Center for Geopolitics | JPMorganChase

May 2026

Helping Clients Navigate Global Challenges

# Beyond the Benchmarks:

## A Systemic View of U.S.-China AI Competition

JPMorganChase

# Key Takeaways:

- **Geopolitically, AI is becoming the whole ballgame.** Until recently, U.S.-China AI competition has been framed mainly as a technological race: who has the most compute, the better model. But increasingly, AI's real-world impacts are turning it into a broader systemic contest, testing each country's governance and economic models. The stakes of this competition will only sharpen, along with the potential economic opportunities and downside strategic risks.
- **This is fundamentally a two-country race.** The United States remains ahead in frontier model development and advanced compute deployment. But China's deep bench of AI engineering talent, low-cost models, control over critical nodes in the hardware supply chain, vast energy infrastructure, and aggressive push to diffuse AI throughout the economy make it a formidable competitor.
- **While companies are driving AI's most consequential advances, their success will increasingly track their country's fortunes.** As they grapple with their own enterprise AI adoption and position for AI's economy-wide disruptions, they cannot ignore the geopolitical contest: in addition to their cybersecurity and regulatory exposures, firm competitiveness is increasingly correlated with the broader success of which country leads in AI.

# Table of Contents

# Introduction

U.S.-China AI competition is one of the defining geopolitical dynamics of the 21st century—and the pace at which its stakes are escalating is itself part of the story. In our October 2025 report, *The Geopolitics of AI*, we surveyed the full geopolitical terrain being reshaped by AI, from Europe's pursuit of tech sovereignty to Middle Eastern capital flows and defense transformation. In this report, **we focus on the competition that anchors all of them: the U.S.-China contest for AI leadership.**

**The release of Anthropic's Mythos and similar models has offered a glimpse of the competition's growing stakes.**[1] Mythos's ability to identify cyber vulnerabilities that have evaded even the most advanced human experts and AI systems represents not only a significant advance in software assurance, but also cyber capability—one that could widen capability asymmetries overnight and push civilian systems across finance, energy, telecommunications, logistics, and health onto the front lines of interstate rivalry and the predations of non-state actors.[2] China will develop comparable technology—and when it does, the competitive landscape will again shift.

As AI improves, **the gap between a technical milestone and a geopolitical inflection point is closing faster than governments or companies are prepared for.** At the same time, by over-indexing on fast-moving developments, decision-makers may underappreciate some of the AI competition's slower moving dynamics. **This report attempts to offer a framework that looks beyond snapshots of current capability to track leading indicators that reveal where this competition is heading.**

**The stakes of U.S.-China AI competition are high.** AI's potential to accelerate innovation and efficiency across technologies and sectors means that leadership here has the potential to translate into broader economic and first-mover advantages—and its applications across military and security domains make competition a national security imperative at a moment when the two countries are increasingly operating as peer competitors.

**And those stakes are not lost on either government.** Policymakers in Beijing and Washington have taken extraordinary steps over the last several years both to foster environments conducive to faster domestic AI innovation and adoption, and in some cases to widen the gap by restricting the other country's access to essential components of the AI supply chain. **Yet despite a shared recognition of AI's strategic importance, the two governments are making fundamentally different bets about where that competition will be won.**

- **The U.S. strategy remains concentrated at the frontier.** It is investing heavily in the most advanced hardware, backing the development of the world's leading models, and deploying abundant private capital in pursuit of raw capability leadership.
- **China's strategy is increasingly centered on deployment at scale.** China is competitive at the model level, but its real emphasis is broader diffusion: rapid adoption at scale across the domestic economy, development of lower-cost and "good enough" open-source models, aggressive exports of Chinese technology stacks into global markets, and the data advantages, standards influence, and geopolitical leverage that come from being the infrastructure provider of choice across much of the developing world.

These are not just different resource allocations—**they are different theories of how AI competition is won**. The U.S. bet is that frontier capability translates into unassailable strategic advantage. China's bet is that adoption, scale, ecosystem embedment, and integration of AI into physical systems will ultimately prove more consequential. **The two wagers are also manifestations of the two countries' defining approaches: a U.S. commitment to exquisite systems, intangibles, and services; and a Chinese commitment to the practical, tangible, and manufactured.**

[1] This assessment, and this report as a whole, is based entirely upon open source material and does not reflect any non-public insights gleaned via JPMorganChase's testing of Mythos or other similar models.

[2] Goldstein, Gordon, "Six Reasons Claude Mythos Is an Inflection Point for AI—and Global Security," Council on Foreign Relations, 15 April, 2026, https://www.cfr.org/articles/six-reasons-claude-mythos-is-an-inflection-point-for-ai-and-global-security.

## The Assessment Challenge

That basic difference is itself part of what makes assessing the state of U.S.-China AI competition so difficult. That difficulty is compounded further by an opaque information environment requiring specialized forms of expertise ranging from semiconductor manufacturing to power generation and uncertainty about what factors to privilege and how they interact. These dynamics, combined with the lack of incentive for either government—or the companies operating within them—to disclose either progress or constraints, means that the true state of the competition is rarely visible until a development forces it into the open.

The release of Chinese company DeepSeek's R1 model in January 2025 was exactly that kind of forcing event. The model's advanced performance and training efficiency revealed a far narrower gap between U.S. and Chinese AI capabilities than many had anticipated—and critically exposed the limits of the prevailing analytical framework. DeepSeek's newer V4 model, reportedly trained on Huawei's chips, remains behind the frontier but outperforms open-source models in the U.S., and is only one among many leading open-source models in China. **Simply tracking hardware comparisons, where the U.S. maintains a clear lead, and existing model architectures doesn't capture the full competitive landscape**. China's AI competitiveness has begun to unsettle the pervasive assumption in some U.S. circles that China can imitate but not innovate—a fallacy that had given some policymakers and companies a false sense of security.

**DeepSeek did not change the underlying competition; it revealed how much of it had been happening out of view. Now Mythos is being characterized by some in China as a "reverse DeepSeek."**

## Our Framework, and Its Findings

**What's needed is a framework that looks more holistically at the challenge, and further over the horizon**—one that tracks leading indicators across the full scope of the AI landscape to assess how U.S. and Chinese capabilities compare, where each country is directing its efforts, and where the next shift might emerge.[3] This paper attempts to do exactly that. Rather than evaluating AI competition through a narrow technological lens, **we examine the evolving structural foundations on which future AI dominance will be built, across seven primary dimensions:**

1. Policy
2. Hardware
3. Models and Software
4. Finance
5. Socioeconomics
6. Energy
7. Military and Security

Companies can use such a framework to better anticipate event risks, articulating what would constitute a material step change in any one dimension.

[3] The authors would like to acknowledge the tremendous work of the Harvard Kennedy School Belfer Center, The Stanford Institute for Human-Centered AI, and others in tracking components of this competition. This framework aims to build on this foundational research.

Across those dimensions, several findings stand out:

- **U.S. hardware leadership remains the clearest and most consequential advantage** in the competition—but it is constrained by supply chain exposure and a widening energy gap relative to China.
- **At the model level, the gap is narrower than benchmarks suggest,** and Chinese open-source models, developed by a deep bench of engineering talent, are capturing global market share in ways that are not just commercial, but increasingly about standards-setting, data governance, and ecosystem influence.
- On finance, the **U.S. capital advantage is overwhelming—but the productivity gains that justify the investment thesis have not yet materialized** at a macroeconomic level, and concerns about a spend bubble persist.
- And in the domains where AI-enabled competition is already most active—cyber, influence operations, and the early stages of military application—**established norms are lagging dangerously behind capabilities.**

**Our core conclusion is that the U.S. maintains a meaningful overall lead—but that lead is narrower, more contested, and more fragile in key areas than conventional wisdom suggests. Critically, the dimensions where China is gaining ground are precisely those that are hardest to reverse.**

This scorecard *(Figure A)* offers a snapshot of U.S.-China competition across the AI ecosystem today and where it is heading. By its nature, any such snapshot involves qualitative judgment and simplification; we offer it as a guidepost to contextualize the analysis in the report.

*Figure A.* **AI Ecosystem Scorecard**

*Assessing U.S. vs. Chinese advantage and trajectory across different layers of the AI ecosystem.*

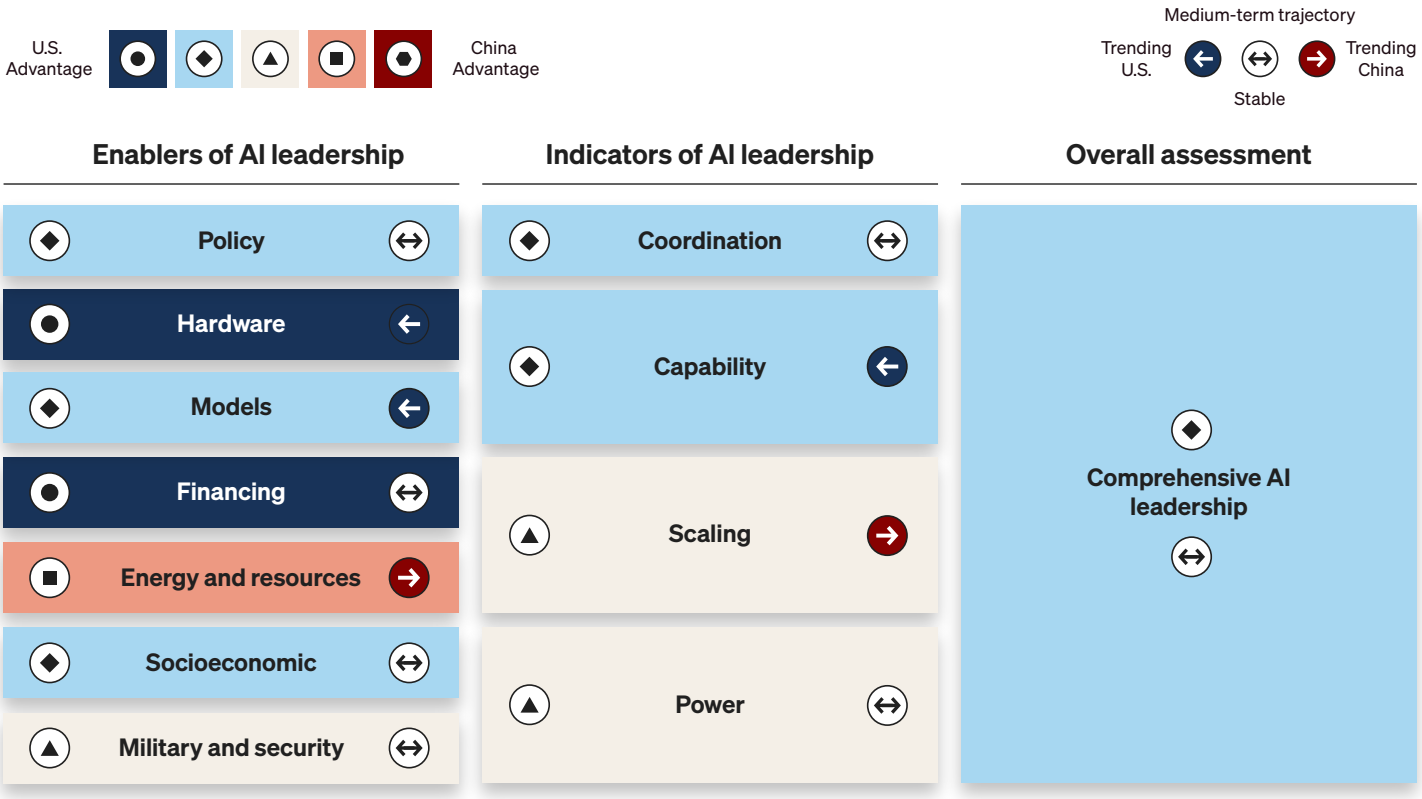


*Advantage and trajectory assessments based on the authors' qualitative assessments of the state of play in each area.*

# Section One:

## Why This Matters for Companies

U.S.-China AI competition is not a dynamic that companies can observe from a distance. The infrastructure businesses run on, the capital markets they access, the regulatory regimes they navigate, and the threat environment they operate within are all being actively shaped by this contest. **The firms best positioned in the coming decade will be those that understand not just where the competition stands, but where it is heading.**

- **Geopolitically-driven policies are already reshaping access to hardware, capital, and data.** Export controls on advanced semiconductors have constrained which chips companies can acquire and where they can deploy AI infrastructure. Investment restrictions are redrawing the boundaries of cross-border capital flows. And the data dimension is sharpening: the forced restructuring of TikTok—a platform with 170 million U.S. users—illustrated how concerns about which jurisdiction governs user data can translate into existential business disruption. As AI makes data more strategically valuable, regulatory pressure on cross-border data flows will intensify.
- **AI-enabled cyber threats are entering a new phase—and companies are on the front lines.** Mythos and other frontier models have demonstrated that frontier AI models can identify vulnerabilities in critical systems at an unprecedented speed. That capability will not remain exclusive to the United States. As China develops comparable models, companies already targeted by geopolitically-motivated threat actors—in finance, energy, defense supply chains, and critical infrastructure—will face a more dangerous environment. Companies that have historically viewed nation-state cyber threats as a government problem are finding that the line between corporate and national security infrastructure is blurring faster than their security postures have adjusted.
- **Most importantly, the competitive trajectory of leading AI companies matters for everyone.** The largest U.S. and Chinese AI firms now account for a significant share of global market capitalization, capital expenditure, and growth. J.P. Morgan Research estimates that U.S. firms spent $490 billion on AI in 2025, and that figure is projected to climb further.[4] **Investment at a scale that makes the health and competitiveness of a small number of companies a macroeconomic variable, not just a technology sector story.** But even this may understate their influence: the sweeping economic impact of AI—from how work gets done to how entire industries operate—means that **a company's prospects are increasingly tied to whether it is built on a winning AI platform.** As U.S. and Chinese AI systems spread across the globe, they will redraw the map of economic opportunity, determining which companies, in which markets, get early and advantaged access to the next wave of growth.

[4] Lakos-Bujas, Dubravko, et al., "Global equity strategy: market update, AI & security resurgence, earnings accelerate amidst geopolitical uncertainty," J.P. Morgan, 21 April 2026

# Section Two:

## The Framework: Seven Dimensions

U.S.-China AI competition cannot be understood through any single lens—and for that reason, neither can it be meaningfully assessed through any single metric. **We organize our analysis across seven dimensions because the competition is playing out simultaneously across distinct, but deeply interconnected arenas, each of which offers its own signals about where the balance of advantage is shifting.**

We begin with **policy,** which occupies a category of its own: it both **shapes and responds to every other dimension in this framework,** setting the rules that govern technological development, capital allocation, and strategic competition. **We divide the remaining six dimensions into two groups.** The first is **technological—hardware, models and software, and energy**—covering the physical and computational substrate on which AI capabilities are built. The second is **systemic—finance, socioeconomics, and military and security**—covering the broader environments that determine how those capabilities are resourced, absorbed, and ultimately applied.

*Figure B.* **Framework for Assessing U.S.-China AI Competition**

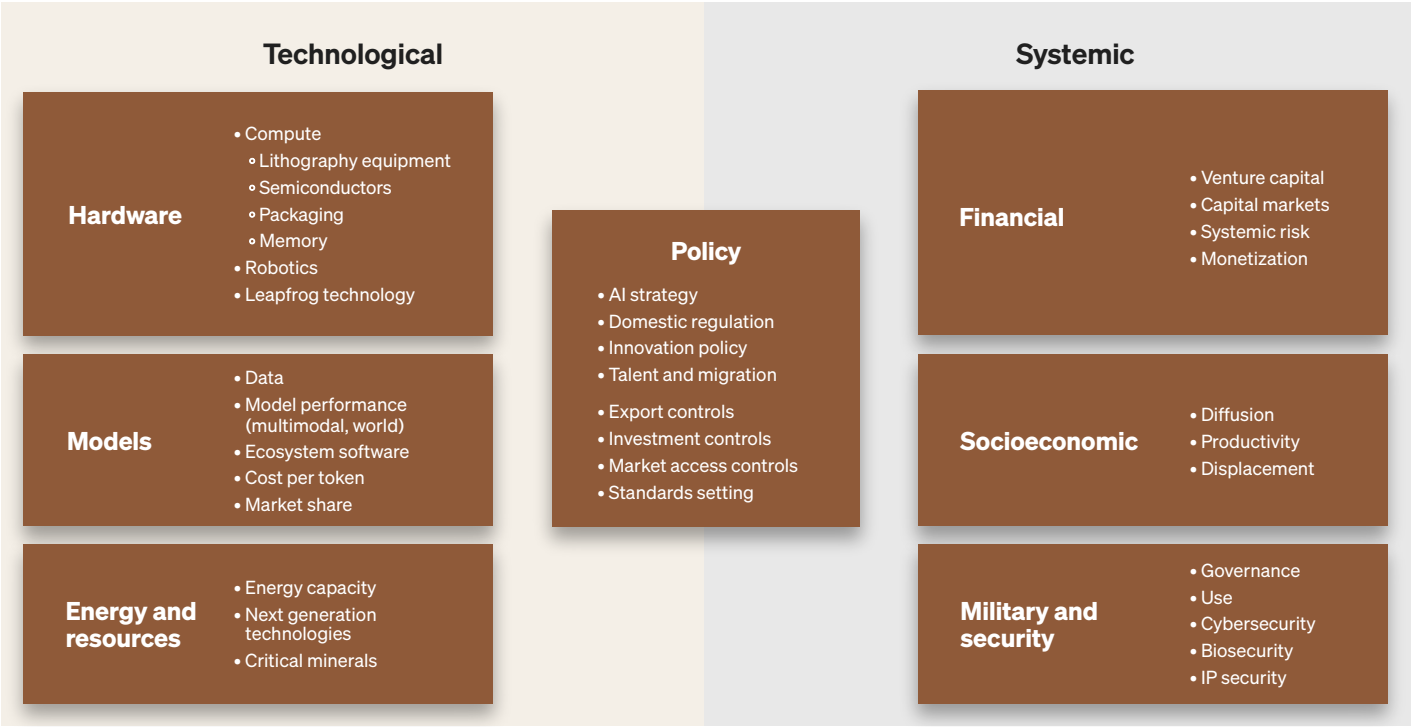

Some dimensions, like model benchmarks, generate regular headlines. Others—energy infrastructure buildout, capital market access, productivity trends in national accounts data—move quietly but may prove more consequential as leading indicators of where the competition is heading before that trajectory becomes visible in the capabilities themselves. A shift in China's grid capacity or the pace of U.S. permitting reform will not make the front page, but it will shape the frontier model race years before the next benchmark does.

**Critically, the dimensions do not operate independently.** Energy constraints shape what hardware scaling is possible. Model diffusion drives data accumulation, which feeds future model development. Policy choices around export controls determine which hardware is available to build on. Each section closes with a short watchlist of the leading indicators we consider most consequential for tracking how that dimension evolves.

*Figure C.* **Interconnected Enablers of AI Leadership**

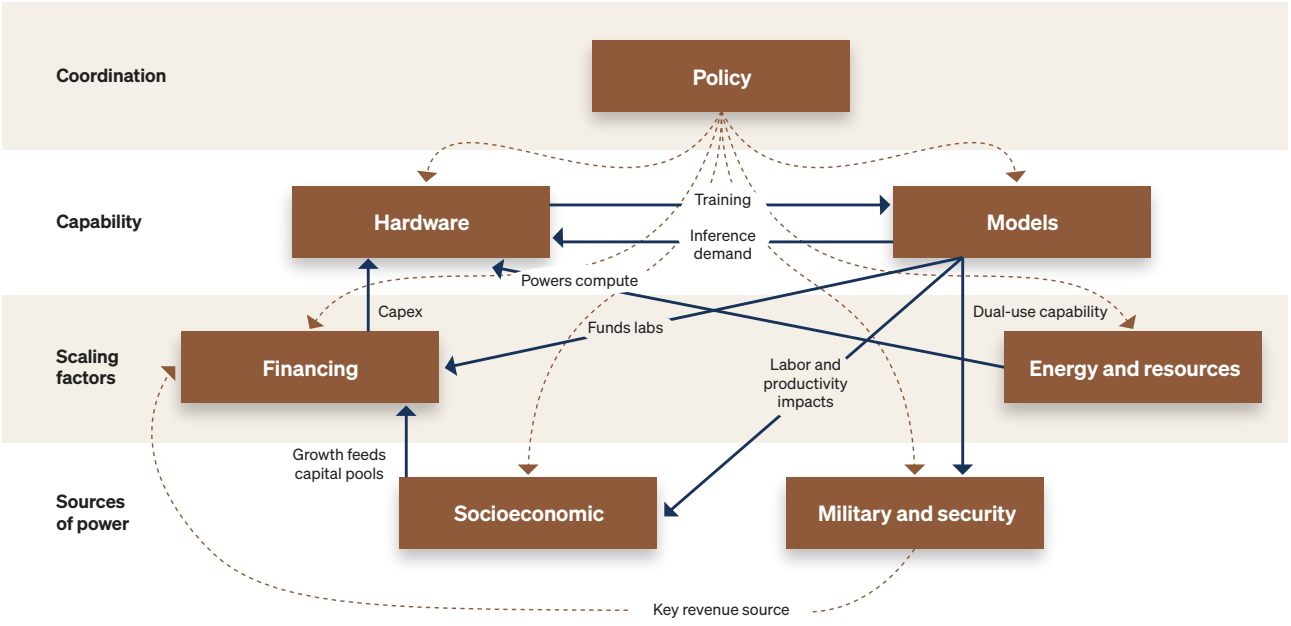


# Policy: Balancing Offense and Defense

Policy is the dimension that touches everything else in this competition—shaping where companies can sell hardware, how they can access data to train models, and the flows of energy, talent, and capital available to players across the AI ecosystem. **Understanding the policy landscape requires holding two distinct categories simultaneously:**

- **Policies aimed at *running faster,*** which create conditions for domestic actors to innovate, scale, and deepen AI adoption;
- **Policies aimed at *playing defense,*** which seek to preserve technological advantage by constraining the other side's access to critical inputs.

**Key underlying assumptions can have a major impact on the relative balance between those two approaches.** In the Biden administration, the assumption that artificial general intelligence was possible and that there would be irreversible consequences if China achieved it first tipped the balance towards defense.[5] With the notable exception of significant industrial policy commitments to partially reshore advanced semiconductor production, the Biden administration placed much more emphasis on the defensive side of AI competitiveness.

**The Trump administration has reoriented U.S. strategy, placing greater weight on commercial leadership as a geopolitical asset.** China's stranglehold on critical minerals supply chains has exposed the limits of a denial-based approach—Washington has been reticent to retaliate when it remains dependent on critical Chinese inputs.

The AI industry has reinforced this shift, with leading firms making the case in Washington that export and investment restrictions undercut the revenues needed to stay at the frontier. **The net result is a U.S. policy no longer anchored to a single governing assumption, but one evolving into a more flexible, "all-weather" approach—built to accommodate multiple scenarios for how AI development unfolds.**[6,7]

## Running Faster

**China has a significant head start when it comes to strategic planning in AI. Its 2017 "New Generation Artificial Intelligence Development Plan" set explicit, phased targets:** catch up with leading AI nations by 2020, achieve major breakthroughs by 2025, and become the world's primary AI innovation center by 2030. Beijing's recently passed 15th Five-Year Plan doubles down on that trajectory, positioning AI as the fundamental enabler of industrial upgrading and prioritizing self-reliance in semiconductors, computing power, and "embodied intelligence"—with the **goal of integrating "AI+" across the economy, science, and governance.**

Washington's approach looks different. **The U.S. AI Action Plan, released in July 2025, prioritizes accelerating frontier innovation**, building energy and infrastructure, and securing supply chains. Where China plans for diffusion and scale, the United States is betting on speed and breakthrough capability.

> **"Where China plans for diffusion and scale, the United States is betting on speed and breakthrough capability."**

[5] Froman, Michael, "China, the United States, and the AI Race," Council on Foreign Relations, 10 October, 2025, https://www.cfr.org/articles/china-united-states-and-ai-race

[6] Sullivan, Jake and Tal Feldman, "Geopolitics in the age of artificial intelligence," Foreign Affairs, 27 January 2026, https://www.foreignaffairs.com/united-states/geopolitics-age-artificial-intelligence.

[7] Some scholars argue that the U.S. can yet go further, adopting a "rideout" strategy that will position it regardless of whether the U.S. achieves AGI, if it is indeed possible, first. This entails taking steps to mitigate the risk that an adversary might attack to disrupt U.S. progress towards AGI and also taking steps to mitigate the harm were an adversary to gain an AI-enabled military advantage. See Frelinger, David and Karl Mueller, "The AGI rideout strategy for reducing strategic risk and promoting stability in the transition to Artificial General Intelligence," RAND, April 2026, https://www.rand.org/pubs/perspectives/PEA4347-1.html

**That distinction matters. Both governments have made fostering a pro-innovation environment a priority, but with meaningfully different emphases.**

- **The United States has focused primarily on enabling frontier development—**through energy permitting reform designed to accelerate data center buildout, executive action to reduce regulatory friction for AI companies, and the Stargate initiative committing hundreds of billions in AI infrastructure investment.
- **China's approach prioritizes diffusion and scale.** Its industrial policy, reflected in successive Five-Year Plans, pushes AI adoption across entire sectors of the economy rather than concentrating support at the frontier, with the goal of compounding adoption-driven advantages in data, cost, and standards.

History suggests this is not a trivial distinction. Previous periods of rapid technological change—including the Industrial Revolution—have shown that **widespread adoption and integration of new tools can matter as much as, or more than, leading-edge discovery in determining which societies benefit most.**[8]

On balance, the U.S. regulatory environment remains more pro-innovation than China's, where state oversight of algorithms, content, and data flows imposes constraints that have no direct U.S. equivalent—but the gap is narrower than it is sometimes portrayed.

## Playing Defense

**Even as both sides observe a period of informal restraint in the wake of a destabilizing period of export control escalation in 2025, neither is standing still.** Washington and Beijing are each building out defensive frameworks designed to protect strategic advantage—but the harder questions about where to draw those lines, and how to update them as the technology frontier shifts, remain unresolved.

- **In the United States, limitations on private U.S. investment in Chinese advanced AI and national security-related technology.** Translating that consensus into durable policy has proven more difficult. Congress is currently considering legislation that would expand oversight of chip sales and equipment to China, mandate on-chip location tracking to prevent diversion of advanced semiconductors, and tighten regulation of remote access to AI compute—measures that reflect the growing complexity of controlling technology that moves faster than the regulatory apparatus designed to govern it.
- **China is similarly building out its defensive toolkit: it has increased scrutiny of innovative domestic companies seeking to relocate abroad to access U.S. capital, blocked U.S. models from operating within China—limiting their access to Chinese data and foreclosing a major market—and is strengthening its own export control regime.** Beijing demonstrated in 2025 how that regime can be wielded as a strategic instrument, imposing restrictions on critical mineral exports that sent a clear signal about China's leverage over upstream inputs essential to the entire AI supply chain.

[8] Ding, Jeffrey, "The innovation fallacy," Foreign Affairs, 19 August 2024, https://www.foreignaffairs.com/china/innovation-fallacy-artificial-intelligence

## Background on U.S.-China AI Diplomacy

Experts in the U.S. and China have engaged in informal "track II" dialogues on AI safety and governance for over a decade. However, only after generative AI became widely adopted did the two governments begin to devote meaningful diplomatic attention to the topic. The Biden administration viewed AI as one area in which diplomacy was essential—to manage potential strategic and security risks—despite the otherwise competitive nature of the bilateral relationship. At their November 2023 summit, President Biden and Xi Jinping agreed to launch working-level dialogues to discuss "the risks associated with advanced AI systems." The first formal bilateral meeting was held in May 2024, culminating in an agreement between the two leaders later that year on "the need to maintain human control over the decision to use nuclear weapons." However, the overall discussions did not advance very far.

The second Trump administration has focused on cooperating with China on counternarcotics and immigration enforcement but has viewed AI more as a domain of strategic friction than one of alignment. In February 2026, both the United States and China attended the Responsible AI in the Military Domain (REAIM) summit but did not sign its outcomes document outlining 20 principles for military uses of AI.[9] Following the release of Mythos in April, and the revelations about its implications for cybersecurity, many analysts have speculated whether AI risks will once again become a focus of bilateral diplomacy.

After the summit between President Trump and Xi Jinping in Beijing in May 2026, China's Foreign Ministry announced the establishment of an intergovernmental bilateral AI dialogue. It remains to be seen whether this channel results in tangible changes to either side's AI development or serves primarily as a confidence-building mechanism and forum to exchange views on respective approaches to AI policy.

## What We're Watching: Policy

- **The "Mythos moment":** Will the advent of Mythos and comparable models make the risks of next-generation AI visible enough to catalyze a shift in momentum in the U.S., galvanizing support for export controls and on-chip restrictions? Will it compel U.S. and Chinese leadership to pursue a degree of AI cooperation to manage risks?
- **National AI legislation:** This summer, the U.S. Congress is expected to focus on developing a legislative framework for AI. The contours of that legislation will shape the innovation environment in the U.S. in meaningful ways.

[9] Waldersee, Victoria, "US, China opt out of joint declaration on AI use in military," Reuters, 5 February 2026, https://www.reuters.com/business/aerospace-defense/us-china-opt-out-joint-declaration-ai-use-military-2026-02-05/.

# Section Three:

## Technological Dimensions

### Hardware: The Clearest U.S. Advantage, and Its Limits

Advanced computing power is the physical foundation of AI progress. **Model improvements have generally followed a pattern in which superior hardware translates to stronger performance, meaning that access to cutting-edge chips and the data centers that house them is not a secondary question, but a primary determinant of who leads.**[10] Access to compute may be even more important for running frontier models than for training them, since serving millions of users over a model's lifetime can consume more resources than building it.

For every gigawatt of data center capacity China has installed, the United States has nearly 1.7 *(See Figure D)*. China's advanced processing capacity is further constrained by U.S. export controls on the most sophisticated semiconductors and the equipment used to manufacture them. Washington is determined to maintain that gap.

*Figure D.* **The Cloud Divide: Where the World's Data Centers Actually Are**
*Installed data center capacity in gigawatts*

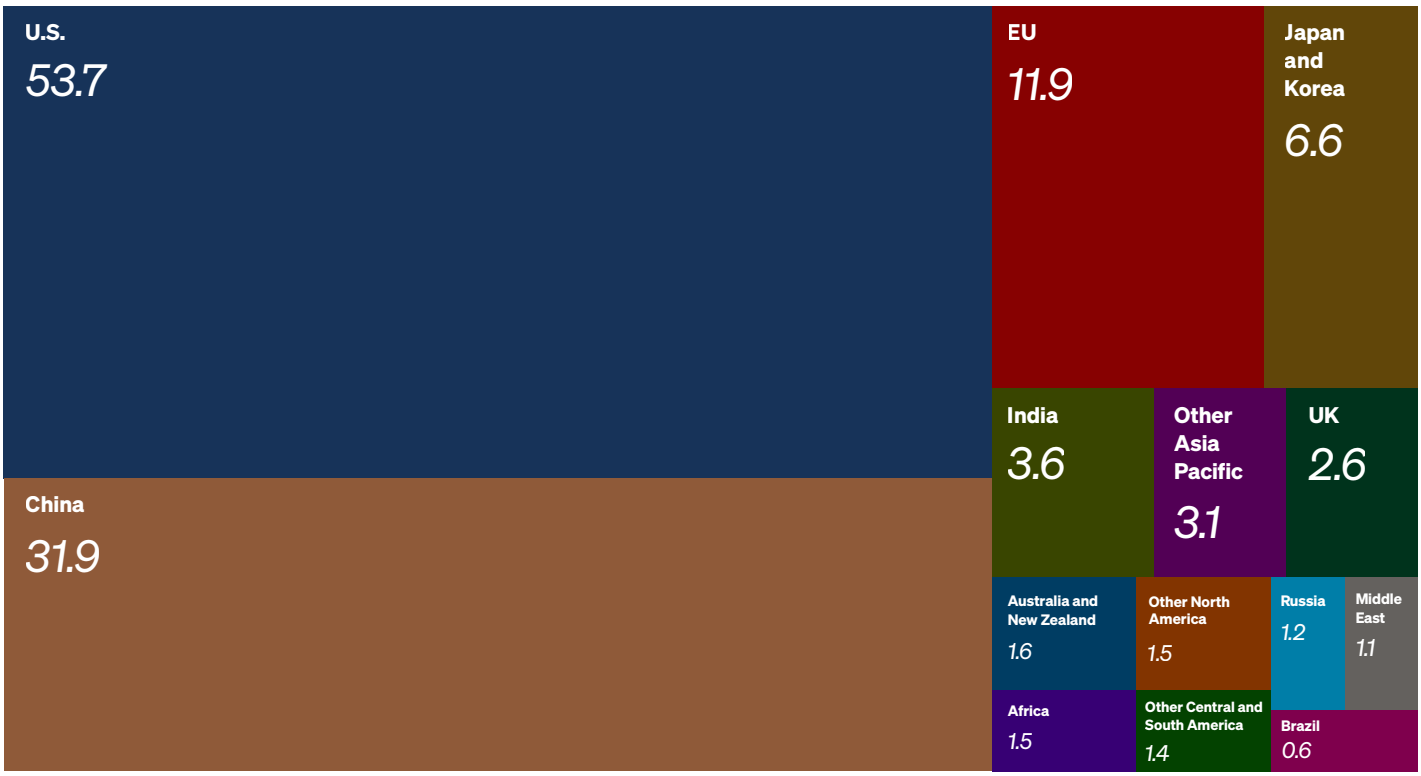


Source: Digital Information World, International Energy Agency, link

[10] Briski, Kari. " How Scaling Laws Drive Smarter, More Powerful AI AI Scaling Laws," NVIDIA Blog, 12 February 2025[Date], https://blogs.nvidia.com/blog/ai-scaling-laws/.

**The AI hardware value chain is one of the most complex global supply chains.** One chip requires thousands of components designed and manufactured by thousands of companies distributed across dozens of countries *(See Figure E).*[11] Despite this economic reality, governments around the world, including both Washington and Beijing, are leveraging industrial and trade tools to grow their share of the supply chain.

**The United States, along with its allies and partners, maintains a lead over China in many layers of the hardware stack powering frontier AI models.** American firms have long occupied leading positions in chip design, the specialized software used to create those designs, and the manufacturing equipment—such as etching, deposition, and testing tools—required to produce them.

But the U.S. was not always absent from manufacturing itself. American companies began adopting "fabless" models in the 1990s, offshoring capital-intensive fabrication and back-end packaging to specialize where margins and intellectual property were highest. That strategic choice delivered enormous commercial returns—but also created dependencies. In recent years, multinationals have begun scaling foundry capacity back onto U.S. soil, partly incentivized by shifts in tax, grant, and trade policy designed to reshore critical production.

**Crucially, America's advantage in hardware is not America's alone. U.S. allies and partners occupy key nodes in the AI accelerator supply chain, making their strategic positioning a decisive factor in U.S.-China AI competition:**

- **Taiwan** leads in advanced fabrication, substrate, printed circuit boards, and packaging
- **Japan** leads in wafers, chemicals, photoresist, lithography, etching, deposition, and substrate manufacturing
- **The Netherlands** leads in lithography
- **South Korea** is a leader in advanced fabrication and packaging, particularly for memory
- **Germany** is a leader in high-quality wafer production and lithography components
- **The UK** is a leader in semiconductor design.

*Figure E.* **Global Semiconductor Supply Chain: Key Node Leadership**

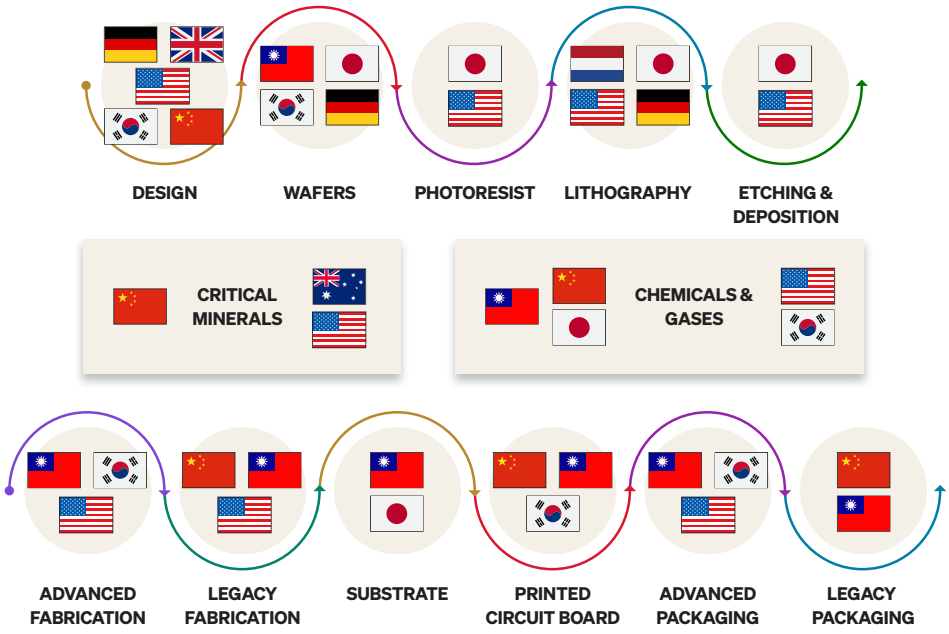


[11] Klimek, Peter, et al., "Mapping of the global semiconductor supply chain: embedding Austria in the global semiconductor inter-firm network," ASCII, May 2024, https://ascii.ac.at/wp-content/uploads/Mapping-of-the-global-semiconductor-supply-chain_policy-brief_ASCII.pdf.

The Netherlands, Japan, and South Korea have joined the U.S. in imposing multilateral export controls on China, constraining Beijing's ability to acquire or produce its own leading-edge chips.

## China's Hand: Strengths, Constraints, and Workarounds

China may have fewer friends, but it is not without leverage. As we noted in our recent Center for Geopolitics Report, *Breaking the Critical Minerals Chokepoint*, **China maintains a dominant position in many of the upstream critical minerals and chemicals that go into today's AI chips. Chinese manufacturers also dominate global production of many of the unglamorous but essential components that keep data centers running**—transformers, switchgear, batteries, and printed circuit boards. These are not frontier technologies, but they are the connective tissue of the entire AI hardware ecosystem. Finally, China maintains a large share of global printed circuit board manufacturing and legacy packaging capacity, which fuel its downstream consumer electronics, appliance, automotive, and other manufacturing sectors.

In recent years, Chinese firms have made strides in AI accelerator design, as evidenced by Huawei's Ascend 910C and next generation 950PR and 950DT. **Huawei has also pioneered innovative rack-scale architectures that interconnect hundreds of AI accelerators together to form powerful computing clusters.** On paper, these systems approach the computational performance of the worlds best hardware, but they operate at lower power efficiency and face manufacturing bottlenecks from export controls.[12]

**China can design competitive AI chips, but it cannot yet produce them in the volumes that frontier AI demands.** China's semiconductor manufacturers remain years behind their international peers in the advanced logic and memory processes that would be required to scale Huawei's AI accelerator deployment.

- **Much of the reason comes down to a single, critical piece of equipment: the extreme ultraviolet lithography (EUV) machine, the most advanced tool in semiconductor manufacturing, which China cannot access due to export controls.** Without it, Chinese manufacturers have found creative workarounds—using older equipment with complex multi-step processes to produce chips that approach the frontier—but at higher cost, lower reliability, and significantly lower output than their international peers.[13] Some in the U.S. Congress have proposed extending export controls to older generation lithography equipment to further obstruct China's ability to produce advanced AI accelerators.[14]
- **Beijing is well aware of this vulnerability—and is investing heavily to close it. The Chinese government has poured resources into building a domestic semiconductor equipment industry,** with firms like Shanghai Micro Electronics Equipment (SMEE) and Advanced Micro-Fabrication Equipment (AMEC) developing homegrown alternatives to Western-made tools. But indigenizing an equipment supply chain of this complexity is a generational undertaking, not a five-year plan. These firms have made progress on older-generation processes, but their ability to manufacture at the advanced levels required for frontier AI chips remains limited.[15] **JPMorgan Research estimates that a local lithography manufacturer in China will not pose a competitive threat to Netherlands-based global leader ASML for *at least a decade*.[16]**

[12] Patel, Dylan, "Huawei AI CloudMatrix 384 - China's answer to Nvidia GB200 NVL72," SemiAnalysis, 16 April 2025, https://newsletter.semianalysis.com/p/huawei-ai-cloudmatrix-384-chinas-answer-to-nvidia-gb200-nvl72.

[13] Fedasiuk, Ryan and Julia Torres, "The lithography loophole: how China is printing its way to chip self-sufficiency," AEI, 21 April 2026, https://www.aei.org/research-products/report/the-lithography-loophole-how-china-is-printing-its-way-to-chip-self-sufficiency/.

[14] Baumgartner, Michael, "Baumgartner Introduces Bipartisan Bill to Tighten Controls on Sensitive Chipmaking Equipment," 2 April 2026, https://baumgartner.house.gov/2026/04/02/baumgartner-introduces-bipartisan-bill-to-tighten-controls-on-sensitive-chipmaking-equipment/.

[15] Wentz, Jacob and Anita Lin, "Breakthroughs or boasts? Assessing recent Chinese lithography advancements," CSIS, 24 September 2025, https://www.csis.org/blogs/strategic-technologies-blog/breakthroughs-or-boasts-assessing-recent-chinese-lithography.

[16] Deshpande, Sandeep, et al., "Tech snippets," J.P. Morgan Research, 18 December 2025.

China maintains advantages in building data center infrastructure such as warehouses, power, water, telecommunications, and grid interconnects. However, as we noted, China's capacity to scale the "guts" of AI clusters, the advanced AI server racks, remains limited. As a result, estimates suggest that as many as 80 percent of China's data centers may be sitting idle.[17]

## What We're Watching: Hardware

- **U.S. export controls:** Whether the United States tightens or relaxes export controls on advanced chips and equipment, and whether U.S. allies and partners follow suit.
- **Chinese EUV lithography:** Whether Chinese chipmakers can scale advanced-node production without EUV lithography, as Chinese firms target 7nm and 5nm process nodes and aim to boost output.
    - Whether China succeeds in developing advanced lithography equipment or leveraging alternative semiconductor platforms—such as silicon photonics—that reduce dependence on cutting-edge lithography.

# Models and Software: Looking Beyond the Benchmarks

The AI competition between the United States and China is often framed as a race to build the most powerful model. That framing is incomplete—and increasingly misleading. **Raw capability, as measured by standardized benchmarks, is converging:** top models from both countries now cluster near the frontier on many tasks *(Figure F)*, and the gaps that remain on leaderboards often overstate real-world performance differences.[18] **What increasingly separates winners from losers in this competition is not who builds the best model, but who deploys it most effectively, most widely, and most durably.**

**Beyond leading model performance, several other dynamics will be decisive:**

- **Real World Deployment:** Benchmark gaps overstate real-world differences. Top models may converge near the frontier, but outcomes increasingly depend on prompt quality, tools, and domain integration. Chinese labs produce "good enough," affordable models to win at adoption.
- **Model Diffusion:** Widespread deployment of models is often underweighted relative to frontier capability. China's more open-model strategy, exported digital infrastructure, and large domestic rollout make diffusion a likely advantage multiplier.
- **Open vs. Closed Releases:** Open-weight models (those that share more of their back-end architecture publicly) can accelerate innovation and diffusion, while closed models result in slower uptake due to prioritizing security through controlled access and deliberate releases. However, because trust is a key factor in widespread adoption, it can turn out to be the more sustainable option over time.

[17] Chan, Caiwei, "China built hundreds of AI data centers to catch the AI boom. Now many stand unused," MIT Technology Review, 26 March 2025, https://www.technologyreview.com/2025/03/26/1113802/china-ai-data-centers-unused/.

[18] "2026 AI Index Report," Stanford University Human-Centered Artificial Intelligence, https://hai.stanford.edu/ai-index/2026-ai-index-report.

- **Ecosystem Advantage:** The U.S. leads not just in frontier models, but in the surrounding technology stack. This lowers deployment friction and creates stickiness with users. However, China is innovating and investing in alternative stacks, which could erode the U.S. edge.

*Figure F.* **Chinese Models Closing the Performance Gap**
*Comparison of aggregate capabilities over time of the most capable publicly released U.S. and PRC models according to a suite of benchmarks covering five domains (cyber, software engineering, natural sciences, abstract reasoning, and mathematics).*

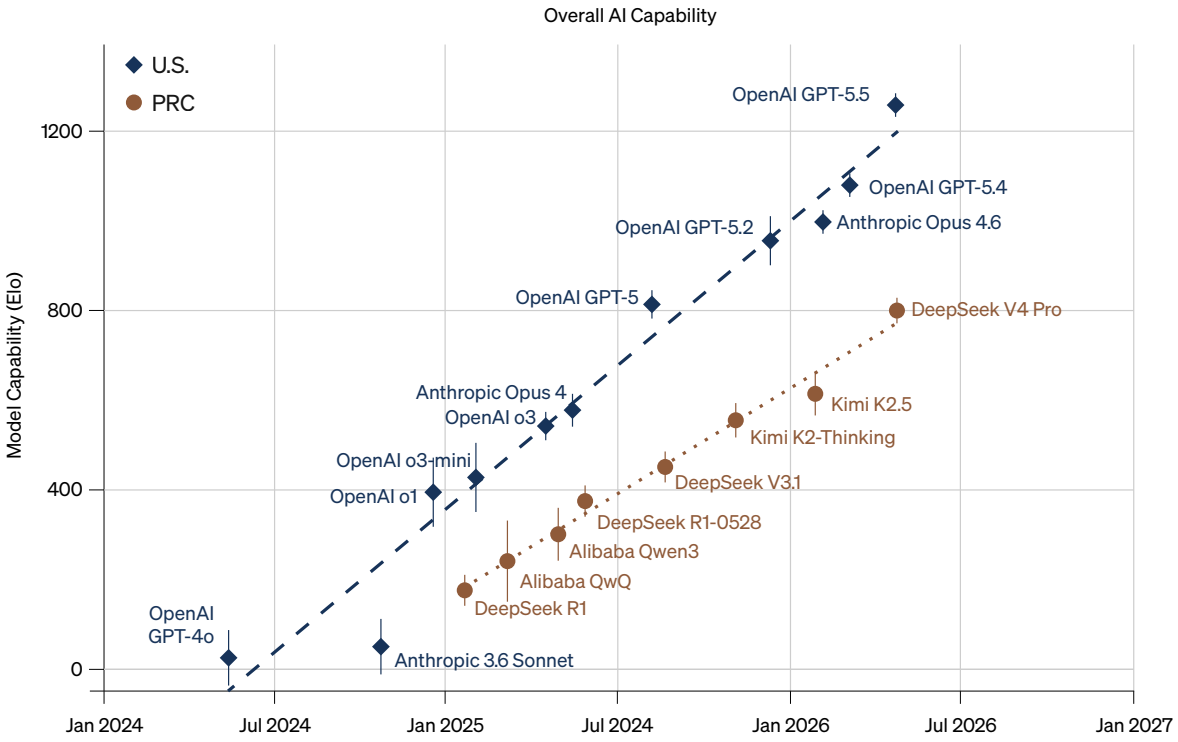


Source: U.S. Center for AI Standards and Innovation, link

## Deployment and Diffusion

**China's model strategy is built on a simple but powerful insight: in most real-world applications, a model that is 90 percent as capable but dramatically cheaper to run will outcompete a frontier system that fewer users can afford.** Chinese labs—including DeepSeek, Alibaba, and Baidu—have rapidly improved model quality while competing aggressively on the cost of running their systems.[19] That matters because in AI, cost is what determines how far a model travels. A cheaper model reaches more firms, more government agencies, and more markets—and each deployment generates data and feedback that makes the ecosystem stronger. Even where U.S. frontier models retain a technical edge, Chinese models are winning deployments where "good enough" at the right price is what matters.

**Diffusion, in other words, how widely and quickly models are adopted across geographies and sectors, is an underemphasized variable.**[20] China's approach combines open or partially open model releases with international digital infrastructure investments that give its models built-in distribution channels. The Digital Silk Road, a key component of the Belt and Road Initiative, has already expanded Chinese cloud services, telecommunications infrastructure, and digital platforms across parts of Southeast Asia, Africa, and Latin America.[21]

[19] "American labs say China's AI tigers are copycats," Economist, 25 February 2026, https://www.economist.com/china/2026/02/25/american-labs-say-chinas-ai-tigers-are-copycats.

[20] Egan, Janet and Spencer Michaels, "Five objectives to guide U.S. AI diffusion," CNAS, 29 April 2025, https://www.cnas.org/publications/commentary/five-objectives-to-guide-u-s-ai-diffusion

[21] "Assessing China's Digital Silk Road Initiative," Council on Foreign Relations, https://www.cfr.org/china-digital-silk-road/

**AI models deployed through these channels do more than serve users—they shape data governance practices, technical standards, and regulatory expectations in recipient countries. Estimates suggest that global downloads of Chinese and U.S. models reached parity in 2025.**[22] China's domestic market amplifies the effect further: large-scale deployment within its own economy generates feedback loops that strengthen both model performance and the application ecosystems built around them.

**The United States, by contrast, has generally prioritized frontier capability over mass diffusion—a bet that the best models will ultimately command the most valuable markets.** That bet may prove correct. But history suggests it is not self-evident: the technologies that reshape economies are not always the most advanced—they are the ones that get adopted.

## Open vs. Closed: The Trust Tradeoff

One of the sharpest strategic divergences between the two countries lies in how they release their models—and the downstream consequences of that choice.[23]

Although this trend may be shifting now, **Chinese releases have generally favored openness**[24]—making model weights more accessible to developers, researchers, and foreign governments *(See Figure G)*. DeepSeek claims that its recently released V4 model outperforms all other open-weight models on key metrics but trails state-of-the-art frontier models by approximately 3 to 6 months.[25] The open-weight approach accelerates experimentation, lowers barriers to adoption, and fuels the diffusion strategy described above. However, it also brings lower returns to model developers, which is one reason some Chinese firms have begun to tighten access to model weights. **U.S. firms, by contrast, have generally favored more controlled access,** citing safety, security, intellectual property protection, and the risk of misuse.[26]

**The result is a structural tension: openness drives faster uptake, but controlled access offers greater safeguards and higher potential profits.** Both countries are actively searching for the right balance—and how each calibrates that tradeoff will do more than shape their respective AI industries. It will influence the global governance norms that emerge around transparency, auditability, and security.

That tension, however, is more nuanced than it first appears—and the conventional wisdom that "open wins" may not hold over time. Adoption does not correlate linearly with openness. Where users distrust the ecosystem (e.g., unclear provenance, weak security guarantees, uncertain liability) adoption can begin to slow even when models are widely available. In that framing, **while tighter controls in the near-term may delay broad experimentation and diffusion, it helps to build the trust needed for more durable adoption, suggesting that "slow and steady" could outperform "fast and open" over time.**

> **"slow and steady" could outperform "fast and open" over time.**

[22] Qian, Z. China's AI Landscape: a free-for-all, not a central plan: what 6000+ filings with regulators reveal," 2026, ChinaTalk, https://www.chinatalk.media/p/chinas-ai-landscape-a-free-for-all.

[23] OECD, "Artificial intelligence and competitive dynamics in downstream markets," OECD Roundtables on Competition Policy Papers, 14 November 2025, https://www.oecd.org/en/publications/artificial-intelligence-and-competitive-dynamics-in-downstream-markets_ccf0624a-en/full-report/component-5.html.

[24] https://hai.stanford.edu/policy/beyond-deepseek-chinas-diverse-open-weight-ai-ecosystem-and-its-policy-implications

[25] McGuire, Chris et al. "DeepSeek v4 Signals a New Phase in the U.S.-China AI Rivalry," 29 April 2026, https://www.cfr.org/articles/deepseek-v4-signals-a-new-phase-in-the-u-s-china-ai-rivalry.

[26] Anthropic, "Responsible Scaling Policy: Version 3.0," 24 February 2026, https://www.anthropic.com/news/responsible-scaling-policy-v3.

## Ecosystem Depth

Looking beyond frontier models, **the U.S. retains a significant advantage in the broader AI software ecosystem.** Elements such as training frameworks, developer tools, and cloud platforms continue to be dominated by U.S. firms such as NVIDIA, Amazon Web Services, Microsoft Azure, and Google Cloud. Once organizations build their AI systems on these platforms, they are less likely to switch—giving the U.S. an advantage that deepens over time and extends well beyond any single model.

- However, the U.S. lead here is not insurmountable. **Chinese firms are investing heavily in alternative software stacks and AI infrastructure designed to reduce reliance on U.S. technologies.**[27] If these efforts reach functional parity, the durability of the U.S. ecosystem lead could erode.

**One of the most significant threats to U.S. leadership is distillation—a set of techniques through which smaller, cheaper models are trained to replicate much of the performance of larger frontier systems by learning from their outputs.**[28] Distillation can significantly reduce compute requirements while preserving task performance, effectively compressing technological advantage, and potentially accelerating China's ability to develop advanced AI despite U.S. export controls on advanced hardware. In practical terms, **leadership at the frontier may translate only into temporary advantage if competitors can replicate capabilities at a lower cost.**[29]

*Figure G.* **Chinese Companies Have Made the Most Capable Open AI Models**
*How each company's best open-weight model ranks on LMArena, a crowdsourced AI benchmarking platform*

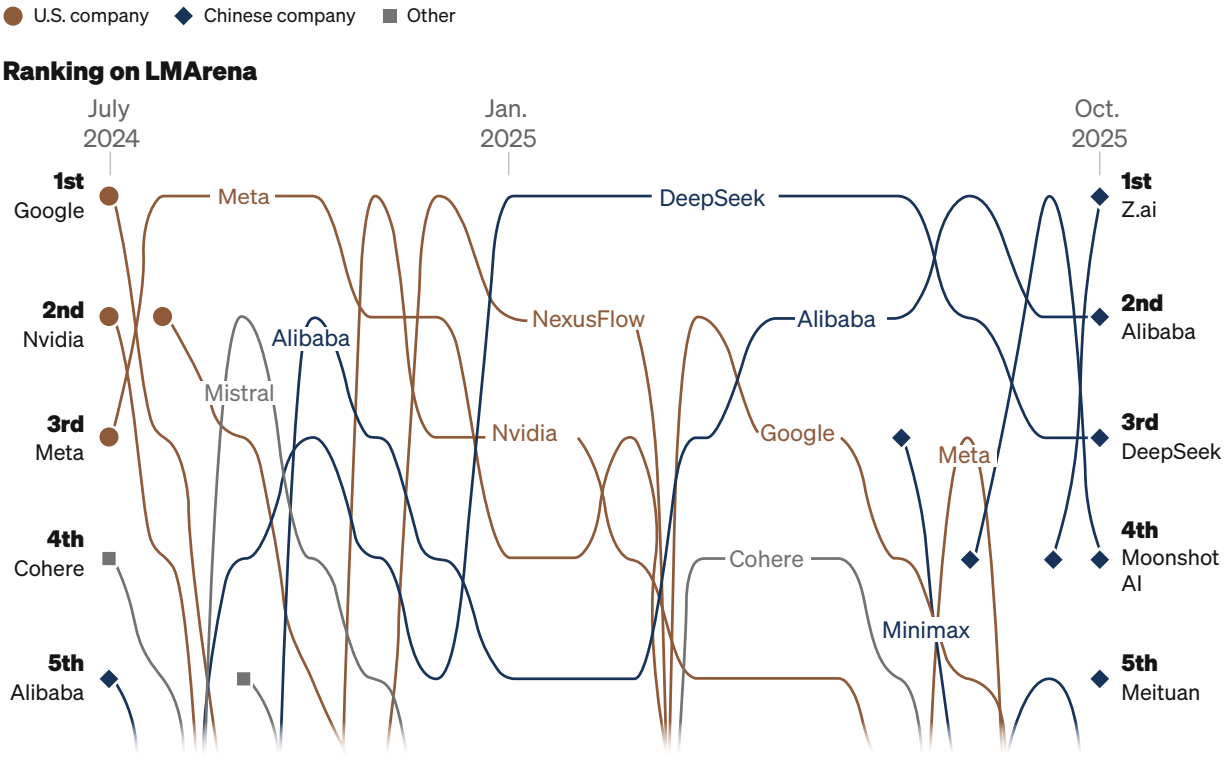


Source: LMArena

[27] Shivakumar, Sujai, "China's localization drive in semiconductors gains impetus from allied chip export controls," CSIS, 24 March 2026, https://www.csis.org/analysis/chinas-localization-drive-semiconductors-gains-impetus-allied-chip-export-controls.

[28] Luong, Ngor, "Two loops: how China's open AI strategy reinforces its industrial dominance," U.S.-China Economic and Security Review Commission, 23 March 2026, https://www.uscc.gov/research/two-loops-how-chinas-open-ai-strategy-reinforces-its-industrial-dominance.

[29] Sevastopulo, Demetri and Cristina Criddle, "White House accuses China of 'industrial-scale' theft of AI technology," Financial Times, 23 April 2026, https://www.ft.com/content/abde4e1e-c69a-4cc4-ad96-d88308314298.

At the same time, a new competitive frontier is opening up—one where the U.S. starts with an edge because of its ecosystem depth, but where the outcome is far from determined. Agentic AI—systems that go beyond answering questions to taking actions on a user's behalf—is widely seen as one of the next major leaps in AI capability. But building reliable agents depends less on the underlying model than on the infrastructure, governance, and institutional trust required to deploy them safely at scale. **The U.S. currently holds meaningful advantages in each of these areas. But the edge is narrow.** Standards for how agents operate and who is accountable when they fail are still being written—and whoever shapes those rules will hold a significant advantage in what may become a consequential layer of AI competition.

## What We're Watching: Models and Software

- **More controlled releases by Chinese firms:** Chinese AI firms are increasingly adopting hybrid release strategies, maintaining some open-weight models for ecosystem development while restricting access to more advanced systems through application programming interfaces (APIs) and enterprise channels. This reflects the limitation of the open-source approach, including monetization pressure, safety considerations, and the strategic value of frontier capabilities.[30] The convergence toward controlled openness could also be a signal of rising confidence in domestic capability.
- **DeepSeek V4:** What will pressure testing of DeepSeek's newly released V4 reveal about shifts in the frontier of Chinese model capability? First looks indicate it has achieved key architectural innovations in memory efficiency and attention mechanism. Moreover, it is the first model optimized for China's Huawei chip, the Ascend, which could signify the loosening of China's dependence on Nvidia.[31]
- **Inference costs:** Inference costs are a key diffusion indicator, representing the ratio of compute capacity needed for AI model usage. Will U.S. model developers be able to lower inference costs with future releases to more closely track with Chinese competitors on this metric?

[30] Bloomberg News, "Alibaba Targets $100 Billion of AI Revenue in Five Years", 19 March 2026
https://www.bloomberg.com/news/articles/2026-03-19/alibaba-revenue-disappoints-as-ai-profit-push-grows-urgent

[31] Chen, Caiwei, " Three reasons why DeepSeek's new model matters", MIT Technology Review, 24 April 2026
https://www.technologyreview.com/2026/04/24/1136422/why-deepseeks-v4-matters/

# Energy: China's Expanding Edge

The AI revolution has transformed energy from a background input into a frontline strategic asset. Data centers—essential for training and deploying AI models—are extraordinarily power-hungry, and **the race to build AI infrastructure is now inseparable from the race to secure reliable, abundant electricity.** Depending on the region, electricity can account for as much as 40 percent of data-center operating costs.[32]

In the United States, data centers currently represent 4 to 8 percent of total power demand but are projected to drive two-thirds of future load growth.[33] McKinsey, BCG, and S&P Global expect U.S. power demand to grow 13–15 percent per year through 2030, driven primarily by AI and electrification. China, meanwhile, has been expanding generation capacity at a staggering pace: **since 2019, it has added over 2,500 TWh of new generation—more than twice Japan's entire annual electricity supply—compared with just 221 TWh in the U.S., roughly equivalent to what Florida uses in a year. In 2025, China's total electricity generation reached 10,707 TWh—more than double the United States's 4,670 TWh.** That divergence in raw energy expansion sets the stage for a fundamentally asymmetric competition at the base layer of the AI stack.

*Figure H.* **Great "Power" Competition**
*Total electricity generated in each country, measured in terawatt-hours*

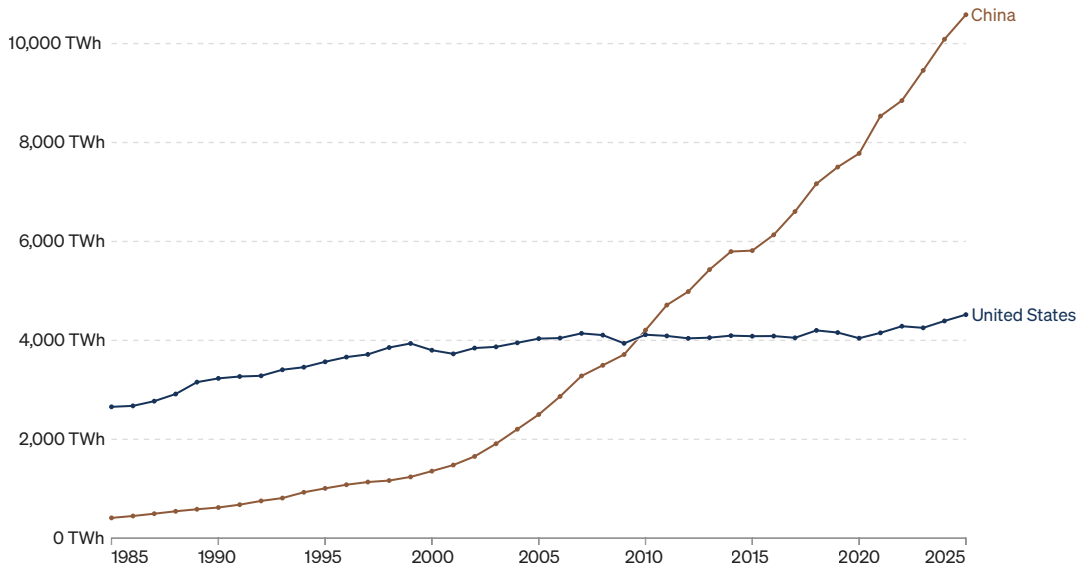


Source: Ember (2026); Energy Institute - Statistical Review of World Energy (2025)

**Price differentials widen the gap further.** In China's three largest data-center regions—Guangzhou, Shanghai, and Beijing—electricity costs average just $0.07–$0.09 per kWh. The National Development and Reform Commission (NDRC) sets benchmark prices and permits fluctuations within a ±20 percent band; when energy prices surge, state-owned utilities often absorb the cost.[34]

[32] U.S. Department of Energy, "Energy efficiency in data centers," https://www.energy.gov/cmei/femp/energy-efficiency-data-centers.
[33] Ibid.
[34] Zhang, Z. "Energy Price Reform in China," May 2018, Econstor. https://www.econstor.eu/bitstream/10419/191360/1/ndl2018-018.pdf

In the U.S., by contrast, prices vary widely across major data-center corridors:

- In Silicon Valley, rates can reach $0.15–$0.22 per kWh—nearly double China's.
- Other markets are more competitive: Dallas-Fort Worth, for instance, averages just $0.06–$0.08 per kWh.

But on a national basis, the disparity is stark. **China's average power rate reached US$0.08/kWh in the fourth quarter of 2025, more than 50 percent lower than the U.S. average of US$0.18/kWh.**

That gap is poised to widen further in the coming years. **China is expected to add more than 3.4 terawatts of generation capacity by 2030—nearly six times the U.S. total, according to BloombergNEF.**[35] Permitting backlogs, grid interconnection issues, workforce shortages, and supply chain bottlenecks all contribute to delays in bringing new U.S. power generation online.

In the long run, rising compute efficiency and improving energy performance of AI hardware could, in theory, temper overall power demand. But most experts believe those gains will be more than offset by surging aggregate AI usage—a dynamic in which cheaper energy only accelerates adoption.

## What We're Watching: Energy

- **U.S. energy politics:** Data centers and their perceived impacts on local energy costs and resources have become a hotly debated political topic in parts of the U.S. A recent Gallup survey found that seven in 10 Americans oppose constructing data centers for AI in their local area.[36] Will negative perceptions of data centers' energy impact and public pushback hamper U.S. efforts to increase generation capacity?
- **Efficiency:** How will the power efficiency of future AI models and hardware impact overall energy demand?

[35] Pike, L. et al. "China's Energy Buildout Is Its AI Superpower in Race With the US." February 5, 2026. Bloomberg, https://www.bloomberg.com/news/newsletters/2026-02-05/china-s-energy-buildout-is-its-ai-superpower-in-race-with-the-us.

[36] Jones, Jeffrey M. "Americans Oppose AI Data Centers in Their Area," Gallup, 13 May 2026, https://news.gallup.com/poll/709772/americans-oppose-data-centers-area.aspx.

# Section Four:

## Systemic Dimensions

## Financial: Capital, Risk, and the Question of Monetization

AI is commanding extraordinary amounts of investment, directly testing both countries' ability to marshal resources. We are watching venture and capital markets financing; concerns regarding systemic risk; and the extent to which AI is being monetized.

**The U.S. has outpaced China in the deployment of capital across the AI technology stack, from venture to mega-cap debt raises.** This is a reflection both of an absolute U.S. advantage (e.g., supply of capital) and the relative emphasis it is placing on costly frontier model development (e.g., demand for capital). The OECD estimates that the U.S. captured 75% of global AI venture capital investment in 2025 compared to 5% by China.[37] The U.S. also leads in capital expenditure: J.P. Morgan Research estimates U.S. firms spent $490 billion in 2025 compared to $98 billion by Chinese firms.[38] China competes asymmetrically through direct government investment and various forms of state support across the AI technology stack, but the scale of the U.S. private capital advantage remains a defining feature of the competition.[39]

**In AI, however, capital rarely travels alone. Investment brings with it access to intellectual property, research talent, strategic partnerships, and influence over where returns and capabilities ultimately flow. That makes AI financing inherently geopolitical, and both countries are increasingly moving to exercise control over it.**

- **The U.S. has long had a regime in place to screen inbound foreign investment in strategic sectors, including AI and other emerging technologies.** Over the last several years, Washington has also increased restrictions on outbound investment, limiting American investors ability to invest in Chinese AI, semiconductors, and quantum computing.

[37] OECD, "Venture capital investments in artificial intelligence through 2025,"17 February 2026, https://www.oecd.org/content/dam/oecd/en/publications/reports/2026/02/venture-capital-investments-in-artificial-intelligence-through-2025_3bcb227f/a13752f5-en.pdf.

[38] Lakos-Bujas, Dubravko, et al., "Global equity strategy: market update, AI & security resurgence, earnings accelerate amidst geopolitical uncertainty," J.P. Morgan, 21 April 2026; and Aziz, Jahangir, et al., "Ten questions about China in 2026," J.P. Morgan, 19 January 2026.

[39] See Sun, Luna and He Huifeng, "China launches venture capital fund to lead the charge on tech investment," South China Morning Post, 26 December 2025, https://www.scmp.com/economy/china-economy/article/3337798/china-launches-venture-capital-fund-lead-charge-tech-investment; Cao, Ann, "New AI fund in China to pour US$8 billion into early-stage projects," South China Morning Post, 11 April 2025, https://www.scmp.com/tech/policy/article/3306047/new-ai-fund-china-pour-us8-billion-early-stage-projects

- **China is now reportedly developing its own restrictions on inbound AI investment—and going further to prevent domestic talent and companies from leaving to access foreign capital.**[40] The case of Manus, a Chinese AI agent startup, illustrates the intensity of this dynamic. After the company's founders relocated from China to Singapore—in part to access the global marketplace—and Meta subsequently acquired the firm, Beijing imposed exit bans on the founders and has ordered the companies involved to unwind the transaction.[41] The intensity of the reaction is a demonstration of the competitive impact of America's financing advantage.

The deeper question is whether the current pace of investment is sustainable—and what happens if it is not. **AI capital expenditures continue to outpace revenues in both countries.** The costs of maintaining a competitive model has prompted some Chinese providers to shift away from open source models—a shift that could partially blunt the global adoption momentum described earlier in this report.

**Concerns about U.S. overinvestment in AI infrastructure have been joined in recent weeks,** first, by concerns about credit risks in the software-as-a-service sector,[42] and, now, about potential systemic risks if the accelerating capabilities of offensive AI are directed towards critical infrastructure.[43] So far, the U.S. financial system appears to be able to absorb current AI investment levels. J.P. Morgan Research estimates that tech sector bond issuance for 2026—a key measure of how much of the AI investment wave is being financed through corporate borrowing—will be 4% of total, rising to a "manageable" 10% by 2030.[44] But the trajectory matters: the more the AI boom is financed by debt, the more exposed the broader financial system becomes if returns disappoint. **The U.S. capital advantage is real and significant—but it is ultimately sustained by confidence that the investment will pay off.**

## What We're Watching: Financial

- **The Gulf's evolving role:** how does the Iran conflict and economic impacts across the Gulf impact one of the largest sources of investment in AI, which has generally favored U.S. firms?[45]
- **Signs of pullback or repricing in AI infrastructure investment:** Hyperscaler capex commitments have driven the current buildout, but those commitments are not irreversible. Guidance revisions, project delays, or shifts in data center contracting activity would be early signals that the investment thesis is softening—particularly if earnings calls begin to reflect shareholder pressure around near-term AI returns.

[40] Bloomberg News, "China to curb U.S. investment in tech companies after Meta deal," 24 April 2026, https://www.bloomberg.com/news/articles/2026-04-24/china-to-curb-us-investment-in-tech-companies-after-meta-deal.

[41] Tan, Rebecca and Lyric Li, "Beijing tightens its grip on AI firms that try to shed their Chinese ties," Washington Post, 21 April 2026, https://www.washingtonpost.com/world/2026/04/21/china-ai-competition-manus-meta/.

[42] Pitcher, Jack and Matt Wirz, "Private credit's exposure to ailing software industry is bigger than advertised," Wall Street Journal, 29 March 2026, https://www.wsj.com/finance/investing/private-credits-exposure-to-ailing-software-industry-is-bigger-than-advertised-d80da378.

[43] Sims, Tom, et al., "Banking industry scrambles for Anthropic's Mythos as global regulators review risks," Reuters, 20 April 2026, https://www.reuters.com/business/finance/banks-close-contact-with-european-regulator-anthropics-mythos-banker-says-2026-04-20/.

[44] Panigirtzoglou, Nikolaos, et al., "Flows & liquidity: AI financing looks more manageable following stronger cash flow growth," J.P. Morgan Research, 22 April 2026.

[45] Bhat, Divsha, "Can the Gulf buy its way to AI supremacy?," Rest of World, 24 June 2025, https://restofworld.org/2025/gulf-ai-investment-us-china-race/.

# Socioeconomic: The Productivity Promise and the Displacement Risk

AI's potential to transform economic productivity is enormous—but so is the potential disruption that comes with it. McKinsey & Company has estimated that currently demonstrated technologies could, in theory, "automate activities accounting for 57 percent of U.S. work hours today." [46] Concerns about socioeconomic disruption due to AI are growing in both the U.S. and China even as evidence of impacts on productivity and employment remain debated.

If these impacts materialize, **the competitive question will be whether each country's political and economic system can absorb the resulting dislocation fast enough to capture the gains without triggering a backlash that slows adoption.**

**In the U.S., tech firms are confronting bipartisan political discontent that ranges from concern over rising electricity bills caused by data center construction to fears of job loss.**

In response, they are actively seeking to shape debate on the reforms needed to deliver society into the AI era. OpenAI has proposed an "industrial policy for the intelligence age," which endorses a shift in tax burden from labor to capital, a public wealth fund, adaptive safety nets, and portable benefits, among other proposals. The company has also encouraged firms to experiment with four-day weeks.[47] Anthropic has been less specific in its social policy agenda, but has forecasted even greater labor market disruption. CEO Dario Amodei has repeatedly stated his belief that AI will wipeout up to 50% of entry level white collar jobs in the near term.[48] **That the architects of the technology are themselves proposing wealth redistribution to cushion its impact is a striking signal about what the industry expects is coming.**

> **"That the architects of the technology are themselves proposing wealth redistribution to cushion its impact is a striking signal about what the industry expects is coming."**

**China's approach has been markedly different. Beijing has discouraged open debate about AI's negative impacts, even as underlying pressures are no less acute.** Even prior to the AI boom, China was wrestling with elevated levels of youth unemployment, which officially stands at 16.9%,[49] compared to 8.5% during the same period in the U.S.[50]

[46] Yee, Lareina, et al., "Agents, robots, and us: skill partnerships in the age of AI," McKinsey Global Institute, 25 November 2025, https://www.mckinsey.com/mgi/our-research/agents-robots-and-us-skill-partnerships-in-the-age-of-ai.

[47] OpenAI, "Industrial policy for the intelligence age," 6 April 2026, https://cdn.openai.com/pdf/561e7512-253e-424b-9734-ef4098440601/Industrial%20Policy%20for%20the%20Intelligence%20Age.pdf.

[48] VanDeHei, Jim and Mike Allen, "Behind the Curtain: A white-collar bloodbath," Axios, 28 May 2025, https://www.axios.com/2025/05/28/ai-jobs-white-collar-unemployment-anthropic

[49] Reuters, "China's youth jobless rate rises to 16.9% in March" Reuters. 20 April 2026, https://www.reuters.com/world/asia-pacific/chinas-youth-jobless-rate-rises-169-march-2026-04-21/

[50] U.S. Bureau of Labor Statistics, Unemployment Rate - 16-24 Yrs. [LNS14024887], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/LNS14024887, May 5, 2026.

**The Chinese government has taken a more active role in mitigating the potential negative societal impacts of AI.** Two Chinese rulings have held that employees cannot be laid off because of AI.[51] Chinese AI labs have nonetheless echoed the U.S. sentiment that, as a DeepSeek spokesperson put it, "humans will be completely freed from the work in the end, which might sound good but will actually shake society to its core."[52] Even if China imposes a relative constraint on the upper bound of AI's productivity potential, even modest AI-driven productivity improvements to its considerable state-owned enterprise sector would have significant economic impact. China also continues to build out its social safety net amid considerable local fiscal pressure; officials and tech companies have recognized the need for reform to address AI's impacts.

**The extent to which each society is able to mitigate the negative impacts of AI may have important implications for adoption.**

**U.S. and Chinese concerns over AI are distinct.** Surveys indicate Americans are generally less enthusiastic about AI compared to people from other countries.[53] At the same time, China's political system places considerable emphasis on the ideological reliability of AI systems. Both types of concerns could weigh on AI progress; China's more centralized system, however, poses a greater risk of abrupt policy shifts than America's decentralized systems.

**Whether AI delivers greater benefits to services or manufacturing economies will also be key to the trajectory of AI in each country.** China's lead in robotics positions it to capture AI-related gains in manufacturing. Concerns in China often center on displacement of lower-wage workers, whereas U.S. concerns gravitate towards impacts on knowledge workers. The extent to which automation can deliver efficiencies in key sectors such as healthcare, which accounts for nearly a fifth of U.S. GDP, will impact the country's economic competitiveness and fiscal trajectory.

## What We're Watching: Socioeconomic

- **Degree of AI optimism:** Do the politics of possibility or fear over AI win out in each country?
- **Employment data:** When will productivity gains and/or clear employment impacts emerge in each country's employment data, and what are the political or societal impacts of the sectors and demographic groups where those impacts are concentrated?

[51] Swezey, Victor, "Chinese court rules firms can't lay off workers on AI grounds," Bloomberg, 2 May 2026, https://www.bloomberg.com/news/articles/2026-05-02/chinese-court-rules-firms-can-t-lay-off-workers-on-ai-grounds; G., Stephen, "China on AI job loss: 'no 'Matrix for us, thanks,' ChinaTalk, 1 April 2026, https://www.chinatalk.media/p/china-on-ai-job-loss-no-matrix-for.

[52] Chow, Vincent, "China's DeepSeek makes rare public comment, calls for AI 'whistle-blower' on job losses," South China Morning Post, 9 November 2025, https://www.scmp.com/tech/big-tech/article/3332086/chinas-deepseek-makes-rare-public-comment-calls-ai-whistle-blower-job-losses.

[53] Poushter, Jacob, Moira Fagan, and Manolo Corichi, "How People Around the World View AI," Pew Research Center, 15 October 2026, https://www.pewresearch.org/global/2025/10/15/concern-and-excitement-about-ai/.

# Military and Security: Norms Lagging Capabilities

AI is becoming embedded across military workflows, including cyber operations, targeting, weapon autonomy, and time-sensitive decision support. This has major implications for how governance, safety practices, and norms evolve alongside these new integrations to support stability. For example:

- **Steadily compounding military advantage:** Rather than a step-change weapon, AI integration is showing up in militaries through intelligence, surveillance, and reconnaissance (ISR), logistics, and decision support, creating gradually compounding shifts in operational effectiveness.
- **Cyberspace scaling:** AI-enabled cyber operations are scaling fast, increasing in speed and sophistication, vulnerability discovery, social engineering, and adaptive malware.
- **Ambiguous norms increase risk:** With limited sustained U.S.-China military dialogue and fragmented, non-binding norms, the pace of AI adoption is outstripping shared safety and crisis-management expectations.
- **Not your mother's arms control models:** Private sector innovation and cross-border diffusion complicate traditional arms-control models. Emerging agentic systems that can plan and act over longer horizons further intensify challenges around human control, accountability, and escalation thresholds.

**AI has not yet produced a decisive, observable shift in the military balance between major powers.** As of yet, there is no clear analogue to prior step-change technologies like nuclear weapons that fundamentally altered the calculus of great power conflicts. However, this risks understating what is already changing, which is that AI is meaningfully re-shaping the "kill chain"—how militaries gather information, decide, and operate—from intelligence, surveillance, and reconnaissance, to logistics optimization and decision support. These operational AI enhancements can yield cumulative advantages even absent a visible shift in battlefield dominance.[54] To wit, the result is a gradual, but compounding shift in operational dynamics, rather than a major inflection point.

**A major area of concern is that this transformation is unfolding alongside a governance gap.** Despite increased strategic stakes, military-to-military dialogue between the United States and China on AI risk remains very limited. The absence of more structured engagement could increase the risk of miscalculation, particularly in fast-moving crisis scenarios where AI-enabled systems may generate ambiguous or unexpected signals.[55] The history of the evolution of emerging technologies suggests that period of rapid capability development without parallel norm-building are especially prone to instability.

- **While there are emerging efforts to establish norms, such as the 2020 U.S. Department of Defense AI principles and ongoing United Nations discussions on autonomous weapons, these efforts remain fragmented and non-binding.** Key concepts like "meaningful human control" lack consistent definition, and states retain strong incentives to preserve operational flexibility. This fragmentation also reflects both the early stage of AI governance and the strategic competition underpinning it. The result is a widening gap between rapidly advancing capabilities and the norms intended to constrain them.

[54] Probasco, Emelia, "AI for military decision-making," CSET, April 2025, https://cset.georgetown.edu/publication/ai-for-military-decision-making/.

[55] Chivvis, Christopher S., and Jennifer Kavanagh, "How AI might affect decisionmaking in a national security crisis," Carnegie Endowment, 17 June 2024, https://carnegieendowment.org/research/2024/06/artificial-intelligence-national-security-crisis.

# AI's Impact on Cybersecurity

Cybersecurity is already the most active arena for AI-enabled competition, where capabilities are being deployed at scale. AI enhances the speed, sophistication, and reach of offensive cyber operations—automating vulnerability discovery, generating convincing social engineering content, and enabling adaptive malware that evolves in response to defenses. The offense-defense cycle is compressing, and AI is lowering the barrier to entry for both state and non-state actors.[56]

The competitive implications are asymmetric. Critical American infrastructure—energy grids, financial systems, healthcare networks—is managed by scores of entities of varying cyber capability and subject to limited central influence or support. In some cases, this can amplify vulnerability.[57] China's more centralized approach to infrastructure defense presents a different, though not necessarily smaller, set of vulnerabilities. If defensive AI does not keep pace with its offensive counterpart, the ability to absorb and recover from attacks—national cyber resilience—could become as important as the ability to launch them.

**Nearly three quarters of cybersecurity leaders are seeing the impact of AI-powered threats**

AI in its various forms can be exploited by attackers at every stage of the killchain—from initial intrusion, through to privilege escalation and data exfiltration. A survey conducted by Darktrace, a cybersecurity platform,  shows this is no longer science fiction.

With AI, attackers can launch novel / unknown attacks at scale, and traditional security solutions trained on historical attack data simply weren't designed to deal with this reality.

| | |
|---|---|
| 73% | say AI-powered cyber-threats are already having a significant impact on their organization |
| 92% | agree that AI-powered cyber-threats are forcing them to significantly upgrade their defenses |
| 87% | agree AI is significantly increasing the sophistication and success rate of malware |
| 87% | say AI is significantly increasing the number of threats requiring attention by the security operations team |

Darktrace: The State of AI in Cybersecurity, 2026
https://www.darktrace.com/resource/the-state-of-ai-cybersecurity-2026-threat-landscape

Note: These figures are based on an October-November 2025 survey of 1,549 cybersecurity leaders and practitioners from 14 different countries

[56] Metz, Cade and Kate Conger, "A.I. is on its way to upending cybersecurity," New York Times, 6 April 2026, https://www.nytimes.com/2026/04/06/technology/ai-cybersecurity-hackers.html.

[57] Neuberger, Anne, "China is winning the cyberwar," Foreign Affairs, September/October 2025, https://www.foreignaffairs.com/china/china-winning-cyberwar-artificial-intelligence.

## Decision Speed, Attribution, and the Risk of Miscalculation

**Beyond cyber, AI is reshaping the conditions under which the U.S. and China would manage a crisis between them.** AI systems can synthesize vast amounts of information and generate recommendations under significant time pressure, compressing the window available for human judgement. In high-stake scenarios, the risk is that decision-makers on either side over-rely on AI outputs when time for deliberation is limited—a dynamic sometimes described as "flash war," in which the speed of AI-generated analysis outpaces the ability of leaders to verify, contextualize, and decide.

**Compounding this is a growing ambiguity problem. AI-enabled operations, particularly in the cyber domain, can further obscure both attribution and intent.** Synthetic media, automated attack infrastructure, and adaptive systems make it more difficult to determine whether an action was deliberate, accidental, or an emergent behavior from the system itself. In a relationship already characterized by strategic distrust, this ambiguity increases the likelihood of misinterpretation and overreaction.

## The Dual-Use Challenge

**These challenges are intensified by the fundamentally dual-use nature of AI.** Unlike nuclear technologies, AI development is largely driven by the private sector, with commercial innovation flowing rapidly into military applications. Traditional arms control approaches, which rely on clear distinctions between civilian and military technologies, are poorly suited to a world in which the same model can power a customer service chatbot and a military targeting system. Companies such as OpenAI and Anthropic are central to frontier model development for everyday civilian use, yet their technologies also have clear national security implications. **This dual-use dynamic complicates governance, as capabilities diffuse rapidly across sectors and borders.**

- Anthropic's current position illustrates the resulting tensions. Anthropic is simultaneously the developer of Mythos—currently the most strategically consequential AI model to date—and currently designated as a supply-chain risk to U.S. national security. That a single firm can be both so important to and ostracized from the national security apparatus captures the governance contradictions that dual-use AI creates, and that existing institutional frameworks have yet to resolve.

**Looking ahead, the emergence of agentic AI systems introduces an additional layer of complexity.** Unlike traditional models, these systems can plan, use tools, and operate over extended time horizons. As such, these systems could enable continuous, adaptive operations in both the cyber and physical domains.[58] **In a military context, this raises new questions about control, accountability, and escalation thresholds, particularly if systems act in ways that are difficult for human operators to fully anticipate.**

**The broader implication is that AI is reshaping the conditions of military competition without yet redefining its outcomes.** The most immediate risks arise not from a single breakthrough capability, but from the interaction of accelerating adoption, compressed decision timelines, and weak or fragmented norms. Stability in this domain will depend less on technological leadership alone and more on whether major powers can establish shared expectations for how AI systems are developed, deployed, and governed. Until such frameworks mature, **the gap between capability and governance will remain a central source of strategic risk.**

[58] https://www.computer.org/csdl/magazine/co/2026/01/11321042/2cTQGI6RjO0

## What We're Watching: Military and Security

- **Space competition:** China and U.S. firms are both actively exploring the potential for space-based data centers as the economics of launch continue to improve, potentially injecting AI into yet another domain of competition between the two countries.[59]
- **Anthropic v. Pentagon:** The U.S. government's announcement earlier this year that it would ban use of Anthropic's technologies for Pentagon contracts and designate the company as a supply chain risk to national security raised significant questions in the U.S. about the relationship between AI developers and the defense ecosystem. How will ongoing legal proceedings play out, and how will those results impact the balance of power in the U.S. defense technology ecosystem?
- **The blurring of AI's dual-use dynamics:** The scale of model deployment and dual-use dynamics will re-shape the threat landscape. Advances in technology will lower the cost and skill required to deploy AI, including in robotic systems, enabling wider proliferation across both state and non-state actors. Civilian platforms can be rapidly re-purposed, making governance and export controls harder to define and enforce.

[59] "Data centers in space: less crazy than you think," The Economist, 2 March 2026, https://www.economist.com/science-and-technology/2026/03/02/data-centres-in-space-less-crazy-than-you-think; Chen, Laurie, "China plans space-based AI data centers, challenging Musk's SpaceX ambitions," Reuters, 29 January 2026, https://www.reuters.com/science/china-vows-develop-space-tourism-explore-deep-space-it-races-us-2026-01-29/.

# Conclusion

**The U.S. and China unmistakably seek to achieve and maintain leadership in AI, seeing it as central to their national security and economic competitiveness.** The competition between the two countries will continue to evolve as each country registers technological advances and makes adjustments to its respective systems.

**The U.S. has clear and sustainable leads in frontier models, semiconductor design, and AI financing; China's advantage is most pronounced in energy. China also brings considerable counterweights: in its support for open-source models, its influence, through geographic proximity, over a semiconductor supply chain concentrated in Taiwan, and the extraordinary talent of Chinese origin competing on both sides of the Pacific.**

**What is most unclear at this stage is the political economy that will decide which country most embraces AI's potential, while maximizing societal gain and minimizing harmful disruption.** Both countries' militaries face considerable challenges in scaling and embedding AI in their operations; the United States' advantage in combat experience is increasingly offset by China's control over key supply chains from critical minerals to drones.

Both countries will continue to deploy a mix of tactics to outcompete and obstruct each other and get out of their own way. The U.S. maintains leads in frontier AI model development and capital deployment, and the political pendulum may again be swinging back towards stricter semiconductor export controls. In energy, the U.S. does not need to lead—just avoid self-constraint. China will double down on its strengths in physical AI, keep its medium-term chokepoint on critical minerals at the ready, and aggressively seek to work around limits on its access to compute.

**Even as the U.S. and China compete, other countries retain agency and a degree of influence.** Their markets, talent pools, resources, and policy environments are all opportunities to capture AI's potential benefits. The technological alignments their markets make will both reflect and shape relative U.S. and Chinese leadership, not only in AI, but in the broader sweep of their political, security, and economic ties.

For companies, which are racing to leverage AI within their own organizations and navigating industry-wide disruptions, the geopolitical AI contest may seem to be of secondary importance. **But it should not be ignored: the U.S. and China's relative standing in AI will influence the geography of economic growth and the balance of power.**

**Looking ahead, the base case for businesses is one of continued progress and fragmentation** *(see Figure I)***.** In this scenario, the global technology landscape continues to fragment due to policy-driven access to hardware and models. The returns from AI justify continued investment and responses to AI's socioeconomic impacts stay ahead of disruption. But there are upside and downside cases too.

- **In the upside case,** companies benefit from a policy competition oriented towards innovation and adoption instead of denial and control. AI's gains are broad-based and sustainable.
- **In the downside case,** security eclipses economics. AI progress slows and access to hardware and software becomes more uneven. Financing pulls back even as AI's disruptive impacts are met without sufficient socioeconomic reform.

The decisions of governments, AI innovators, adopters, and civil society will all influence the balance between these possible outcomes. This framework can provide a language that moves discussion of the AI competition beyond the technological or geopolitical lens and more fully reflect its societal importance.

*Figure I.* **Scenarios of U.S.-China AI competition and implications for global businesses**

| | Potential scenarios | | |
|---|---|---|---|
| **Today** | **Base case** | **Upside case** | **Downside case** |
| **Racing towards uncertainty** | **Progress and fragmentation** | **Prosperity delivers its own stability** | **Security transcends economics** |
| • Policy tensions between innovation and commercial leadership and security<br>• Hardware subject to controls and capacity constraints depending on geography<br>• Softening innovation owing to fragmentation with uneven access to advanced chips<br>• AI financing both driving growth and potentially crowding out other investment<br>• Recognition of the socioeconomic implications of AI, but competing for course correction<br>• Military AI adoption and cyber-biosecurity risks accelerating | • Global technology landscape continues to fragment due to policy differences and hardware access<br>• Energy proves a limited practical constraint as firms can access lower-cost energy in energy-abundant geographies<br>• Gap between technological frontier and real-world deployment persists, but narrows slightly<br>• Productive debates on socioeconomic response to AI stays ahead of blowback from disruption<br>• Self-interest guides responsible military and security use of AI even in the absence of robust norms | • Policy continues to support innovation<br>• Geopolitical emphasis shifts from hardware and software leadership and control to support for private sector adoption<br>• Policy and private sector reduce energy as a bottleneck<br>• AI delivers broad-based financial gains<br>• AI drives socioeconomic reforms, provides a foundation for sustainable continued growth<br>• Military and security norms and governance ensure AI is not geopolitically destabilizing | • Policy and regulation slows AI innovation<br>• Uneven access to hardware and software capabilities or energy-driven constraints<br>• Gap between investment requirements and revenues prompts financing pullback, while second-order impacts create a drag on growth<br>• Adoption increases and drives displacement without stabilizing reforms to socioeconomic governance<br>• Emphasis shifts from AI's economic potential to military and security priorities |

# About the JPMorganChase Center for Geopolitics

The JPMorganChase Center for Geopolitics harnesses the firm's vast global network of expertise and know-how to help clients successfully seize opportunities and weather the trends transforming our global landscape. Access to the Center's offerings can be facilitated through bankers or other client advisors. For more information or to contact the JPMorganChase Center for Geopolitics, please visit: www.jpmorganchase.com/geopolitics

# About the Authors

**Deborah Eng**
Executive Director, AI/Data Policy, JPMorganChase

**Kyle Hutzler**
Executive Director, JPMorganChase Center for Geopolitics

**Todd O'Boyle**
Executive Director, AI/Data Policy, JPMorganChase

**Nathaniel Sher**
Vice President, JPMorganChase Center for Geopolitics

**Emily Sullivan**
Senior Associate, JPMorganChase Center for Geopolitics

The authors appreciate the feedback of colleagues from across the firm.

# Disclaimers