

FOREIGN AFFAIRS

NOVEMBER/DECEMBER 2025 • VOLUME 104 • NUMBER 6

The AI Grand Bargain

What America Needs to
Win the Innovation Race

BEN BUCHANAN AND TANTUM COLLINS

The AI Grand Bargain

What America Needs to Win the Innovation Race


BEN BUCHANAN AND TANTUM COLLINS

The United States' lead in artificial intelligence might seem unassailable. U.S. companies—Anthropic, Google, OpenAI, and xAI—are out in front across almost all assessments of the technology's general capabilities. American AI models are outperforming doctorate-level scientists on challenging questions in physics, chemistry, and biology. Just a few American AI and chip giants are worth more than the entire Chinese stock market, and investors from across the world are plowing ever more resources into the American AI ecosystem.

This breakneck progress is, in many ways, a testament to the strengths of the model of American AI development that has dominated for the last decade: letting the private sector operate on its own, with remarkably

BEN BUCHANAN is the Dmitri Alperovitch Assistant Professor at Johns Hopkins School of Advanced International Studies. From 2021 to 2025, he served in a variety of roles in the White House, including as Special Adviser for Artificial Intelligence.

TANTUM COLLINS was Director for Technology and National Security on the National Security Council from 2023 to 2025.



little direct government meddling or resourcing. This approach is quite different from those that ushered in past breakthrough technologies. Nuclear weapons and power, space travel, stealth systems, personal computing, and the Internet emerged either directly from U.S. government efforts or on the back of significant public funding. AI also has roots in government-funded science, including in personal computing and the Internet, and it benefits from ongoing government-supported research. But scaling up AI has been essentially a private-sector activity.

Yet there is reason to think the American way of developing AI is reaching its limits. Those limits will likely become increasingly evident in the coming months and years, and they will start to erode—and perhaps even end—U.S. dominance. Eventually, they will place the United States at a disadvantage against China, which has an alternative approach to the AI contest.

To avoid that outcome, Washington will need to embrace new ways of advancing AI development, ones that demand much tighter mutual support between the private sector and the state. Further progress now depends on resources and capabilities that only the government can provide or facilitate: the energy to power ever-larger data centers, a pipeline of international talent, and effective defenses against sophisticated foreign espionage efforts. The U.S. government, for its part, will need the cooperation of the private sector to integrate AI into the national security apparatus and to make sure the technology does not undermine democracy across the world.

The new American model of AI, in other words, must rest on a grand bargain between the tech industry and the government. The tech sector can help the state make sense of and deploy AI. The state can help the tech sector continue to grow in a way that advances everyone's interests.

MAXING OUT

It is easy to see why Washington's light-touch approach to AI has, by and large, paid dividends. Past revolutionary technologies, such as nuclear weapons and space flight, did not have immediate commercial applications. But the business case for modern AI is already highly compelling. AI firms have found huge user demand, resulting in skyrocketing revenues, and they have promised to automate myriad valuable tasks, such as coding. As a result, capital markets are funding AI projects at scales that would historically have required government

resources. Moreover, the computation-centric nature of today's AI means that it builds neatly on the cloud computing infrastructure that the private sector, not the government, has mastered.

The sufficiency of private-sector capital in enabling AI advances is wonderful for taxpayers, but the limits of this approach are becoming apparent. To see why, look at infrastructure. The vast fleets of computer chips needed to develop and use today's AI require extraordinary amounts of energy, so U.S. companies will need more power to fuel the data centers they plan to build in the coming years. An analysis by Anthropic estimated that the United States will need to produce 50 gigawatts of new power just for AI by 2028—roughly equivalent to what the entire country of Argentina uses today. (One of us, Buchanan, advises AI and cybersecurity companies, including Anthropic.) By then, data centers could consume up to 12 percent of American electricity production. Without more electricity, the AI build-out will stall. Amazon's CEO, Andy Jassy, for example, has labeled power the "single biggest constraint" to AI progress. And building this level of new infrastructure will require government help.

For too long, Washington did too little to add new power to its grid. From 2005 to 2020, the United States added close to zero net new power. After U.S. President Joe Biden took office, in 2021, and passed a law subsidizing the construction of clean energy infrastructure, the country added more than 100 gigawatts in new capacity. In the last days of his term, he signed an executive order specifically aimed at further expediting the AI and clean energy build-out. But although his successor, Donald Trump, has said the right things about building new energy infrastructure for AI, he has not delivered. He signed an executive order to accelerate federal permitting for data centers, but implementation remains nascent. Worse yet, his signature "One Big Beautiful Bill," passed in July, and other executive actions gutted key parts of Biden's energy expansion efforts, such as vital transmission projects. An area that could have been a bipartisan success fell prey to politics and has now become a major concern for business and AI competitiveness.

Executed well, an AI-fueled energy boom would have benefits far beyond AI development itself. Leading AI companies are spending

There is reason
to think the
American way of
developing AI is
reaching its limits.

hundreds of billions of dollars on infrastructure build-outs, creating employment opportunities. They have committed to carbon-free operations and demonstrated a willingness to pay higher prices for clean energy. These massive investments can accelerate the domestic development of better energy sources, many of which have bipartisan appeal, such as advanced geothermal power and next-generation nuclear facilities. Powerful AI models could also accelerate climate-related research.

If the United States does not construct more energy capacity, however, American AI firms will feel pressure to outsource the development of strategically critical facilities—likely to oil-rich regions such as the Gulf that run on dirtier fuel. For Washington, any prospect of offshoring AI should set off alarm bells. An American company shifting advanced AI training to a foreign country, especially an autocratic one, would pose huge risks as AI begins to power more of the U.S. economy and to play an integral role in defense. If a host country became unhappy with American behavior, it could punish Washington with the flick of a switch. A failure to build domestic energy capacity would thus echo outsourcing mistakes of past decades in other important industries, such as semiconductors, in which the United States is now dependent on foreign suppliers.

The United States has the technology and industrial capacity needed to build new energy facilities. But it remains inhibited by a thicket of government and utility regulations and by procedural delays—some backed by good reason, some not. These restrictions impose huge delays in interconnection (the process of connecting a new power source or data center to the grid) and require years-long environmental assessments. On top of federal and utility hurdles, state and local policies can be cumbersome, especially for projects that cross multiple states, such as transmission lines. Companies—not citizens—should pay for the energy build-out, but government policies must make it possible for them to undertake these projects on reasonable timelines.

Infrastructure is not the only domain in which American policies hold back the AI sector. AI progress depends as much on talented people as it does on technology and computing power, which is why the Biden administration acted aggressively to enable people of extraordinary technical ability to come to the United States from all over the world. This effort included elevating AI and other high-tech fields as priority areas for visas, as well as updating eligibility criteria to welcome top-tier scientists.

Here, too, Trump has sometimes said the right things but always fallen short in practice. During his 2024 presidential campaign, he declared that when foreign students graduate from U.S. colleges, they should “automatically” get a green card. Instead, under his leadership, the United States has begun shutting itself off to foreign workers and students, intimidating even those who have visas and are in the country legally. In September, the administration even said it would slap a \$100,000 fee on applications for H1-B visas—the visas most commonly given to high-skilled immigrants. Such actions are already having adverse effects. For example, preliminary research by NAFSA, a nonprofit association of international educators, suggests that in 2025, American universities will suffer a 30 to 40 percent reduction in international enrollment.

If Washington cuts itself off from foreign-born scientists or sends them back home, the consequences will be catastrophic. The United States leads the AI race in large part because it attracts experts from across the world. According to a Georgetown University study that looked at AI research from 2010 to 2021, 70 percent of top U.S.-based AI researchers were born abroad. Sixty-five percent of leading U.S.-based AI companies, as ranked by *Forbes*, have at least one immigrant co-founder. Before the current Trump presidency, 70 percent of the students enrolled in American AI graduate degrees hailed from abroad. Historically, the vast majority of these students have stayed, often making critical contributions to American industry and academia. But thanks to Trump’s policies, many of today’s students might instead return home. Some could head to China, which has spotted an opportunity to recruit AI experts—and which poses the most significant challenge to the United States’ AI interests.

THE CHINA CHALLENGE

Protecting American AI leadership is not just a matter of pride. It is essential to U.S. national security and economic competitiveness. China has been making strides in AI development, and although none of its firms can yet match the best American ones, they do not lack for technical talent.

China faces one devastating disadvantage in this competition: its inability to make large quantities of advanced AI chips, a weakness exacerbated by U.S. export controls that began in the first Trump term and that Biden greatly expanded. But after heavy lobbying by

industry, the second Trump administration has started dismantling this area of bipartisan consensus. In July, for example, it reversed its April decision to cut China off from newer AI chips, and the president indicated plans to undo other Biden-era controls, as well. Such moves will almost certainly accelerate Chinese AI development.

Beijing has acted decisively to ensure that, if it gains access to such chips, the rest of its AI ecosystem will supplant the United States'. Consider energy. Beijing has made extraordinary investments in power plants,

If Washington cuts itself off from foreign-born scientists, the consequences will be catastrophic.

energy storage, and energy transmission. As a result, China now produces more than twice as much electric power as the United States, and its lead is expanding. In some individual months, the country has installed over 90 gigawatts of new, clean energy capacity—almost double the amount of energy American AI firms will need in the next several years.

Beijing has also gained an edge by fusing its AI industry with its national security apparatus. The U.S. Department of Defense

has said that major Chinese AI companies, such as Tencent, are key pillars of China's military-civil fusion strategy. AI systems like these have broad utility to military and intelligence agencies: they can support weapons development, cyber-operations, and domestic surveillance, among other tasks. In exchange, the Chinese government has provided technology firms with extensive policy and security support. Historically, this assistance has included defense services and passing along industrial secrets stolen from American businesses.

Losing AI leadership to China would also cause tremendous global harm. Today, consumers across the world benefit from the rigor and transparency of U.S. regulations and standard setting, often developed in concert with other democracies, in many technological domains. For example, new technologies such as electric vehicle charging have required global collaboration on standards. AI will require similar partnerships, and it is in Washington's interest to take the lead. Otherwise, there is a risk that autocracies will unilaterally set the standards. If Washington cannot establish more and better relations with its AI sector and secure its global AI leadership, the Chinese national security apparatus might shape global standards to adhere to Chinese censorship rules.

THIS FOR THAT

Right now, strong partnerships between AI firms and U.S. national security agencies are few and far between, and those that do exist are in the early stages. To address this deficiency, the government will need a better understanding of what AI is and how it functions. The government can help U.S. industries, but only in sectors it comprehends deeply, and today AI is not among them. In our time in government, we worked with some truly tech-savvy civil servants and military officers who drove bureaucratic change. The Biden administration hired hundreds of AI experts to bolster their ranks. But many have been fired or have left government in recent months, including many of its top technical staffers. Washington will have to do a better job of securing AI talent and reverse the trends that Trump has set in motion.

Industry leaders, meanwhile, must make sure U.S. officials understand their work—and they will have to be more responsive to Washington’s needs. AI leaders in Silicon Valley might be reluctant to cooperate more closely with the government, given its technical ignorance and bureaucratic sluggishness and given that they have enjoyed such success flying solo. But executives should remember that collaboration between industry and Washington often works out well for everyone. U.S. President Franklin Roosevelt’s administration teamed up with Ford Motor Company in the run-up to World War II to produce the B-24 Liberator heavy bomber, generating both revenue for Ford and much-needed aircraft for the military. The Manhattan Project would not have succeeded without DuPont, General Electric, Chrysler, and other corporate stalwarts, all of which profited by providing help. And the invention and refinement of radar, satellites, jet aviation, microprocessors, and the Internet all flowed from corporate-government teamwork.

One critical area in which the government can help companies is security assistance. Because AI is increasingly essential to national security, foreign intelligence services are redoubling their efforts to steal innovations from U.S. tech businesses. In March 2024, for example, the U.S. Department of Justice indicted Linwei Ding, a Google software engineer, for allegedly passing Google’s AI chip designs to China. Beijing has also worked hard to smuggle in advanced American technology products, such as AI chips themselves.

Chinese spies will grab whatever AI secrets they can from American businesses. But they have good reason to focus on stealing model

weights—the numbers that encode a trained AI model. Chinese companies can then avoid the training costs, as well as cut down on the time it takes to develop models. For example, AI companies work hard to discover algorithmic tricks that enable them to use their computing power more efficiently. For China, with its severe constraints on computational power, such multipliers are almost certainly incredibly valuable. These multipliers are likely far less well defended than the core secrets of past eras, such as during the atomic and space ages, simply because the government has been largely uninvolved in their development.

AI companies bear the primary responsibility for defending their networks and organizations. But despite the recent loss in expertise, the U.S. government has cybersecurity capabilities that businesses cannot match, and it should provide significant assistance, as it does today to firms in the defense industrial base and critical infrastructure sectors. Such aid could include intelligence about foreign hacking attempts, support in vetting international talent, and guidance on security procedures. Companies that work directly with the U.S. government on national security should meet stringent standards, akin to those imposed on other defense contractors.

To meet their side of the grand bargain, AI businesses should help the United States incorporate their technologies into the national security apparatus. This lack of integration is a persistent area of weakness for Washington. American companies may lead the world in inventing AI, but without such cooperation, the country will fall behind in adapting AI for military purposes, which could prove devastating in a conflict. Time and again throughout military history, states that fail to integrate new technology into their armed forces have wound up suffering. France and the United Kingdom, for example, invented the tank during World War I but paid a heavy price when the Germans were the first to master its use, enabling their powerful blitzkrieg offensives of World War II. To guard against a similar outcome, the Department of Defense and intelligence agencies must procure and use frontier AI systems, and they will need hands-on guidance from private-sector technical experts to do so effectively.

Biden created an outline for how such cooperation might function in October 2024, when he signed a national security memorandum directing the government to bolster its use of AI for national security purposes. This document included strict guardrails to ensure that the

technology does not enable human rights abuses, domestic surveillance, or other unethical activity—protections essential to earning the trust of private-sector AI developers and the general public. Unlike with many of Biden's other executive actions, Trump has not yet repealed it. But the president has made little progress on its many important provisions and has ousted some of the senior nonpolitical experts vital to its implementation.

As government and industry develop a better national security partnership, they will need to pay particular attention to how AI can solve pressing problems in Washington's competition with Beijing. There are, for example, many ways the United States could use AI in cyber-operations, as illustrated by the AI Cyber Challenge conducted by the Defense Advanced Research Projects Agency. The country that more quickly and effectively integrates AI into the cyber-domain will likely prove better able to protect its own networks and penetrate others', unlocking a huge intelligence advantage. AI could also improve Washington's capabilities in other security domains, such as geospatial intelligence, signals intelligence, logistics, and weapons design. But nothing will happen without clear direction from government and meaningful engagement by American AI firms.

STRANGER DANGER

There is a final reason why the U.S. government and leading AI companies need to build closer ties: together, they will have to weigh tradeoffs between upsides and risks. Within the technology world, almost everyone agrees that although AI could offer tremendous benefits to humanity—curing diseases, advancing clean technology, eliminating grunt work—it could also cause massive harm. Some of these risks, such as the use of AI by authoritarian states to reshape the global order, can be avoided by preserving and expanding Washington's lead over competitors. But other dangers are trickier to mitigate. Several top AI thinkers, for instance, believe it is plausible, or even likely, that a single malicious user could someday harness powerful AI to engineer a deadly novel pathogen. Others worry that even in benevolent hands, powerful algorithms might cause catastrophic accidents by taking actions that their creators do not intend. Then there are less fantastical but still dire outcomes, including massive unemployment, acute concentration of economic power, and discrimination in sectors such as health care because of biased models and training data.

From a policy perspective, any one of these scenarios would pose a challenge of historic proportions and force difficult trade-offs. For example, in a hypothetical world in which a single AI user can cause catastrophic harm, the government will have to consider sweeping regulations on the development and use of sophisticated systems, even if that slows innovation. And if AI automates away a large share of human work, the government may have to spend large sums of money retraining the workforce, or it might have to facilitate a restructured

Beijing has gained an edge by fusing its AI industry with its national security apparatus.

economy. Given the pace of AI progress, policymakers will have to make these consequential decisions under exceptionally tight timelines—all concerning a technology that government is not inventing and about which it knows worryingly little.

Deeper collaboration between the public and private sectors, as well as with civil society, does not guarantee that the state will make the right calls. But it does give Washington a fighting chance of securing a net-positive outcome.

With a stronger technical foundation, officials can better understand how reliably AI systems follow instructions, how they handle dangerous tasks, in which areas they can replace human labor, and to what extent they favor offense versus defense in security and safety domains.

The rise of the Center for AI Standards and Innovation at the Department of Commerce (founded as the AI Safety Institute under the Biden administration) represents a valuable initial step to build meaningful collaboration. Since its inception, CAISI has brought together government officials and companies to collaborate on safety issues. It has also aided in the development of standardized testing mechanisms for AI. CAISI has worked alongside other agencies with domain-specific expertise to carry out additional voluntary testing on particularly critical topics, such as partnering with the Department of Energy and the AI company Anthropic to assess whether frontier AI models have dangerous knowledge about nuclear weapons. CAISI featured prominently in Trump's AI Action Plan, and the administration must empower it to carry out voluntary collaboration with companies, to set standards, and to conduct safety testing.

Thanks to CAISI's work and the voluntary commitments that leading AI companies made to the Biden White House, AI firms have

already promised to conduct independent safety testing of their models, often based on CAISI guidance. In some cases, companies have even agreed to grant CAISI access to new systems before they are released and have praised the government for the national security-specific expertise it has offered in return. Both sides should deepen this collaboration, spending more time and resources building high standards and conducting rigorous assessments of new models.

FROM THE GOVERNMENT, HERE TO HELP

Grand bargains often work better as tag lines than as policy, and getting the right kind of deal when it comes to AI is easier said than done. The technology, after all, is rapidly progressing along an unpredictable path. As AI improves, ever-larger amounts of infrastructure, power, and money will be required; the need for improved security from foreign intelligence threats will increase; and the urgency of collaboration with the defense apparatus will grow. So will the risks of misuse, prompting new policy tradeoffs. More startups will arrive on the scene, and legacy companies that today look unstoppable may fall by the wayside. Everyone involved in the AI world should prepare for constant renegotiation and rebalancing. U.S. officials, for their part, will almost certainly have to remain agile, experimenting with different AI policies as time goes on.

But amid this uncertainty, it is imperative that Washington take a more active role in enabling and shaping the American AI ecosystem. The technology does not need to develop as nuclear weapons did—under strict state control—but Washington cannot sit this one out. Instead, AI should perhaps evolve as the American railroads did in the 1800s. The private sector handled most planning and construction, but the government played a vital role, as well. It organized laws and permits for building the infrastructure. It passed carefully calibrated, common-sense government safety requirements—such as standardized track gauges, rules for the use of air brakes, and requirements for car coupling—which all helped make trains both faster and safer. The collaboration was not perfect, but it worked: American railroads became a national asset that increased the United States' security and prosperity. Advanced AI, too, can promote U.S. power and interests, provided it is developed in the right way and under the right set of arrangements. Now, as before, it is time for the public and private sectors to stand shoulder to shoulder. 🌐