

# AI and Grand Strategy

## *The Case for Restraint*

Erica D. Lonergan and Benjamin Jensen

---

### *Introduction*

Conventional wisdom holds that an AI arms race will define the twenty-first century and could be decided as early as **2030**. The second Trump administration’s **National Security Strategy** proclaims that AI “will decide the future of military power,” echoing Russian President Vladimir Putin’s **warning** in 2017 that whoever leads in AI “will become the ruler of the world.” But what if the defining technology of the twenty-first century actually rewards the most nineteenth-century of strategies: a cocktail of strategic parsimony and geopolitical fatigue, served neat and called “restraint”?

The **AI arms race** is well covered, but it is still unclear what it means for American grand strategy. Since the end of the Cold War, U.S. grand strategy has largely oscillated between variants of liberal internationalism and efforts to ensure that the United States remains the world’s dominant military, economic, and political actor. The new race concerns which groups—not restricted to states—can best mobilize and deploy the resources required to build AI **infrastructure** and **foundation models**. It also concerns who defines a new set of rules about AI governance and the prevailing research system.

Indeed, how AI shapes **grand strategy**, a state’s theory for how to best secure itself, remains poorly understood in the emerging technology competition. AI is reshaping how economies generate wealth, how militaries fight, how states collect and process information, and how leaders engage in diplomacy. Just as important, AI is likely to reshape the role private actors like Big Tech companies play in the strategic policy arena.

*Restraint appears the logical grand strategy for the AI era. It prioritizes butter over guns and focuses on core economic interests rather than unsuccessfully attempting to deter war across the globe.*

In this new world, and despite the **new interest in offensive realism** at play in Washington, restraint appears the logical grand strategy for the AI era. It prioritizes butter over guns and focuses on core economic interests rather than unsuccessfully attempting to deter war across the globe. And it sees utility in substituting traditional approaches to political-military alliances with new economic networks that connect states, firms, and sovereign wealth funds to meet the insatiable demand for infrastructure investments. The net result is a different set of foreign policy priorities that diverge from the series of overseas raids and threats of land grabs that have marked the first year of the second Trump administration, while aligning with the administration's push to **change the regulatory environment** supporting AI development.

### *From Liberal Hegemony to Restraint*

American grand strategy has been mired in stale debates between competing paradigms, largely anchored in different forms of **liberal internationalism**. Washington has often privileged strategies based on **primacy**, emphasizing military might and unilateral action. At other times, this approach has leaned more toward cooperative security, embracing international institutions and the legitimating power of liberal values, such as the **Biden** administration's framing of American strategy as a struggle between democracy and autocracy. Historically, **restraint** has existed on the periphery of grand strategy, though it has gained increasing prominence during the first and second Trump administrations (even as some elements of Trump's grand strategy are clearly at odds with restraint).

These paradigms emerged in the context of a radically different international system, one in which the United States enjoyed unrivaled **global hegemony** and possessed the power and influence to shape international institutions in ways that aligned with its interests and values. In the ensuing decades, there have been profound changes in the distribution of power, especially with China's rise. There have also been extraordinary advancements in the character of transformative, **general purpose technologies** like **AI**. States believe that mastering this technology will confer tremendous strategic advantages. Both trends are creating new strategic imperatives and challenges for the United States: how to simultaneously navigate both a changing *balance* of power and an evolving *character* of power.

Perhaps counterintuitively, some of the core assumptions of a grand strategy of restraint are better suited to securing American interests in the AI age, for two reasons. First, restraint is not fundamentally vexed by changes in the balance of power. China's pursuit of AI dominance is not an existential threat. Given Russia's relative decline and U.S. nuclear deterrence, the United States is secure. While the United States should strive to be the global leader in AI, there will be multiple centers of power in the coming decades. In this sense, perhaps, AI is similar to nuclear weaponry.

*In 2025, major technology firms made investments in AI infrastructure that dwarf the national security budgets of the vast majority of states.*

Second, AI holds transformative implications for economic power in particular. This is naturally compatible with restraint's inclination to focus on privileging "butter" over "guns" in the proverbial

trade-off. Instead of military force, restraint privileges diplomatic engagement and international trade over hard power, similar to **Washington's Farewell Address**.

At the same time, while restraint prioritizes diplomatic engagement and international trade, how it approaches these issues would need to be reimagined for the AI age. New domestic and transnational networks are connecting powerful firms (e.g., OpenAI, Google, AWS, Microsoft, and NVIDIA), state interlocutors (e.g., sovereign wealth funds in Qatar and the United Arab Emirates pouring capital into AI firms), and politicians in a race to **broker** AI deals.

## *When it comes to grand strategy, AI is changing the nature of the game and the players.*

Access to the money and supply chains required to field more advanced AI models increasingly defines strategy. In 2025, major technology firms made investments in AI infrastructure that dwarf the national security budgets of the **vast majority of states**. **Amazon** invested \$100 billion, a \$20 billion dollar increase from 2024. **Google** invested \$85 billion. **Meta** invested \$72 billion and plans to invest **\$600 billion** over the next three years in a bid to compete with **Stargate**, the joint venture between Softbank, OpenAI, and MGX (an Emirati state-owned investment firm). **NVIDIA's** market capitalization, \$5 trillion, makes it technically larger than the German economy.

These investments are chasing the promise of AI's potential to transform the economy. A 2023 McKinsey **study** estimates that generative AI alone (i.e., a form of AI that leverages deep learning and neural networks to produce new, high-quality, human-like content) could add between \$2.6 and \$4.4 trillion annually in value to the global economy. Similarly, a September 2025 JPMorgan **report** notes that in the first half of 2025, capital expenditures related to AI contributed 1.1 percent to U.S. GDP growth, more than the American consumer. At the same time, AI is poised to cause some economic dislocation; with widespread AI adoption, **experts** project that AI will displace approximately 6-7 percent of the American workforce.

In other words, when it comes to grand strategy, AI is changing the nature of the game and the players. The “game” of grand strategy is increasingly focused on innovating and adopting AI at scale in ways that are changing how great powers achieve their security and interests. And “who plays” the game not only includes the state, but also the private actors building the latest models and furnishing the capital and resources to power them. AI's systemic nature makes the influence of the private sector more pervasive and significant than the historical influence of powerful firms, like the **United Fruit Company**, on strategy and policy. This is as much a risk as an opportunity. Altogether, this necessitates updated strategic thinking to articulate new theories for the AI age.

## *Restraint in the AI Age*

Although there are many **variations**, restraint tends to be anchored in the logic of **defensive realism**. It is premised on the idea that the United States enjoys a tremendously advantageous security position and the additional protection of nuclear deterrence. As a result, military engagements are unnecessary at best and dangerous at worst. To maximize its security, the United States does not need to—nor should it—overextend itself to preserve its superpower status. **Restraint** is deeply concerned about the

implications of entangling alliances for drawing the United States into military intervention. It privileges diplomacy and economic engagement over the use of force.

Many of these tenets easily extend to an AI context, though much of the restraint literature has not explicitly made these connections. While much of the U.S. AI policy debate is vexed about China's aspirations to be an AI **superpower**, a restraint grand strategy would be less concerned. From a restraint perspective, gaps between the United States and a great power rival only matter if they are sufficiently large to undermine the United States' security. The core consideration would be whether an AI-powered China poses a fundamental threat to the integrity of the U.S. homeland.

In turn, an updated strategy of restraint for the AI age would more unequivocally focus on the economic bases of national power. It would refocus U.S. priorities around developing American domestic capacity, especially economically, through investing in core AI inputs: data, compute, enabling infrastructure, and supply chains. The pathway to cultivating U.S. security and interests in an AI age would ideally come through domestic investment to the greatest extent possible. What follows from this is a more permissive domestic regulatory environment that enables U.S. technology firms to continue to grow and expand, together with **government-supported investments**, and that encourages the domestic development of AI infrastructure and policies that support **low-cost domestic energy production**. In many respects, the Trump administration has led in this area through its **AI Action Plan** and novel ideas like leveraging **private equity** to build data centers on Army bases.

That said, not all inputs to the AI supply chain can be domestically sourced. When it comes to the production, manufacturing, and fabrication of advanced semiconductors, firms in Taiwan, South Korea, the Netherlands, Germany, and Japan have developed critical expertise and capabilities for the global AI **supply chain** that cannot be easily on-shored. Critical minerals are also essential inputs for nearly every aspect of the AI technology stack. Currently, China **controls** 98 percent of the global gallium supply and processes 90 percent of all rare earth minerals. Recently, China imposed (and then **walked back**) export controls on critical minerals in the context of a trade dispute with the United States. This episode reveals the risks associated with a U.S. AI strategy premised on wielding trade restrictions and other coercive economic instruments to shape the global AI marketplace in ways that favor U.S. technology while restricting Chinese development. Because restraint would not be *inherently* challenged by China's pursuit of AI power, it could turn to less punitive economic strategies, prioritizing economic and diplomatic relationships to enable American access to the global AI supply chain.

AI also holds important implications for the role of alliances and regional priorities. Restraint sees Europe's pursuit of "strategic autonomy" as a net positive for the United States. In this view, the United States should not be solely responsible for underwriting European security. The pursuit of AI power will further shift the strategic importance of different regions. A restraint strategy would likely push for the United States to reorient international relationships around partnerships that facilitate the transfer of critical AI inputs, especially in those areas where the United States lacks domestic capacity (e.g., semiconductor fabrication) or simply lacks the inputs (e.g., critical minerals). New, looser "alliance" networks are likely to emerge around the political-economic coalitions required to sustain large-scale investments in AI infrastructure. These may not be alliances in the traditional sense, but instead an emphasis of diverse, cross-cutting economic partnerships that connect states, wealth funds, and businesses.

AI holds the most significant potential for restraint's approach to the United States' force posture. If AI-enabled military capabilities can effectively **substitute** machines for manpower and **algorithms** for military intelligence officers, the result could be significant reductions in the size of the military, which is key given that rising personnel costs have **constrained force modernization**. It could also lead to less bureaucratic red tape as well as streamlined administrative and logistical processes within the Pentagon, resulting in a better return on investment to the American taxpayer. Greater private sector partnerships will likely speed up the delivery and development of lighter, cheaper, AI-enabled military capabilities. The U.S. Air Force's **Collaborate Combat Aircraft** system, which relies on uncrewed, AI-powered aircraft alongside crewed fighter jets, is just one example. In this scenario, the U.S. military could substitute drones for personnel at an increasing rate and also reduce the defense budget without sacrificing its ability to generate combat power.

But restraint will also need to be reimagined and updated for the AI age. Restraint strategies tend to be more focused on the military dimension of power, making the case for reducing the United States' military engagements. They have not grappled as much with the implications for **economic policy** and the role of the private sector. Yet, these are precisely the issues that will be especially salient for any AI grand strategy.

Regarding economic policy, in some ways restraint's domestic focus is congruent with the increasing popularity (on both the American right and left) of economic nationalism and populism. Applied to AI, this could include policies that promote the **onshoring of supply chains**, limit global data exchange, and implement export controls and tariffs to support domestic reindustrialization and compensate for job losses. But a wholesale embrace of economic nationalism is in tension with elements of building the United States' AI power. For example, restraint will need to navigate difficult trade-offs between inclinations to limit global data sharing in the interests of maintaining American sovereignty and autonomy, on the one hand, and the need to ensure the widest possible access to data to improve the quality of AI models. Analogous trade-offs will emerge in other areas, particularly regarding tariffs and other coercive economic and financial instruments, as well as immigration and the global fight for talent.

Similarly, the private sector will play an outsized role in a restraint-based grand strategy focused on building the United States' AI power. Coalitions among private sector actors will likely push for deregulation across relevant sectors (i.e., not only the AI tech sector, but also related industries such as the energy sector). This will give rise to another key source of tension: Washington's effort to promote economic policies to enhance domestic production and autonomy versus the private sector's interests in access to global markets, talent, data, minerals, infrastructure, and the like. Altogether, a restraint grand strategy that results from the economic interests and preferences of domestic coalitions that converge around generating AI power is likely to have different characteristics than the original vision of restraint.

Any AI grand strategy will have to contend with who makes the rules and the larger policy and research environment most conducive to advancing the technology and its transformative potential. The strategic landscape is shaped by competing visions of control, risk, and architectural evolution. Thinkers like **Stuart Russell** argue that the decisive variable is not resource mobilization, but ensuring systems remain provably aligned with human preferences to avoid a catastrophic arms-race dynamic. **Nick Bostrom** expands this perspective to focus on governance. While compute and talent grant near-term leverage, the lack of global governance structures creates a "vulnerable world" where competition among private and state actors accelerates existential tail risks. Conversely, Yann LeCun

challenges the necessity of centralized control. He **advocates** for open research and decentralized ecosystems, suggesting that current foundation-model paradigms are merely transitional and that power will ultimately stem from interoperable, world-model-based architectures rather than the sheer hoarding of infrastructure.

### *Implications for the Future of American Grand Strategy*

Historically, grand strategy has often been an **illusion**, and states have struggled to articulate a clear **vision** amid the chaos and friction of the moment. What matters then are the policy incentives and resource bets a state mobilizes to change its position in world politics.

### *Nativism is counterproductive to sustaining the AI transformation.*

Seen in this light, a grand strategy of **restraint** for the AI era aligns with some, but is in tension with other, aspects of Trump’s “America First” approach to world politics. The **National Security Strategy**, along with calls for a new **Unified Command Plan**, indicate a pivot toward a different world order, while **deregulation plans** for AI and the energy sector promote the type of domestic economic incentives required for twenty-first century economic growth. And while the administration talks about a desire to **reduce the burden** of the United States’ traditional alliances and security commitments, military operations seem to be increasing, risking overextension and another generation of endless wars. Restrictive trade and immigration policies alongside **threats to research and higher education** are fundamentally at odds with a restraint-driven effort to empower the United States in the AI age. Over the next three years, there will need to be a debate.

Adopting any grand strategy demands making difficult choices, accepting trade-offs, and setting priorities. If the second Trump administration truly seeks to position the United States for a new technological age, it should forego elements of its current strategy that aim to perpetuate overseas military commitments and wall off the United States from the world in ways that will inevitably undercut the development of American AI power. Nativism is counterproductive to sustaining the AI transformation. ■

*Erica D. Lonergan is an assistant professor in the School of International and Public Affairs at Columbia University. Benjamin Jensen is director of the Futures Lab and a senior fellow for the Defense and Security Department at the Center for Strategic and International Studies.*

*This report is made possible by general support to CSIS. No direct sponsorship contributed to this report.*

**This report is produced by the Center for Strategic and International Studies (CSIS), a private, tax-exempt institution focusing on international public policy issues. Its research is nonpartisan and nonproprietary. CSIS does not take specific policy positions. Accordingly, all views, positions, and conclusions expressed in this publication should be understood to be solely those of the author(s).**

© 2026 by the Center for Strategic and International Studies. All rights reserved.