than ordinary chatbot queries, encouraging metered or usage-based pricing. Customers may be charged according to tokens—small units of text processed or generated by a model—or computing resources consumed. Anthropic and OpenAI already charge enterprise customers partly on this basis.[36]

This makes the cost of computation more visible to customers. Businesses must judge whether additional AI use produces productivity gains sufficient to justify its cost, while providers must recover large infrastructure expenses without pushing customers towards cheaper models. Some firms have therefore limited employee AI spending or routed routine tasks towards less expensive models.[37]

Investment and the Bubble Question

By April 2026, Amazon, Alphabet, Microsoft and Meta were planning as much as $725 billion in combined annual capital expenditure—long-term investment in assets such as chips, data centres and computing equipment—primarily associated with AI infrastructure.[38] Large technology companies are consequently becoming more capital intensive.

The scale of investment does not by itself establish an AI bubble. Investors have become less tolerant of rapidly rising expenditure: larger-than-expected capital budgets contributed to sharp falls in major technology stocks, with approximately $950 billion in market value erased following one series of spending forecasts.[39] Yet there is also evidence of substantial commercial demand. One 2026 estimate found that quarterly AI revenues for hyperscalers and specialised cloud providers had begun exceeding depreciation—the accounting cost reflecting the declining value of infrastructure as equipment ages or becomes obsolete—although margins remained narrow after other costs.[40]

---

[34] OpenAI. "Accelerating the Next Phase of AI." 2026; Anthropic. "Anthropic Raises $65B in Series H Funding at $965B Post-Money Valuation." 28 May 2026; Bloomberg. "Anthropic Eclipses OpenAI with Valuation of $965 Billion." 28 May 2026.

[35] OpenAI. "Accelerating the Next Phase of AI." 2026; Anthropic. "Anthropic Raises $30 Billion in Series G Funding at $380 Billion Post-Money Valuation." 12 February 2026; Bloomberg. "Anthropic Files Confidentially for IPO in Race with OpenAI." 1 June 2026; Bloomberg. "OpenAI Weighs IPO in 2027 after Expected Anthropic Public Debut." 26 June 2026.

[36] Bloomberg. "The End of Tokenmaxxing? How AI Is Moving to Pay-as-You-Go." 18 June 2026; KPMG. "Global AI Pulse Q2 2026." 2026.

[37] Ibid.

[38] Bloomberg. "US Big Tech Ratchets Up AI Spending Past $700 Billion This Year." 30 April 2026.

[39] Bloomberg. "How Much Is Big Tech Spending on AI Computing? A Staggering $650 Billion in 2026." 6 February 2026.

[40] Azhar, A., et al. "The State of the AI Economy." Exponential View, 25 June 2026.

The distinction between tasks and jobs is important for interpreting this evidence. Generative AI may automate particular activities within an occupation without eliminating the occupation itself, while the remaining tasks may become more productive or change in character. Entry-level roles may be especially sensitive if they contain routine research, drafting or analytical tasks that can be partly automated. At the same time, new demand can emerge for workers who supervise, integrate or evaluate AI systems. This makes short-term employment effects difficult to translate into a definitive prediction about the long-run labour market.

The scale of private financing also connects corporate strategy to the semiconductor and infrastructure issues discussed earlier. Frontier developers require access to advanced processors, data-centre capacity and energy, meaning that model development increasingly involves large fixed costs. High valuations therefore reflect expectations not only about software demand but also about whether firms can convert expensive computing infrastructure into sustainable revenue. Enterprise adoption is especially important because business customers may provide larger and more predictable revenue streams than consumer subscriptions.

Usage-based pricing makes this relationship more explicit. As models perform longer or more complex tasks, the cost of serving a customer can rise with the amount of computation used. Metering therefore allows providers to pass some of that cost directly to customers, but it also gives businesses a clearer reason to compare AI expenditure with measurable productivity gains. The commercial success of increasingly capable models may consequently depend as much on cost efficiency as on technical performance.

The bubble question should therefore be separated into several issues. High valuations can prove excessive even if AI becomes economically important, and some infrastructure projects can generate poor returns even if aggregate demand continues to grow. Conversely, market corrections do not establish that the underlying technology lacks value. The relevant economic question is whether the revenues and productivity gains generated over time are sufficient to support the capital committed to developing and operating AI systems.

The economic picture is therefore unsettled rather than contradictory. AI may raise productivity while disrupting particular workers; frontier firms may generate growing revenues while depending on exceptional capital requirements; and infrastructure spending may prove necessary even if some present valuations do not. Markets are pricing competing expectations about how quickly technological capability will translate into durable economic value.

VI. Law Catching Up: Intellectual Property and Antitrust

The economic expansion of AI has required courts and regulators to apply legal doctrines developed before generative AI to new forms of production and competition. Rather than creating entirely new regimes, courts and agencies have generally attempted to fit AI training, AI-assisted invention and changing digital markets within existing copyright, patent and antitrust law.

Copyright and Fair Use

Copyright litigation has centred on the large quantities of copyrighted material used to train generative AI. Rights holders argue that developers copied protected works without permission or compensation, while AI companies argue that their reliance on fiar use—which permits certain uses of copyrighted materials without the owner's permission—is justified, since training is transformative: models learn from source materials to generate different outputs.[41]

In *Bartz v. Anthropic*, US District Judge William Alsup held in June 2025 that Anthropic's use of books to train Claude was sufficiently transformative to constitute fair use. The court nevertheless distinguished training from acquisition: Anthropic's creation and retention of a library containing millions of pirated books was not protected merely because some copies might later serve a transformative purpose.[42] Anthropic subsequently agreed to a $1.5 billion settlement concerning the pirated works.[43]

A separate case against Meta demonstrated that the boundaries remain unsettled. Judge Vince Chhabria ruled for Meta because the plaintiffs had not sufficiently established market harm, but warned that AI training could fail the fair-use test in other circumstances, particularly where generated material competes with human creators.[44] Existing fair-use analysis can therefore accommodate some AI training without automatically legalising the acquisition of training data or resolving whether particular outputs infringe copyright. Licensing agreements are meanwhile developing alongside litigation.[45]

Patent Inventorship

Patent law raises the related question of who receives legal recognition when AI contributes to invention. In *Thaler v. Vidal*, patent applications named the AI system DABUS as inventor. The USPTO rejected them, and the Federal Circuit held that the Patent Act's reference to an "individual" inventor requires a natural person; the Supreme Court declined to hear the appeal.[46]

In November 2025, the USPTO revised its guidance for AI-assisted inventions, reaffirming that no separate inventorship standard applies simply because AI was involved. AI systems are treated as tools, while inventorship remains limited to humans.[47] Patent law has

---

[41] Reuters. "AI Copyright Battles Enter Pivotal Year as US Courts Weigh Fair Use." 5 January 2026.
[42] Bartz v. Anthropic PBC. Order on Fair Use. U.S. District Court for the Northern District of California, 23 June 2025; Reuters. "Anthropic Wins Key US Ruling on AI Training in Authors' Copyright Lawsuit." 24 June 2025.
[43] Reuters. "Anthropic Agrees to Pay $1.5 Billion to Settle Author Class Action." 5 September 2025.
[44] Reuters. "Meta Fends Off Authors' US Copyright Lawsuit over AI." 25 June 2025.
[45] Reuters. "AI Copyright Battles Enter Pivotal Year as US Courts Weigh Fair Use." 5 January 2026.
[46] American Bar Association. "Can an AI Be an Inventor? In the US, the Answer Remains No." 25 May 2023.
[47] United States Patent and Trademark Office. "Revised Inventorship Guidance for AI-Assisted Inventions." 26 November 2025; Aboy, M., and K. Liddell. "Revised USPTO Guidance on Inventorship for AI-Assisted Inventions: A Pro-Innovation Pivot Away from Pannu Factors." Journal of Intellectual Property Law & Practice 21, no. 5 (2026): 272–275.

therefore accommodated increased AI involvement principally by maintaining its existing human conception of inventorship.

Antitrust and the Google Search Case

Antitrust law, which seeks to protect competition by restricting monopolistic and other anticompetitive conduct, reveals a different problem: the relevant market can change while litigation is underway. The Department of Justice's case against Google began in 2020 and focused on agreements making Google the default search engine on browsers and mobile devices. In August 2024, Judge Amit Mehta concluded that Google unlawfully maintained monopolies in general search and search text advertising through exclusionary distribution agreements.[48]

By the remedies stage, generative AI had begun changing the competitive environment. The government sought remedies including the sale of Google's Chrome browser. In September 2025, the court instead required Google to share certain search data with qualified competitors and restricted aspects of default-placement arrangements without ordering a Chrome divestiture. Judge Mehta observed that generative AI had "changed the course of this case".[49] Companies including OpenAI, Anthropic and Perplexity were providing alternative ways to retrieve information, while Google itself was incorporating generative responses into search.

The distinction drawn in the copyright cases is important because "AI training" is not treated as a single legal act. A court may consider the use of a work for training transformative while reaching a different conclusion about how the training copy was obtained or whether a generated output infringes copyright. For a non-specialist reader, this helps explain why apparently conflicting decisions can coexist: fair use is a context-dependent legal doctrine rather than a general exemption for AI developers.

Patent law currently presents a clearer boundary because the statutory question of inventorship remains tied to a human individual. The unresolved issue is therefore less whether AI may assist inventive activity than how much human involvement is required for a person to qualify as the inventor. As AI becomes more capable within research and design, that distinction may become increasingly important even if the underlying statutory rule remains unchanged.

The Google litigation shows a different timing problem. Copyright and patent cases ask how established rights apply to new methods of creation, whereas antitrust remedies attempt to shape competition prospectively after unlawful conduct has been identified. If the market changes during years of litigation, a remedy designed for the market that existed when the case began may no longer address the most important competitive conditions when it is imposed. Generative AI therefore affects not only the subject matter of technology law but also the environment in which existing legal remedies operate.

---

[48] United States et al. v. Google LLC. Memorandum Opinion. U.S. District Court for the District of Columbia, 5 August 2024.
[49] New York Times. "Google Avoids Harshest Penalties in Landmark Search Monopoly Ruling." 2 September 2025.

The case illustrates the retrospective nature of antitrust enforcement: courts assess past conduct and then design remedies for future competition. In rapidly changing technology markets, that environment may evolve significantly during the process. Across copyright, patent and antitrust law, existing doctrine remains relevant, but technological capabilities and markets can develop more quickly than the rules applied to them.

VII. Conclusion

The developments examined in this paper do not produce a single model of how governments and markets are responding to AI. Frontier models have become objects of procurement control, government vetting and export restriction, while semiconductor policy extends this security logic into the physical infrastructure that AI depends on. Regional responses vary according to existing economic positions, and market exposure is at least as concentrated internationally as it is domestically: Samsung and SK Hynix together hold roughly 80 per cent of the global high-bandwidth memory market, and TSMC alone accounts for some 44 per cent of Taiwan's stock market capitalisation.[50] The market consequences of this concentration, like those of the US-centered investment boom, remain unsettled.

Two patterns recur across this research. The first is a mismatch of speed: AI capabilities, investment and commercial practice change faster than the administrative procedures, statistics, litigation and legislation meant to keep pace with them—visible in uncertain labour-market measurement, successive revisions to export controls, and the Google case, where the competitive environment itself shifted while litigation was still under way.[51] The clearest economic instance of this mismatch is the labour-market signal: a widening employment gap for young workers in AI-exposed occupations, which its own authors present as an early, descriptive indicator rather than a causal estimate, precisely because the underlying data cannot yet separate task-level automation from broader displacement.[52] Much the same caution applies to export-control effectiveness, infrastructure investment and legal doctrine—each still resting on early signals rather than settled proof.

The second pattern is, more directly, a political-economy tension: governments seek to restrict access, diversify supply and reduce strategic dependencies, while firms operate through global supply chains, international customers and commercial incentives that reward scale and efficiency. Together, these patterns suggest that the present period is defined less by any settled equilibrium than by the ongoing friction between institutions built for a slower pace of change and a technology that does not wait for them.

[50] Financial Times. "Samsung and SK Hynix Plan $600bn Chipmaking Expansion." 29 June 2026.
[51] United States et al. v. Google LLC. Memorandum Opinion. U.S. District Court for the District of Columbia, 5 August 2024.
[52] Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence." Working paper, Stanford Digital Economy Lab. Revised 12 August 2026.

This also means the evidence base itself is a moving target: labour-market studies rely on different and evolving measures of AI exposure, corporate investment plans are revised within months, and several of the legal and regulatory arrangements discussed above remain unsettled. Even the basic economic question—whether AI investment is justified by the revenue it generates—cannot yet be answered with confidence: hyperscaler AI revenues have only recently begun to exceed depreciation costs, and margins remain narrow.[53] The findings presented here should therefore be read as a picture of an ongoing technological and institutional transition, not a final assessment of its effects.

Future research should examine whether export controls produce lasting technological advantages or instead accelerate substitution by rivals; whether continued geographic concentration in semiconductor manufacturing represents a durable strategic advantage for South Korea and Taiwan or an increasingly acute vulnerability as production diversifies and global demand cycles turn; how AI's labour and productivity effects develop as more reliable data become available; whether the current scale of AI capital investment is validated by durable revenue growth or instead reflects a period of overbuilding disconnected from realised returns; and whether existing intellectual-property and competition doctrines remain durable as AI capabilities continue to change. The significance of the present period may ultimately lie less in any single policy or market outcome than in the process through which institutions are redefining their relationship with technologies that have become, simultaneously, commercial assets, economic infrastructure and strategic resources.

Acknowledgements

I would like to thank my supervisor, Dr William Vlcek, for his guidance and support throughout the development of this research project. I am grateful for his time in helping me brainstorm countless ideas, as well as his consistent guidance throughout the research and writing period.

I am also grateful to Lord Laidlaw and the Laidlaw Foundation for their support of my research and upcoming leadership development, which has made this paper and my other outputs, past and future, possible. A special thank you to the Laidlaw team at the University of St Andrews for this opportunity.

[53] Azhar, Azeem, et al. "The State of the AI Economy." Exponential View, 25 June 2026.

Bibliography

Aboy, Mateo, and Kathleen Liddell. "Revised USPTO Guidance on Inventorship for AI-Assisted Inventions: A Pro-Innovation Pivot Away from Pannu Factors." Journal of Intellectual Property Law & Practice 21, no. 5 (2026): 272–275. https://academic.oup.com/jiplp/article/21/5/272/8528856.

American Bar Association. "Can an AI Be an Inventor? In the US, the Answer Remains No." 25 May 2023. https://www.americanbar.org/groups/business_law/resources/business-law-today/2023-june/can-ai-be-an-inventor/.

Anthropic. "Anthropic Raises $30 Billion in Series G Funding at $380 Billion Post-Money Valuation." 12 February 2026. https://www.anthropic.com/news/anthropic-raises-30-billion-series-g-funding-380-billion-post-money-valuation.

Anthropic. "Anthropic Raises $65B in Series H Funding at $965B Post-Money Valuation." 28 May 2026. https://www.anthropic.com/news/series-h.

Anthropic. "Statement on the Comments from Secretary of War Pete Hegseth." 27 February 2026. https://www.anthropic.com/news/statement-comments-secretary-war.

Anthropic. "Statement on the US Government Directive to Suspend Access to Fable 5 and Mythos 5." 12 June 2026. https://www.anthropic.com/news/fable-mythos-access.

Anthropic PBC v. U.S. Department of War. Complaint for Declaratory and Injunctive Relief. U.S. District Court for the Northern District of California, 9 March 2026. https://s3.documentcloud.org/documents/27781298/anthropic-v-dow.pdf.

Atlantic Council. "Reading Between the Lines of Trump's New Executive Order on AI." 3 June 2026. https://www.atlanticcouncil.org/dispatches/reading-between-the-lines-of-trumps-new-executive-order-on-ai/.

Azhar, Azeem, et al. "The State of the AI Economy." Exponential View, 25 June 2026. https://www.exponentialview.co/p/the-state-of-the-ai-economy.

Bartz v. Anthropic PBC. Order on Fair Use. U.S. District Court for the Northern District of California, 23 June 2025. https://docs.justia.com/cases/federal/district-courts/california/candce/3:2024cv05417/434709/231.

Bloomberg. "Anthropic Eclipses OpenAI with Valuation of $965 Billion." 28 May 2026. https://www.bloomberg.com/news/articles/2026-05-28/anthropic-raises-at-965-billion-valuation-eclipsing-openai.

Bloomberg. "Anthropic Files Confidentially for IPO in Race with OpenAI." 1 June 2026. https://www.bloomberg.com/news/articles/2026-06-01/anthropic-files-confidentially-for-ipo-as-claude-demand-surges.

Bloomberg. "The End of Tokenmaxxing? How AI Is Moving to Pay-as-You-Go." 18 June 2026. https://www.bloomberg.com/news/articles/2026-06-18/ai-costs-what-shift-from-flat-rate-to-token-pricing-model-means-for-industry.

Bloomberg. "How Much Is Big Tech Spending on AI Computing? A Staggering $650 Billion in 2026." 6 February 2026. https://www.bloomberg.com/news/articles/2026-02-06/how-much-is-big-tech-spending-on-ai-computing-a-staggering-650-billion-in-2026.

Bloomberg. "OpenAI Weighs IPO in 2027 after Expected Anthropic Public Debut." 26 June 2026. https://www.bloomberg.com/news/articles/2026-06-26/openai-weighs-ipo-in-2027-after-expected-anthropic-public-debut.

Bloomberg. "US Big Tech Ratchets Up AI Spending Past $700 Billion This Year." 30 April 2026. https://www.bloomberg.com/news/articles/2026-04-30/us-big-tech-ratchets-up-ai-spending-past-700-billion-this-year.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence." Working paper, Stanford Digital Economy Lab. Revised 12 August 2026. https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/.

The Budget Lab at Yale. "Tracking the Impact of AI on the Labor Market." 15 June 2026. https://budgetlab.yale.edu/research/tracking-impact-ai-labor-market.

Center for Economic Policy Research. "Powering the Digital Economy: The Global Expansion of Data Centers and Its Energy Implications." 28 March 2026. https://cepr.org/voxeu/columns/powering-digital-economy-global-expansion-data-centres-and-its-energy-implications.

Center for Strategic and International Studies. "Delays and Uncertainty in the Export Licensing Process: Findings from a CSIS Survey of U.S. Technology Exporters." 21 April 2026.

https://www.csis.org/analysis/delays-and-uncertainty-export-licensing-process-findings-csis-survey-us-technology.

The Economist. "Don't Blame AI for Your Job Woes." 6 November 2025. https://www.economist.com/finance-and-economics/2025/11/06/dont-blame-ai-for-your-job-woes.

European Commission. "Commission Proposes Tech Sovereignty Package to Strengthen Europe's Digital Autonomy and Resilience." 3 June 2026. https://digital-strategy.ec.europa.eu/en/news/commission-proposes-tech-sovereignty-package-strengthen-europes-digital-autonomy-and-resilience.

Euronews. "European Carmakers Face EU Pressure to Diversify Chip Suppliers." 21 May 2026. https://www.euronews.com/my-europe/2026/05/21/european-carmakers-face-eu-pressure-to-diversify-chip-suppliers.

Executive Order 14409. "Promoting Advanced Artificial Intelligence Innovation and Security." 2 June 2026. https://www.govinfo.gov/content/pkg/FR-2026-06-05/pdf/2026-11415.pdf.

Financial Times. "Anthropic Scrambles after Trump Administration Freezes Its Top AI Models." 14 June 2026. https://www.ft.com/content/f6940d59-28f4-4ae4-a569-c6fc421e52b9.

Financial Times. "China Adds Domestic AI Chips to Official Procurement List for First Time." 10 December 2025. https://www.ft.com/content/83c6521e-fe42-49e2-a9fe-eda97168b316.

Financial Times. "China Boosts AI Chip Output by Upgrading Older ASML Machines." 18 December 2025. https://www.ft.com/content/d10398db-b8b4-40f3-8c6d-b340470f5f3c.

Financial Times. "Chipmaker TSMC to Invest Another $100bn in US Production." 16 July 2026. https://www.ft.com/content/491927e1-1532-486d-94ec-8d6ee2de7bcd.

Financial Times. "Chips, Ships and Guns: South Korea Booms on AI Race and Global Conflict." 7 June 2026. https://www.ft.com/content/208d2ed8-78f3-471f-839a-55d69a3a8a9b.

Financial Times. "EU Wants Crisis Powers to Seize Control of Chip Supplies." 28 May 2026. https://www.ft.com/content/9d7d6204-4fc7-4f1d-af05-473c3649efcd.

Financial Times. "Netherlands to Restrict Chip Exports after US Pressure over China Threat." 8 March 2023. https://www.ft.com/content/e911774c-a048-4ed1-9f90-e4bb684a3156.

Financial Times. "Nvidia's Banned AI Chips Double in Price on China's Black Market." 23 June 2026. https://www.ft.com/content/57fcd3ce-464f-4dc2-8ea2-5712d4972c69.

Financial Times. "Robot Nation: China's Bid to Beat Its Demographic Decline." 24 June 2026. https://www.ft.com/content/c8731833-10ca-4a12-bfe4-8ebb2584ec68.

Financial Times. "Samsung and SK Hynix Plan $600bn Chipmaking Expansion." 29 June 2026. https://www.ft.com/content/86013b7e-41da-445a-981c-075a701dccf6.

Financial Times. "White House Lifts Ban on Anthropic Models." 30 June 2026. https://www.ft.com/content/137ddb71-852f-438c-ad76-25e2dc43486b.

The Guardian. "The White House's Plan to Vet Potentially Dangerous AI Is Cloaked in Secrecy." 7 August 2026. https://www.theguardian.com/technology/2026/aug/07/white-house-ai.

Institut Polytechnique de Paris. "Gaia-X: The Bid for a Sovereign European Cloud." 18 June 2025. https://www.polytechnique-insights.com/en/columns/digital/gaia-x-the-bid-for-a-sovereign-european-cloud/.

KPMG. "Global AI Pulse Q2 2026." 2026. https://kpmg.com/xx/en/our-insights/ai-and-technology/ai-pulse.html.

Lichtinger, Guy, and Seyed Hosseini Maasoum. "Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data." 2025. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5425555.

National Institute of Standards and Technology. "TSMC Arizona." Accessed 17 June 2026. https://www.nist.gov/chips/tsmc-arizona-phoenix.

New York Times. "AI Is Reshaping the Economy. Good Luck Measuring How." 2 July 2026. https://www.nytimes.com/2026/07/02/business/economy/ai-economy-data.html.

New York Times. "Big Companies Aim to Ease AI Transition for American Workers." 25 June 2026. https://www.nytimes.com/2026/06/25/business/economy/ai-work-force-training-job-losses.html.

New York Times. "Google Avoids Harshest Penalties in Landmark Search Monopoly Ruling." 2 September 2025. https://www.nytimes.com/2025/09/02/technology/google-search-antitrust-decision.html.

New York Times. "Is AI Replacing Tech Workers or Providing an Excuse for Job Cuts?" 1 June 2026. https://www.nytimes.com/2026/06/01/technology/ai-tech-job-cuts.html.

New York Times. "Nearly 200 Economists and Tech Leaders Warn of AI Threats." 13 July 2026. https://www.nytimes.com/2026/07/13/business/economists-ai-threat-jobs.html.

New York Times. "Trump Signs Executive Order Seeking Oversight of AI Models." 2 June 2026. https://www.nytimes.com/2026/06/02/technology/trump-executive-order-ai.html.

Nikkei Asia. "How Arizona Is Cultivating Talent for a US Chip Revival." 31 December 2025. https://asia.nikkei.com/business/tech/semiconductors/how-arizona-is-cultivating-talent-for-a-us-chip-revival.

Nikkei Asia. "Singapore Seizes $42m Home in Nvidia Chip Smuggling Case." 1 July 2026. https://asia.nikkei.com/spotlight/society/crime/singapore-seizes-42m-home-in-nvidia-chip-smuggling-case.

Nikkei Asia. "Taiwan GDP Grows Nearly 13% in Q2, Supercharged by AI and US Ties." 31 July 2026. https://asia.nikkei.com/economy/taiwan-gdp-grows-nearly-13-in-q2-supercharged-by-ai-and-us-ties.

Nikkei Asia. "TSMC's $165bn US Expansion Push: 'Inevitable' but Risky." 26 March 2025. https://asia.nikkei.com/business/technology/tech-asia/tsmc-s-165bn-u.s.-expansion-push-inevitable-but-risky.

OpenAI. "Accelerating the Next Phase of AI." 2026. https://openai.com/index/accelerating-the-next-phase-ai/.

Politico. "Judge Pauses Pentagon's Punishment for Anthropic." 26 March 2026. https://www.politico.com/news/2026/03/26/judge-pauses-anthropics-supply-chain-risk-designation-00847682.

Politico. "White House AI Vetting Plan to Exempt Lower-Cost 'Open' Models." 4 August 2026. https://www.politico.com/news/2026/08/04/white-house-ai-vetting-plan-to-exempt-nonproprietary-models-01024816.

Reuters. "AI Copyright Battles Enter Pivotal Year as US Courts Weigh Fair Use." 5 January 2026. https://www.reuters.com/legal/government/ai-copyright-battles-enter-pivotal-year-us-courts-weigh-fair-use-2026-01-05/.

Reuters. "Anthropic Agrees to Pay $1.5 Billion to Settle Author Class Action." 5 September 2025. https://www.reuters.com/sustainability/boards-policy-regulation/anthropic-agrees-pay-15-billion-settle-author-class-action-2025-09-05/.

Reuters. "Anthropic Wins Key US Ruling on AI Training in Authors' Copyright Lawsuit." 24 June 2025.

https://www.reuters.com/legal/litigation/anthropic-wins-key-ruling-ai-authors-copyright-lawsuit-2025-06-24/.

Reuters. "Meta Fends Off Authors' US Copyright Lawsuit over AI." 25 June 2025. https://www.reuters.com/sustainability/boards-policy-regulation/meta-fends-off-authors-us-copyright-lawsuit-over-ai-2025-06-25/.

Reuters. "Pentagon Designates Anthropic a Supply Chain Risk." 5 March 2026. https://www.reuters.com/technology/pentagon-informed-anthropic-it-is-supply-chain-risk-official-says-2026-03-05/.

Taipei Times. "Analysis: How AI Has Rewired Taiwan's Stock Market." 31 July 2026. https://www.taipeitimes.com/News/feat/archives/2026/07/31/2003861680.

U.S. Department of Commerce, Bureau of Industry and Security. "Department of Commerce Revises License Review Policy for Semiconductors Exported to China." 13 January 2026. https://www.bis.gov/press-release/department-commerce-revises-license-review-policy-semiconductors-exported-china.

U.S. Department of Commerce, Bureau of Industry and Security. "Revision to License Review Policy for Advanced Computing Commodities." Federal Register 91, no. 10 (15 January 2026): 1684. https://www.federalregister.gov/documents/2026/01/15/2026-00789/revision-to-license-review-policy-for-advanced-computing-commodities.

United States Patent and Trademark Office. "Revised Inventorship Guidance for AI-Assisted Inventions." 26 November 2025. https://www.uspto.gov/subscription-center/2025/revised-inventorship-guidance-ai-assisted-inventions.

United States et al. v. Google LLC. Memorandum Opinion. U.S. District Court for the District of Columbia, 5 August 2024. https://www.texasattorneygeneral.gov/sites/default/files/images/press/Google%20Search%20Engine%20Monopoly%20Ruling.pdf.

The Wall Street Journal. "Inside China's All-Out Push to Catch Up with American AI Chips." 23 July 2026. https://www.wsj.com/world/china/china-ai-chips-race-949050d0.

"We Must Act Now: A Statement on AI's Transformation of the Economy." 2026. https://www.wemustactnow.ai/.