

# AI Governance as Hegemonic Statecraft: The U.S.–China Contest to Write the Rules of the AI Age

The Economy Research Editorial<sup>1,2</sup>

<sup>1</sup>The Economy Research, 71 Lower Baggot Street, Dublin 2, Co. Dublin, D02 P593, Ireland

<sup>2</sup>Swiss Institute of Artificial Intelligence, Chaltenbodenstrasse 26, 8834 Schindellegi, Schwyz, Switzerland

## Abstract

AI governance is often posed as a narrowly technical or moral question, about how to secure the benefits of powerful systems and steer them away from harms: how to design systems that are safe, fair, transparent and useful. The article maintains that framing is too limited for the geopolitical moment now opening. Using the U.S.-China rivalry as the anchor case, the article demonstrates that from 2023 to 2026, artificial intelligence ascended from a rapidly expanding tech to a key region of strategic competition on the global stage. Rules, norms, infrastructure and institutional authority over global AI development and deployment became increasingly important battlegrounds for control over the distribution of power. The article traces the U.S. and Chinese trajectories through shifting positions over AI governance. To the U.S., the playing field of AI governance has homogenized into a strategy of technological supremacy, export leverage and alliance-centric standards proliferation. To China, it has morphed from a participant in a predominantly Western digital order into an arena of active contest over what kind of AI governance to uphold and how, in U.N.-based diplomacy, standards pursuits and global capacity development. Overall, the movement represents neither a clean transfer of global AI leadership from Washington to Beijing nor a stable balancing of two Great Powers, but a splintered fight to set the rules for the AI age. The paper argues that effective global AI governance will not emerge until states marshal institutions to manage strategic competition without a slide into unbound power politics.

# 1 Introduction - AI Governance as a Struggle for Rule-Setting Authority

The dominant conceptions of governance in the developing field of AI governance still seem too thin for the moment we are witnessing. They tend to assume that governance is an ethical add-on to innovation: an additional layer of safety barriers, voluntary principles, or regulatory practices applied after technological development rather than shaping its preconditions. Such an image seems increasingly inaccurate. In international politics, rule-setting has seldom been neutral technical administration. At least since the British hegemony of the nineteenth century, embodied through the Pax Britannica and certainly since the American-led post-1945 international order, scholars of international order have long argued that states able to shape institutions, standards and norms do more than constrain behavior; they identify what counts as legitimate, acceptable practice in the first place.<sup>[1]</sup> That in the era of AI is all the more significant. Rules of control on the use of computing power, access to massive volumes of data, methods for evaluating model performance, controls on exporting those models, security screening practices, public procurement standards and expectations of interoperability do far more than govern markets- rather they also help define which actors - states, firms, regions and innovators - will be able to build, to buy, to grow and on what terms.

That is why the period between 2023 and 2026 was so important. In that brief span, AI advanced from a rapidly expanding industry to the defining parameter of strategic rivalry. According to the OECD, AI firms generated 61% of the aggregate value of global venture-capital investment by 2025, worth approximately \$258.7 billion, more than twice AI's share in 2022.<sup>[2]</sup> Concurrently, according to the Stanford AI Index, while the United States remains the scale leader in private AI investment and cutting-edge model output, the frontier model performance between the United States and China has been effectively eliminated,<sup>[3]</sup> as of March 2026, Anthropic's leading model was ahead of the top Chinese model by only 2.7 percentage points on the relevant benchmark. The Index also reports that the total number of models being produced continues to be dominated by these two nations, even as other countries develop detailed AI strategies.<sup>[4]</sup> These developments mean that AI governance is no longer a theoretical discussion of the public benefit of emerging technologies, but a struggle over the makings of international power at a time when capital, compute and military relevance are all transforming at the same pace.

The strongest evidence of this shift was not just technological, but institutional. In 2024, the U.S. continued to advance the first U.N. General Assembly resolution on AI <sup>[5]</sup> and signed the Council of Europe's Framework Convention on AI, the first legally binding international treaty in the field.<sup>[6]</sup> By 2025 and early 2026, however, Washington had pivoted away from multilateral rule-making language. The U.S. and UK refused to sign off on the Paris AI Action Summit declaration in February 2025<sup>[7]</sup> and the Trump administration's 2025 AI Action Plan had already brought trade, standards and diplomacy under a strategy specifically designed to promote American hardware, models and standards around the world, while countering Chinese influence in international organisations,<sup>[8]</sup> by 2026, in New Delhi, America's chief White House science adviser, Michael Kratsios, publicly declared that the Trump administration totally rejects global governance of AI.<sup>[9]</sup> China was headed in the other direction: from an agenda of disseminating indigenous use cases and infrastructure toward a global institutional

politics agenda via the UN, the World AI Conference in Shanghai, standards work and a growing campaign of AI capacity-building diplomacy for the Global South.

The core thesis of this paper is that the current AI governance debate is best read as a contest over political authority in a new technological epoch, rather than as a contest over the appropriate level of regulation. The United States has not abdicated governance; it has reconceptualized governance as a matter of competitive advantage, bilateral leverage, export policy and alliance management rather than a universalist public-interest framework. China, on the other hand, has gone from largely functioning within an order shaped elsewhere to scripting the language, institutions and developmental doctrine of AI governance itself. This has not been a straightforward transition from one hegemon to another. The material balance still favors the United States in frontier capital, data-center scale and top-end model production. But China no longer seems happy to be a participant in a market ordering whose core rules were designed elsewhere. It is now challenging the authority to craft the system. This effort is particularly noteworthy because it happens at the very moment when Washington has diminished its appetite for supporting the universalist institutions through which it once wielded agenda-setting power.

This framing matters analytically because the debate tends to be mischaracterized in one of two unhelpful ways. On the one hand is the naïve premise that rule-setting will be a natural corollary of technological dominance, because governance is a downstream consequence of discovery. On the other hand is the naïve premise that the world can arrive at a neutral consensus independent of great-power rivalry, as if AI's stakes aren't tightly bound up in competing national interests, from employment and industrial competitiveness to military capability. Both are unwarranted. Rule-setting advantage isn't just inherited from technical know-how; it must be organized through institutional presence, diplomatic investments, sustained through participation and alliance-building. And governance isn't just guaranteed to be immune from geopolitics, because AI today is at the nexus of gains in productivity, economic position, weapons and information control. The upshot isn't the death of governance, but the transformation of governance into a field of hegemony. The next three sections analyze that transformation through three interconnected trends: American retrenchment from multidimensional projects in cyber and AI rule-setting; China's strategic elevation from rule-adopter to rule-maker; and the structural tensions that render this zero-sum contest more intractable than either side's narrative suggests.

## **2 The U.S. Retreat From Multilateral AI Governance**

This becomes most apparent when reading Washington's path through history, not isolated events. The retreat of 2025–26 did not come from a starting point of disengagement. Leading up to 2024, the U.S. continued to push ambitiously for rulemaking in a variety of contexts. It partnered with other nations in the first-ever U.N. General Assembly resolution on AI, adopted by consensus in March 2024, which supported principles of safe, secure and trustworthy artificial intelligence and called for inclusive development with developing countries; in September 2024 it signed the Council of Europe's Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law, which the Council of Europe widely proclaims as the first treaty of its kind. These were not merely symbolic gestures; although they were important legal milestones, they were also

the kinds of norm-setting milestones strategists seek: establish norms early, actively participate in emerging institutions and preserve leadership by being where the action is.




Figure 1: Selected milestones show the U.S. shift from multilateral participation toward retrenchment, while China concentrates activity in rule-setting and capacity-building diplomacy.

The importance of this reversal rests in the contrast it highlights. The Trump Administration didn’t simply undermine enthusiasm for AI regulation at home. Instead, it changed the venue where governance would take place. Executive Order 14179, signed in January 2025, rescinded the Biden Administration’s prior Executive Order on AI and stated that U.S. policy would be to sustain and improve America’s global AI preeminence.<sup>[10]</sup> By July 2025, the White House’s AI Action Plan had generalized this notion entirely. Its third pillar, Lead in International AI Diplomacy and Security, argued that the U.S. needed to drive the adoption of American AI systems, chips, software, applications and standards around the world.<sup>[11]</sup> It explicitly proposed exporting the entire U.S. AI stack to countries adopting America’s AI alliance, countering Chinese leverage in international institutions like the United Nations, the OECD, the G7, the G20 and the ITU. In other words, this American position shifted from helping construct widely shared rules to one that sought to use technological and market dominance to indirectly influence them.

The diplomatic tone employed by the administration dispelled any lingering uncertainty. At the Paris AI Action Summit in February 2025, the U.S. refused to co-sponsor the summit declaration on inclusive and sustainable AI,<sup>[12]</sup> which sixty nations, including China, France, India and the majority of Europe, signed on to. In February 2026, Michael Kratsios proclaimed to the India AI Impact Summit that the U.S. totally rejects global AI governance and is against top-down control by transnational bureaucracies. The same address redefined AI sovereignty as something to be achieved not through universal rule-making, but by embracing an AI tech stack that enables the fastest AI deployment and then encouraging a global AI supply chain, ready to send the world’s best AI systems as U.S. offerings. This is not anti-governance in any literal sense. It is anti-

universal governance. The U.S. continues to seek influence over global AI outcomes, but through the wielding of market power, selective bilateral arrangements, AI export programs and technology-stack partnerships with selected foreign governments.

To write it off as a lack of understanding would be intellectually lazy. There is a clear and comprehensible strategic rationale for that retrenchment.<sup>[13]</sup> The White House contends that a broad blanket of regulation could have a negative impact on innovation, cement the position of foreign rivals and increase the financial and operational strains on American companies at a moment when their competitive advantages in frontier capabilities are most critically relevant in both the commercial and military worlds. This fear cannot be discounted. Jones Walker's examination of regulatory trends in AI, as well as other documents, points to a broader transatlantic retrenchment trend, consisting of delays to key provisions of the EU AI Act, the collapse of the federally proposed Canadian AI bill and rollback or efforts in U.S. State legislation. It is not simply an argument grounded in deregulation ideology. It is the realization that first-generation, comprehensive AI frameworks were often devised prior to the requisite compliance infrastructure to make their mandates a reality and that governments grew more and more wary of burdening firms at the leading edge of the technology with the operational and financial weight of compliance, when they were being challenged by somewhat less constrained American and Chinese rivals. The U.S. approach, therefore, is largely an assumption that reduction of multilateral and domestic regulatory frictions will help sustain innovation momentum and leverage the country's long-standing lead in capital, hyperscale infrastructure and frontier model development.

There is reason to believe that the United States currently has the material basis to support this gamble. According to the Stanford 2026 AI Index, private investments in U.S. AI totaled 285.9 billion dollars in 2025, more than twenty-three times the 12.4 billion dollars recorded for China<sup>[14]</sup> in that category and U.S. institutions produced fifty-nine notable AI models in 2025, versus thirty-five for China. The same report indicates that the United States has 5,427 data centers, more than ten times as many as any other country<sup>[15]</sup> and that it remains the world's most desirable location for frontier-model development, even though China has significantly caught up to it in that respect. Therefore, an approach to governance based upon the propagation of standards achieved through technical ascendancy rather than treaty-making cannot be said to be without foundation. It hinges on the notion that those in possession of the most refined models, most trusted chips and most coveted software can impose de facto rather than legislated norms.

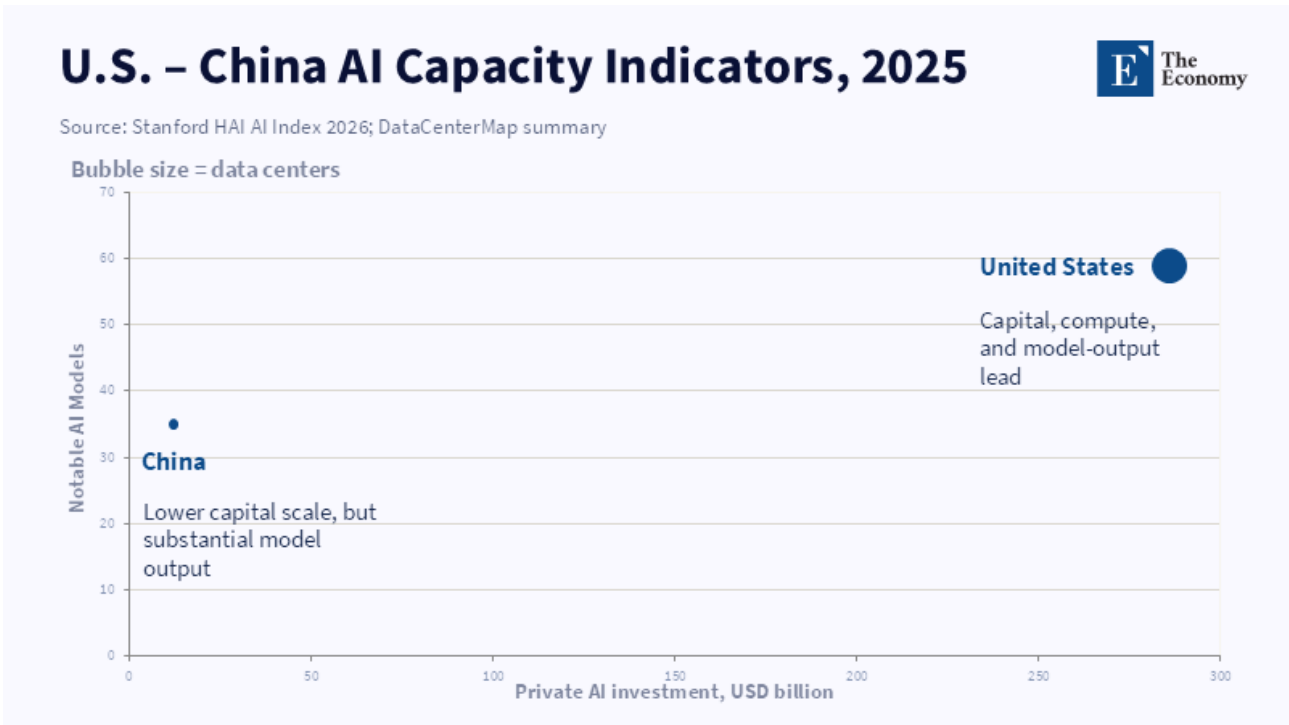


Figure 2: The U.S. retains a large material lead, but China’s model output keeps its governance bid credible.

However, this strategy has a cost, one that is often overlooked in U.S. policy debates. Hegemony is not just the ability to command market share; it is the ability to craft agendas in institutions that others consider legitimate. The Atlantic Council has criticized the Trump administration’s focus on competition with China through international governance bodies as ineffective when the nation’s diplomatic machinery is being eviscerated through publicized staff cuts and internal administration dysfunction.<sup>[16]</sup> The Council on Foreign Relations echoes these concerns, reporting that as of March 2026, the U.S. State Department’s Bureau of Cyberspace and Digital Policy was abolished, its offices vanished, with close to a dozen staffers dismissed,<sup>[17]</sup> the ambassador-at-large position vacant and CISA’s international affairs offices also cut. These trimming efforts matter because international rule-making does not just occur at summits; it depends on diligent, often invisible diplomacy in formal standards bodies, development partnerships and intergovernmental negotiations. When such capacity is diminished, so too is standards-setting momentum.




Figure 3: U.S. capital dominance gives Washington a market-led route to influence, even as its multilateral governance capacity weakens.

The same logic is in evidence in U.S. fiscal signals. The White House’s FY 2027 budget proposed a 30 percent cut to State and other international programs and a \$2.7 billion cut to international organizations and the UN,<sup>[18]</sup> while shutting down funding and privileging only those international institutions that were perceived to be directly advancing American interests. Since this was a budget proposal (rather than a bill), it must not be overstated - indeed, as a policy signal, it speaks volumes. Washington is withdrawing investment from the universalist infrastructure through which liberal rule-setting of the post-Cold War era has always functioned, even while it claims policy primacy in AI. What emerges is a much tighter, more durable, more scalable and transactional vision of order: governance not as a shared framework with which elites must socialize with others, but as an extension of export power and strategic bargaining.

The most profound consequence of this change is thus conceptual. The United States is not abandoning the international arena. It is abandoning a particular kind of international arena. Its retreat is from multilateral cyber and AI governance as a public good, an act of systemic leadership, as opposed to a restriction of the technological focus. Recognizing this distinction explains how the retreat might seem compelling at home but fatally distorted abroad. The short-term advantages may be lower compliance costs for U.S. companies through reduced regulation and a full-stack export strategy; the long-term risk is that the U.S. leadership project weakens the very political legitimacy that made American hegemony resilient, namely, framing self-centered power as a formalized regime of civility. When that legitimating narrative becomes less dense, it becomes not the material abundance of the United States but the members’ relatively high-degree presence outside it that turns them into credible rule entrepreneurs.

### 3 China's Strategic Repositioning From Rule-Taker to Rule-Maker

That is precisely in that opening China has now moved. The importance of the recent Chinese behavior does not lie in the fact that the country has suddenly discovered the governance of artificial intelligence, but in the change of ambition. For a large part of the digital age, China's external strategy was driven by asymmetry. It established a highly unique internet regime at home, while operating within a global order whose profound institutional principles had already been set by others. Carnegie's recent assessment captures this new inflection point in a few words: Beijing's current artificial intelligence diplomacy is shifting away from a focus on exporting infrastructure and related technical standards and is now seeking to overhaul international norms and institutions for Artificial Intelligence governance.<sup>[19]</sup> The change of course is significant because it indicates a move from a position of playing along to one of shaping. China is no longer simply seeking to make sure its own corporate and digital institutions can thrive in an already existing order. It is trying to shape the conceptual terms through which the order itself will be discussed.

The chronology of that shift is now hard to ignore. In October 2023, China's Ministry of Foreign Affairs launched the Global AI Governance Initiative,<sup>[20]</sup> demanding a bigger say for developing nations, upholding the equality between nations in the governance of AI, and advocating for an international institution based at the U.N., to steer AI development, safety and regulation. Then in September 2024, China's government announced an AI Capacity-Building Action Plan for Good and for All,<sup>[21]</sup> directly tying AI diplomacy to the global digital divide, the progression of the U.N. Sustainable Development Agenda, and support for the Global South. The document outlined infrastructure cooperation, open source communities, training exchanges, joint educational centers and ten workshops or seminars focused primarily on developing nations through the end of 2025. At the 2024 World AI Conference, China released a Shanghai Declaration on Global AI Governance.<sup>[22]</sup> A year later, Premier Li Qiang outlined an international AI partnership organization and released a Global AI Governance Action Plan at the 2025 World AI Conference.<sup>[23]</sup> These steps reveal a state institution now eager to build a multi-layered governance platform.

Three aspects warrant close examination. First, the United Nations-centered institutional strategy. Beijing consistently refers to the United Nations as the main channel for global AI governance, to appeal to states concerned that it will be mainly managed by the private sector or Western powers and give them access to preferred venues to rally developing-country support. The U.N.'s own Global Dialogue on AI Governance, established by General Assembly Resolution A/RES/79/325 and due to hold its inaugural session in Geneva on July 6-7, 2026, is promoted as a venue where all countries and all stakeholders can participate in the AI governance discussion.<sup>[24]</sup> For Beijing, venues like these are attractive not because they are guaranteed to succeed, but because they offset the U.S. broad advantage in the private-sector-dominated and alliance-centric venues. However, they fit Beijing's theme that the AI governance agenda should be inclusive of all countries' interests and not just those of the technologically most advanced.

The second aspect is developmental statecraft. Beijing's AI governance diplomacy is not merely rhetorical. It is linked to programmatic offers of infrastructure, training, participating in standards and technical inclusion. The AI Capacity-Building Action Plan says explicitly, in an effort to bring digital and AI infrastructure

connectivity and industrial application in AI to the Global South,<sup>[25]</sup> promote AI literacy and workforce training, develop and share data-governance mechanisms and interoperable risk frameworks and facilitate practical cooperation between different regions and countries through laboratories and workshops. This recommendation for development is exemplified by the Atlantic Council's dissection of Beijing's offering of no-holds-barred technology, project financing, training programs and the more enduring Belt and Road digitization diplomacy. This offering resonates with many in the Global South who are less concerned with the subtle political philosophy behind stakeholder democracy than with guaranteed access to computing power, affordable applications, affordable localization and avoiding technological marginalization. In this regard, Chinese AI governance diplomacy is not merely normative persuasion; it is also distributive and offering. It is, in that regard, an offering of participation instead of lecturing. The third aspect is standards politics. Beijing seems to have thought that agreement on formal institutions alone would not suffice without the technical avenues through which norms may coalesce into practice. An online consultation held in July 2024 by the Chinese government stated that the country would develop more than fifty national and industrial AI standards by 2026,<sup>[26]</sup> join twenty-plus international standards and get over one thousand institutions willing to utilize and further these standards. Similarly, the UK-based news agency Reuters announced at the end of December 2024 that there's a new forty-member AI standardization technical commission and the commission is developing standards that will incorporate LLMs and risk assessments. Why all this emphasis matters is that standards and specifications serve as the most intangible but most durable forms of rule-making. They guide procurement decisions, composition expectations, certification procedures and interface specifications. If the international governance of AI on the whole is likely to stay weak on a law-and-enforcement model, standards and administrative guideposts will be extremely more central in the architecture of influence. Making the push in this area indicates a strategic awareness of the fact that AI mastery will be governed not only through authority of declarations but through sets of technical defaults.



Figure 4: China’s rule-setting strategy relies on standards, technical uptake, and institutional diffusion rather than declarations alone.

This effort would be less important if China were at the periphery of science and technology. It is not. According to Stanford University’s 2026 AI Index, China is the world leader in the number of publications, the number of citations, patents applied for and industrial robot installations<sup>[27]</sup> and while the United States still has the lead in frontier model output and private investment, in 2026, it was second in those categories as well. China’s official statistics claim that the country’s core artificial intelligence industry was valued at 1.2 trillion yuan in 2025 and that the number of Chinese AI firms exceeded 6,200.<sup>[28]</sup> This does not establish Chinese dominance; it does establish that China has built enough technological capacity that its projects for AI governance cannot reasonably be dismissed as just talk. Efforts to set the rules that have actual industrial capacity behind them are completely different from norm entrepreneurs without comparable industrial capacity to produce or make use of frontier systems of their own.

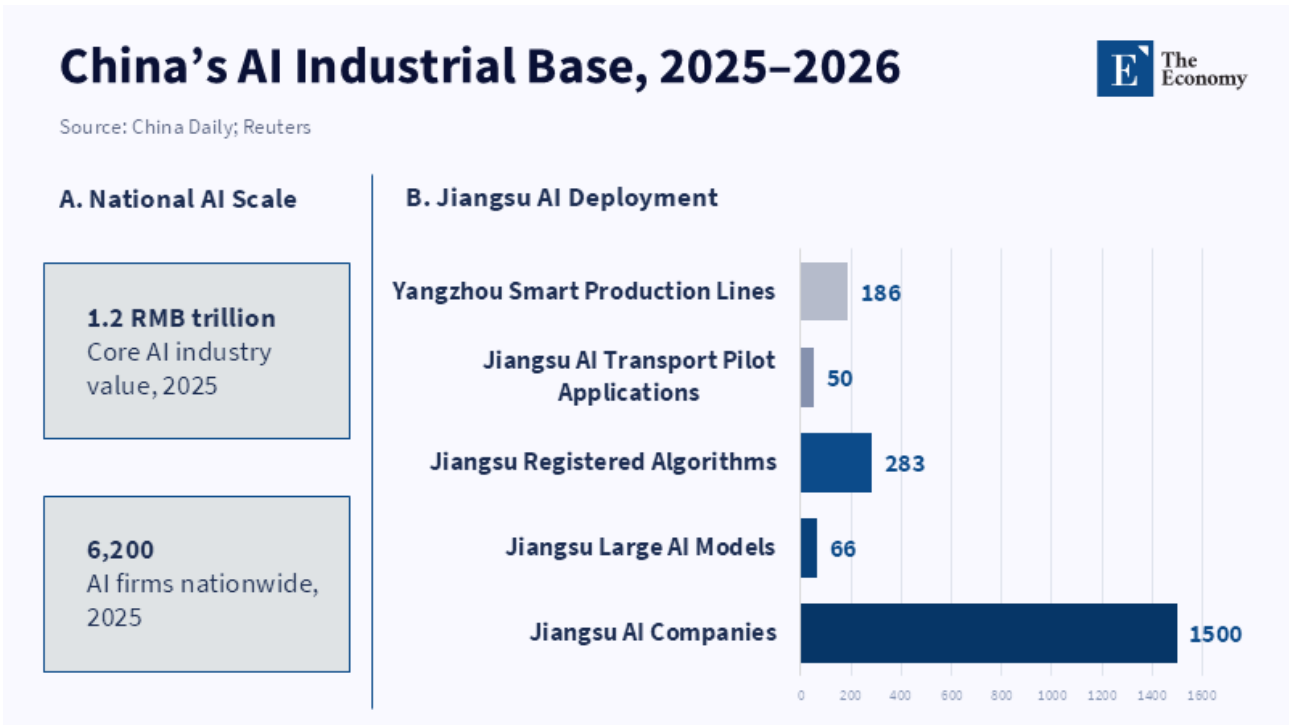


Figure 5: China's governance activism is backed by a large national industrial base and visible provincial deployment capacity.

Even so, to call China the leader on AI governance in advance would be folly. And here the strongest counterargument is also the most salient. The U.S. still maintains a significant edge when it comes to private capital, hyperscale data centers, the most advanced semiconductors and frontier start-ups; AI Index's investment data is fairly clear and U.S. companies still dominate with the most publicly prominent models. In addition, by itself, China's governance agenda is politically contested: China's governance ambitions cannot be explained in the abstract without reference to the specific actions and need to secure the definition of Chinese throughout large swathes of its own society. Carnegie's review of China's shifting perspectives on AI safety as represented through politics reveals that their safety interest is disparate from a naively defined safety and disproportionately integrated with Chinese content suppression, national security strategy and the desire to avoid second-mover status to Western powers. Atlantic Council experts draw a similar conclusion, suggesting that Beijing's close ties between a Sino-centric vision of AI governance, centralized state authority and global institutional change will be exceedingly difficult for many individual countries, corporate players and civil society groups (in liberal democracies, at least) to reconcile with their interests, regardless of China projecting itself as a champion for international cooperation.

But the only sign of China's resilience, its potential against the global skepticism, is not blind optimism. By no means is the mastery of governing the same as the moral authority, nor is it more than the frontier model count. Rather, it requires institutional entrepreneurship, patience and the skill to transmute home capacity into foreign assets. In such a sense, China has become more institutionally assertive. It has spun its AI governance priority to paths of development, sovereignty, anti-hegemony rhetoric and multilateral justice; it has registered them not only at international institutions but also at regional and sub-multilateral forums. For instance, in the 2024 Beijing Action Plan of the Forum on China-Africa Cooperation, it has alluded to China's Global AI

Governance Initiative and pledged both to buttress novel co-production in the realm of AI capacity-enhancement and global cross-border data trade, privacy and digital management rules.<sup>[29]</sup> Such documents do not prove full alignment with China, but they show coalition-building in practice, one issue at a time, one country at a time, and one standard at a time.

That's why this business of from follower to leader has to be nuanced in its application. China has not yet displaced the U.S. as the axis of the global AI ecosystem-perhaps it won't in some of the fields anytime soon-and what has changed is both more understated yet, arguably, more significant. Beijing now operates as though its own rule-setting is a form of geopolitical power worth competing for directly. That is the true turning point. To sell products and services, as the underlying argument suggests, confines a nation to work within someone else's rules. To design standards, host forums, fund capacity enhancement and make a language of sovereignty and progress quotidian are the steps that extend a state's influence to be a part-author of the system. The goal is no longer one of profit alone. It is to be authoritative. And inasmuch as Washington has itself limited its own conceptual range of governance to an industrial policy, China's move to claim the wider value-adding institutional domain now has greater credibility.

## 4 The Structural Limits of an AI Governance Hegemony Battle

All that movement, however, is unlikely to settle into a clear or secure compromise. The most important point is not personal but structural: Chatham House cites the fundamental misapprehension of the most active actors in frontier AI: the United States and China are more interested in national advantage than cooperation; middle powers more interested in expanding capabilities than developing partnerships; public institutions are neither specialized enough nor wield enough power to really impose order; private investors, add together, out-flank the state. Given such circumstances, global governance initiatives are apt to be ambitious in public but weak in practice: a mirror of the inadequacy of symbolic gesture and eager forging. However, between willful self-deceit and proper recognition lies a product of the same overoptimistic mind: that most of the present field of AI governance contains no contractual commitments neither the purpose nor mechanism of any agreement is clear to the imperatives of the United States to maintain frontier advantage and bar it from critical inputs in the future, or to the principles of the unitary soft power for spreading a sovereignty-centric state-centric governance model across the developing world. In effect, this-neither the content nor mechanism of governance- is a rival, an alternative ticket to sovereign control of the digital.

The first level of conflict, then, is conceptual. The American model, at least in its current formulation, assumes that the best way to organize global AI is to diffuse U.S. technologies, to align partners around trusted supply chains and to prevent authoritarian rivals from shaping standards.<sup>[30]</sup> The Chinese model assumes that governance must be anchored in sovereignty, state responsibility, development rights and UN-based inclusivity.<sup>[31]</sup> These positions do not simply reflect rival preferences about administrative design. They encode incompatible answers to more fundamental questions: should private firms or states sit at the center of rule creation; should cross-border data governance stress open flows or sovereign control; should risk management build on liberal rights language or state-led notions of social order and state controllability. This is why compe-

tion is likely to feel more, not less, intense in the coming decades, even if the two superpowers have converged on many words like security, trustworthiness and fairness and even if they have become more similar on those concepts. Even shared vocabulary does not mean shared meaning.

Technological dependence is the second source of conflict. Chatham House observes how governments outside the superpower core are increasingly encountering AI sovereignty as a hard question of ownership and control of their digital infrastructure.<sup>[32]</sup> That challenge is even easier with concentration in the AI stack: Stanford AI Index further informs that 5,427 data centers are in the US, that almost all leading AI chips are produced by a single Taiwanese foundry and that the frontier AI models development happens almost exclusively in the U.S. and China, according to OECD analysis. In that context, governance discussions are virtually indistinguishable from infrastructure decisions. An external actor that adopts American chips, cloud services and models will integrate the same ecosystem of norms, security guarantees and political expectations as the aforementioned state. Conversely, an external actor that partners with Chinese vendors, training programs and integrated solutions will fit a different ecosystem altogether. The fight for rules, therefore, translates into a fight for the installed base. Once dependencies are set, norms tend to follow along supply chains.




Figure 6: Compute concentration shows why AI sovereignty depends on infrastructure ownership as much as formal governance rules.

A third challenge to resolution is institutional fragility. Though the U.N. structures established in 2025 and 2026 to govern AI play an important role, they must be critically examined for their actual impact. The global dialogue on AI governance is described by the United Nations as a forum where all governments and stakeholders can shape global principles and norms on the governance of AI technology and the independent international scientific panel on AI is called the world’s first global scientific body with a mandate to conduct evidence-based assessments for the control of AI. These are advances; they are not mandates. The scientific panel has no coercive authority to control deployment and the dialogue is not a treaty-making platform. The

American vote against global appointment of the forty-member scientific panel in February 2026—the United States was joined by Paraguay, where 117 states voted yes—underscores the flimsiness of these structures.<sup>[33]</sup> Where it can be imagined that open discussion may be helpful, coercive powers are lacking. The problem is that the power of these existing structures is overstated.




Figure 7: Multilateral AI governance attracts broad formal support, but U.S. resistance exposes the limits of universal rule-setting authority.

In this sense, the analogy to the United Nations as a site of symbolism holds to some extent. It is true that high-level meetings can become symbolic instruments when the great powers convene at them without the desire to place a shared set of rules or standards above the strategic profile of individual nations. Chatham House’s grim warning that global governance of AI might only be politically feasible in the wake of a crisis is indicative. It is well familiar to political scientists operating in areas outside AI. States regularly take part in soft-law initiatives while maintaining maximum discretion when it comes to core security and hegemony concerns. AI should be viewed in the same light. The White House Governance Plan emphasizes investment in intelligence, particularly around the collection of knowledge of foreign frontier projects. The United States ensures that the export-control regimes harmonize effectively with the technological limits of U.S.-firm infrastructure. China’s diplomacy regularly invokes national security and sovereignty as the basis for its stance on AI. Given these realities, any credible AI global governance organization will have difficulty overcoming the lowest-common-denominator language of its recommendations if they impinge on strategic autonomy.

This is especially true when considering that private firms wield power on a scale that multilateral institutions have not yet learned to juggle. According to Chatham House, private investment outpaces state capacity; to an even greater extent and based on a far more complex set of variables, the OECD’s venture capital numbers demonstrate how much of AI development is being driven and directed by actors operating across jurisdictions. This complicates classical inter-state governance. The classical institutions of the Bretton Woods era, with the

few existing centrally productive assets, are no longer relevant. They are in the hands of companies whose stake does not necessarily match with national interests or national risk management strategies. However, the U.S. might try for its companies to become the means for diplomatic alliance-building in technology; the Chinese may look for the convergence between state objectives and firms' strategies as an important diplomatic goal. Still, in their hands politics wavers: principles are easier to formulate than the discipline of commercial ecosystems, whose speed of evolution and trans-national spread outpaces what diplomacy alone can realistically constrain.

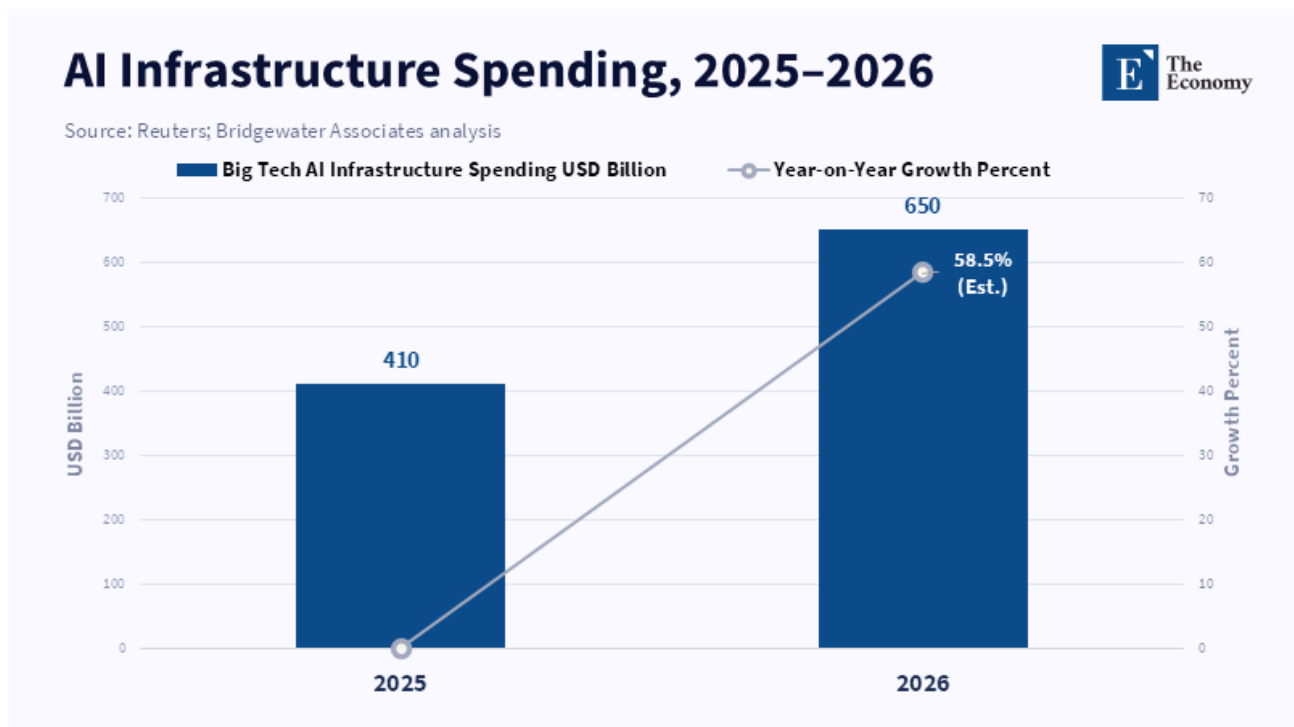


Figure 8: Private AI infrastructure spending is expanding faster than public institutions can realistically govern through diplomacy alone.

This does not mean cooperation is impossible. It does mean that cooperation is likely to be narrower and more technical than many public discourses would lead us to believe. Analysis from the Atlantic Council points out that the 2025 action plans of both Washington and Beijing have some room for joint work on relatively technical standards [34] and for sharing lessons learned from building infrastructure and from limited collaborations on datasets in fields like public health and biotechnology. Brookings also argues that the Trump-Xi summit of May 2026 has provided a platform for the two governments not to leave the global AI governance playing field only on China's terms and for Beijing to express through standards worldwide its concern for risk management.[35] Chinese officials have been reported to have said that there was agreement to establish a government-to-government dialogue on AI during the visit. These are not insignificant openings. They indicate that, despite competition, Washington and Beijing recognize some benefit in talking about risk, standards and incident management. However, they also show the extent of the restrictions. Cooperation in this respect will likely be hampered by dual-use export restrictions, military considerations and suspicion.

This summit illustrates this ceiling. The official Chinese press released summaries of the meeting, noting that the leaders had constructive discussions on AI and agreed to establish bilateral dialogue. But Reuters reports the more expansive trade and aviation agreements announced have only been tentative, while outsiders say the

summit has been largely vague, with no major breakthroughs on issues like Iran, Taiwan, or the AI race. This separation between the summit's ceremonial achievements and the strategic lack of clarity is a pattern: both are likely to have reasons to want an image of responsibility and rule-making on the world stage, but neither side is willing to commit to actions that could limit room for strategic maneuver. A future AI governance forum might end up being another well-known case of high-profile symbolic significance combined with limited concrete results <sup>[36]</sup> and utility as a crisis-management channel rather than a rule-making authority.

Even so, it would be wrong to interpret these new institutions as entirely passive. Weak forums are nonetheless able to generate tangible results, shaping the discourse, procurement expectations, technical literature and coalition practice not of a single state but of many, if they allow a slow but steady process of rule-setting to emerge. Digital standards exchanges are unlikely to lead in the way of formal treaty negotiations, but they can influence certification procedures; a panel of experts can have more sway over what is deemed good evidence than an Asian or European country ever could; a conversation can introduce risk and responsibility categories into the think-tanks and civil services of several nations; a training programme can introduce into several bureaucracies the formats of one country's data standards or modelling practices. Thus, the capacity-building exercise, the list of experts, the standards cooperation and the hegemony of the committee-these may in the future be the defining attributes of digital governance rather than the absence or presence of global agreements or the actions of individual states. Hence, the future is neither full world government nor bifurcation. It is a complex, layered competition. For each function, some forums will be relevant: for imagining the high-end stack and building broad security coalitions, the United States; for naming development sovereignty and standards capacity or activism, China; for universal concern and narrative, the United Nations. The emerging order is therefore likely to be a crude but consequential mosaic.

The policy message is more demanding but also more straightforward. Washington cannot afford to rest on technical hegemony over the longer run if it wishes to deny Beijing governance at the global AI table, where a significant share of developing states perceive arrangements dominated by the West as unrepresentative of their interests. But to strengthen cyber digital means of diplomacy, maintaining a foothold in global institutions and combining are consequently strategic priorities, not optional luxuries. Meanwhile, Beijing will not turn institutional engagement into stable leadership without bringing confidence to other states that its agenda is not merely an instrument to project authoritarianism and technological dependence. For middle powers and other states of the Global South, in turn, possible choices between competing hegemonic orders can be avoided if nations seek interoperable systems, public standards, competing infrastructure and limited technical accords that leave room for national operational autonomy. And the United Nations, finally, should avoid following the mirage of instant global management into pragmatic, incremental services-benchmarking, standards posting, incident registration and aid and capacity building-that can withstand great-power conflict.

## 5 Conclusion - From Responsible AI to Rule-Setting Power

The rivalry over AI governance is portrayed as a moral debate over the responsible use of technology. It is that, but it is also something more intractable and more consequential: a contest over which great power will have

the authority to write the political code of the AI age. The U.S. has not abdicated governance altogether. It has contracted governance into a means of technological leadership, export leverage and alliance management, draining its investment in the universal institutions where it previously commanded agenda-setting authority. China has seen that opportunity in that contraction and commenced a purposeful shift from infrastructure-exporter to budding institutional-builder, knitting standards work and U.N.-centered diplomacy with outreach to the Global South. The outcome, at least for now, is not a war of Chinese hegemony over global AI policy, for U.S. material advantages remain formidable; it is instead an end to the assumption that the rules will simply be authored in Washington. The only order now forming up is some sort of uneven fragmentation-universal forums offering principles and procedures, bilateral and bloc politics that matter materially, vulnerable cooperation until catastrophe precipitates broader alignment. The policy enterprise is not to wait for a perceived impartial consensus but to forge institutions, standards and diplomatic instruments capable of deterring an AI governance breakdown into pure power politics, even as that terrain is today dominated by power politics.

## References

- [1] Ikenberry, G. J. (2011) *Liberal Leviathan: The Origins, Crisis, and Transformation of the American World Order*. Princeton University Press.
- [2] OECD (2026) *Venture Capital Investments in Artificial Intelligence Through 2025*. OECD.
- [3, 4, 14, 15, 27] Stanford HAI (2026) *AI Index Report 2026*. Stanford Institute for Human-Centered Artificial Intelligence.
- [5] United Nations General Assembly (2024) *Seizing the Opportunities of Safe, Secure and Trustworthy Artificial Intelligence Systems for Sustainable Development*. United Nations.
- [6] Council of Europe (2024) *Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law*. Council of Europe.
- [7, 12] Reuters (2025) “US and UK Refuse to Sign Paris AI Summit Declaration.” *Reuters*.
- [8, 11, 30] The White House (2025) *Winning the Race: America’s AI Action Plan*. The White House.
- [9] France 24 (2026) “US ‘Totally Rejects’ Global AI Governance, White House Adviser Tells India Summit.” *France 24*.
- [10] The White House (2025) *Executive Order 14179: Removing Barriers to American Leadership in Artificial Intelligence*. The White House.
- [13] Jones Walker (2026) “The Regulatory Tide Goes Out: What Global AI Governance Retrenchment Means for Organizations.” *Jones Walker AI Law Blog*.
- [16, 34] Atlantic Council (2025) “Reading Between the Lines of the Dueling US and Chinese AI Action Plans.” *Atlantic Council*.

- [17] Council on Foreign Relations (2026) “Trump’s Cyber Strategy Falls Short on China, Iran, and the Threats That Matter Most.” *Council on Foreign Relations*.
- [18] The White House / Office of Management and Budget (2026) *Fiscal Year 2027 Budget Proposal*. The White House / OMB.
- [19] Basu, A. (2026) “China’s Pivot on Global AI.” *Carnegie Endowment for International Peace*.
- [20, 31] Ministry of Foreign Affairs of the People’s Republic of China (2023) *Global AI Governance Initiative*. MFA China.
- [21, 25] Ministry of Foreign Affairs of the People’s Republic of China (2024) *AI Capacity-Building Action Plan for Good and for All*. MFA China.
- [22] World AI Conference (2024) *Shanghai Declaration on Global AI Governance*. World AI Conference.
- [23] State Council of the People’s Republic of China (2025) *Global AI Governance Action Plan*. State Council / Government of China.
- [24] United Nations General Assembly (2025) *Resolution A/RES/79/325 on Artificial Intelligence Governance*. United Nations.
- [26] State Council of the People’s Republic of China / Xinhua (2024) “China to Formulate Over 50 Standards for AI Sector by 2026.” State Council / Xinhua.
- [28] China Daily (2026) “China’s Core AI Industry Reaches 1.2 Trillion Yuan.” *China Daily*.
- [29] Forum on China-Africa Cooperation (2024) *Beijing Action Plan 2025–2027*. FOCAC.
- [32] Chatham House (2025) “AI Sovereignty and the Global Digital Infrastructure Divide.” *Chatham House*.
- [33] United Nations General Assembly (2026) *Independent International Scientific Panel on AI*. United Nations.
- [35] Brookings Institution (2026) “Trump, Xi, and the Prospects for an AI Safety Dialogue.” *Brookings Institution*.
- [36] Carnegie Endowment for International Peace (2026) “Trump, Xi, and the Prospects for an AI Safety Dialogue.” *Carnegie Endowment for International Peace*.