CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

DECEMBER 2024

The AI Export Dilemma: Three Competing Visions for U.S. Strategy

Sam Winter-Levy

The AI Export Dilemma: Three Competing Visions for U.S. Strategy

Sam Winter-Levy

Carnegie Endowment for International Peace
Publications Department
1779 Massachusetts Avenue NW
Washington, DC 20036
P: + 1 202 483 7600
F: + 1 202 483 1840
CarnegieEndowment.org

This publication can be downloaded at no cost at CarnegieEndowment.org.

Contents

Introduction

How widely should the United States share its artificial intelligence (AI) technologies? This question may soon become a defining issue for U.S. foreign policy and economic strategy, yet it has received surprisingly limited public attention. While Washington focuses intensely on constraining AI advances by China, another group of emerging economies—including states like Saudi Arabia and the United Arab Emirates (UAE)—are increasingly positioning themselves as influential players in the AI landscape. U.S. policymakers are only just beginning to grapple with the opportunities and dilemmas posed by such countries' AI aspirations.

On the one hand, the growing global appetite for U.S. AI technology—including advanced chips, massive data centers, and frontier models—can be a source of leverage to court so-called swing states and shore up American influence. And U.S. technology companies, facing increasingly steep capital, land, and energy requirements as they scramble to conduct what one analyst has called "the largest infrastructure buildout that humanity has ever seen," see partnerships with various foreign countries as an answer to many of their prayers. Yet the proliferation of powerful AI systems, even to ostensibly friendly nations, comes with serious risks—including intellectual property theft, misuse by authoritarian regimes, and the siphoning of some of the United States' most advanced technologies to the Chinese military and other adversaries. Balancing these interests will be a central task for U.S. policymakers for years and probably decades to come.

The Joe Biden administration has quietly debated these dilemmas and made a few early calls—for example, expanding the restrictions on exports of top-end AI chips beyond China to other countries, including some in the Middle East, while also signing off on a major deal

between Microsoft and the UAE's leading tech company, G42. (The Gulf states have been a particular focus of discussion given their desire to play a major role in the AI race and the capital and energy resources they bring to the table.) But other proposals, like country-specific caps on computing power and export controls on AI models, remain pending. National Security Advisor Jake Sullivan has promised to release a "new global approach to AI diffusion" in Biden's final months in office, but any lame duck announcements will have muted impact as Donald Trump's new administration takes a fresh look at these novel and weighty issues.

To help advance U.S. deliberations, this article offers a map of key perspectives and debates among the small group of actors who have thought deeply about U.S. AI exports and who are likely to shape policy under the Trump administration—drawing on dozens of interviews over the past three months with private companies, academics, and policymakers from across the American political spectrum, plus foreign government officials and analysts in places including Abu Dhabi, Bangalore, and Riyadh. The article lays out three distinct approaches that Washington might take in determining where and how to export advanced U.S. AI technology.

Those who favor **a strategy of control** believe that the risks of allowing potentially transformative AI technology to proliferate outside a handful of close U.S. allies outweigh the benefits of expanded American AI exports. This approach is particularly influential among national security officials and analysts who forecast dramatic increases in the power of AI systems in the near future. Those who instead favor **a strategy of diffusion** make the opposite calculation: They emphasize the benefits associated with the rapid entrenchment of U.S. AI technology in foreign markets. This view draws its support primarily from U.S. business interests, as well as from analysts with longer and more gradual timelines for the emergence of powerful AI systems. A third approach, **a strategy of leverage**, stakes out a middle ground: It seeks to use U.S. AI exports as a bargaining tool to extract geopolitical and technological concessions. Its supporters include promoters of international AI safety and security standards, along with more traditional foreign policy actors searching for leverage to shore up U.S. influence in swing states.

These are crude, stylized groupings; no attempt to draw such broad categories will be perfect. At times, policymakers will endorse or be able to combine elements of each of them. And some people may align themselves with different strategies depending on the specific technological input under discussion: It is possible, for example, to support controls on chips while opposing controls on software. But on other occasions, policymakers will have to choose between distinct responses to the trade-offs, risks, and opportunities presented by rapid and unpredictable technological progress. This article provides a broad intellectual framework to help situate decisionmakers, expose underlying assumptions, and identify areas of potential convergence or ongoing uncertainty.

A Strategy of Control

A strategy of control lies at one end of the spectrum of possible approaches to international AI policy: It seeks to prevent the diffusion of cutting-edge AI technology to anyone outside a small club of U.S. allies. Such a club would certainly include the other four countries in the Five Eyes intelligence partnership (Australia, Canada, New Zealand, and the United Kingdom). It would also probably involve other major treaty allies with key roles in the AI value chain (such as Japan, the Netherlands, and South Korea) and perhaps cover other advanced democracies as well. But it would exclude countries—like India, Saudi Arabia, Türkiye, or the UAE—that have weaker alignments with U.S. interests, stronger ties to U.S. adversaries, and more transactional foreign policies.

Mechanisms for implementing this strategy include export controls and licensing requirements on elements of the AI supply chain—especially graphics processing units (GPUs), major clusters of computing power, and frontier model weights (the parameters that encode the core intelligence of an AI system). Other possible mechanisms include controls on the activity of U.S. persons supporting the development of these technologies in certain countries. The goal would be to ensure that the United States and its trusted partners retain physical control of the means to develop and deploy the most advanced AI systems.

Onshore Core Strategic Assets

The rationale for this approach was most clearly articulated in a widely read memo published in June by Leopold Aschenbrenner, a former OpenAI researcher. In his view, breakthroughs in AI could lead to extraordinary accelerations in scientific and technological progress, with major implications for the global balance of power—and these breakthroughs could occur within a few years. If AI systems have the potential to drive explosions in economic growth, design new synthetic bioweapons, and develop impressive new cyber capabilities, then AI may become a primary tool of strategic competition. In other words, securing control of advanced AI technology would be the most important national security and economic project of the coming decades.

Consider data centers as an example. Building them in semifriendly countries such as the UAE or Saudi Arabia creates serious risks for U.S. national and economic security. Physical access would make it easier for the host state to steal model weights, training data, algorithmic insights, and more. (For example, Saudi agents in 2014–2015 infiltrated and stole sensitive information from Twitter, now X, and such operations would be even easier to conduct on local soil.) The local government could even seize data centers if officials believed major breakthroughs were imminent. If this sounds far-fetched today, that is because AI infrastructure has not yet become the world-shaping strategic asset that Aschenbrenner and others project it soon will. Once that changes, previous historical episodes—like Iranian and Saudi nationalization of U.S. and British oil interests in the 1950s and Egypt's 1956 seizure of the Suez Canal—become relevant parallels.

In addition, advanced AI systems will almost certainly attract sophisticated cyber attacks from U.S. competitors like China seeking to steal U.S. models. Housing and protecting these systems within the United States or its closest security partners may provide significant defensive advantages, whereas the Gulf states are more hospitable environments for Chinese espionage, and their intelligence sharing and insider threat programs lag far behind those of the United States' Western allies. More fundamentally, Washington faces the question of whether to place its major technological capabilities at the whim of regimes whose interests only partially align with those of the United States. "Do we really want the infrastructure for the [next] Manhattan Project," Aschenbrenner writes, "to be controlled by some capricious Middle Eastern dictatorship?"

This argument is likely to appeal to some—though not all—elements of the incoming Trump administration. Ivanka Trump, the president-elect's daughter and a senior White House adviser in his first administration, has described Aschenbrenner's memo as an "excellent and important read." Influential Republican national security voices in Congress have helped lead bipartisan bills to "keep American AI out of China's hands" by prohibiting American AI companies from releasing open-weight models and have expressed major reservations about Microsoft's decision to build data centers in the UAE. And in announcing his picks for interior secretary and Environmental Protection Agency administrator, Trump emphasized the importance of those roles in helping America "win the A.I. arms race with China (and others)." This suggests that Trump will prioritize permitting reform and energy policy to enable the United States to build key AI infrastructure domestically, rather than offshoring it—and that his team is not focused solely on handicapping China but also harbors concerns about third countries.

Prevent a Multiplayer AI Race

Another argument for the strategy of control relates to so-called AI safety—a loose term for the management of AI risks, in particular catastrophic security threats such as rogue autonomous behavior, AI-powered weapons proliferation, or critical infrastructure failures. There are heated debates about how to balance the benefits of rapid progress in AI capabilities against the risks that ever-more-powerful AI systems will be misused or go rogue. Today, Americans wield outsized influence in such decisions because U.S. companies have a healthy lead in algorithmic progress and control of the preponderance of the world's supply of computing power, or compute. In theory, U.S. companies or regulators could therefore take proactive measures to address AI safety concerns as the technology gets more powerful. For example, U.S. organizations have led the world in publishing relatively detailed voluntary frameworks for risk mitigation, and a pathbreaking regulatory initiative reached the California governor's desk (before ultimately being vetoed).

But if the United States approves the export of major AI compute clusters and advanced model weights, it will wield significantly less control over the technology's evolution. Other actors might enter the race to advanced AI systems, potentially exacerbating competitive

dynamics that undercut AI safety and security efforts. A two-way race between the United States and China is already giving both sides ever-stronger incentives to cut corners and plunge ahead to develop highly capable—yet risky—AI systems. Managing this race-to-the-bottom could become even harder if the United States enables other AI powers to rise and join the frontier. If a two-way race becomes a race among three, four, or more players, then even more complex and competitive dynamics could dramatically shrink the room of all players to conduct rigorous testing and research and apply reasonable guardrails.

Of course, whether the AI safety argument will convince the incoming Trump administration remains to be seen. On the one hand, many Republicans view AI safety as an offshoot of a broader progressive cultural agenda to use big tech companies to censor right-wing speech. The GOP's 2024 platform, which was tightly controlled by Trump, promised to eliminate Biden's safety-oriented executive order because it "hinders AI Innovation" and "imposes Radical Leftwing ideas on the development of this technology." "Republicans," it said, "support AI Development rooted in Free Speech and Human Flourishing."

But Elon Musk, who appears to have the ear of the president-elect, has expressed major concerns about AI risks, describing AI in the week before the election as a "significant existential threat" with a 10 to 20 percent chance of "go[ing] bad." One participant on Project 2025's AI Policy Committee, meanwhile, has said "that Trump's supposed shadow transition takes AGI [artificial general intelligence] and its associated risks seriously." Trump himself has expressed some sympathy with "those people that say [AI] takes over the human race," and he may come to care about polls suggesting that voters, too, are concerned by AI risks. It might therefore be premature to discount the role that AI safety considerations—or at least, a core set of concerns about AI-driven security threats—will play in the Trump administration's international policies.

Preserve a Lead Time

The strategy of control relies on the fact that the United States and its allies currently dominate key elements of the AI value chain. U.S. firms design and own the plurality of the world's high-end GPUs, and their leading-edge compute advantage is, for now, stable. Chinese yields on advanced chip production are low, and since China currently cannot produce enough leading-edge chips for its own purposes, it is unlikely to devote its limited supply to exports. This means Beijing may not have the technological capability or willingness to step in as an alternative supplier if the United States maintains its export controls on countries like the Gulf states. But this may not be true forever—it is a contingent feature of the current economic moment, and forecasts of future Chinese technological capabilities must come with wide error bars. The United States and its allies are unlikely to enjoy such an unchallenged position in the semiconductor supply chain forever.

Still, for advocates of this approach—who belong to a long lineage of national security hawks who have sought to preserve U.S. advantages in militarily useful technologies—preserving a lead time is a more than sufficient strategic justification.

A Strategy of Diffusion

A strategy of diffusion lies at the other end of the spectrum. Its supporters view attempts to lock down compute and AI models within a small club of U.S. allies as counterproductive and futile. Instead, this camp argues that the United States should focus on rapidly developing and diffusing U.S. technologies (and governance standards) at home and overseas. Advocates of this strategy strongly support the development, usage, and export of open-source AI systems. They generally oppose the creation of new authorities that might expand the power of the U.S. government to restrict AI-related technological exports, such as controls on proprietary models. And while some analysts within this broad grouping support certain limited export controls on hardware, they generally emphasize the costs and unintended consequences of those measures; they almost all oppose controls on software.

This strategy will be highly influential in certain parts of the Trump administration, especially when it comes to software. The venture capitalist and Trump adviser Marc Andreessen, for example, has called for the free proliferation of U.S. open source models "to drive American and Western AI to absolute global dominance," and Vice President-Elect JD Vance has expressed similar sentiments. Many major technology companies have also lined up behind a version of this approach. Meta, for example, has strongly opposed controls on model weights, as have influential venture capital firms including Andreessen Horowitz. U.S. chipmakers and semiconductor equipment manufacturers, meanwhile, have expressed skepticism, if not outright opposition, to the expansion of U.S. controls on hardware.

Embrace Competition

Of course, some of this opposition is self-serving and familiar from the long history of U.S. export policy. Just as there have always been national security hawks strongly committed to curtailing technology flows, there have also always been U.S. corporations pushing hard to retain access to foreign business opportunities. But proponents of this view argue that a strategy of diffusion can draw on not only relatively narrow commercial considerations but also a broader economic and political logic. In the words of Dean W. Ball, a researcher at the Mercatus Center, "the broad sweep of history suggests that export controls, particularly on AI models themselves, are a losing recipe to maintaining our current leadership status in the field, and may even backfire in unpredictable ways."

Proponents of the strategy of diffusion emphasize, for example, the possibility that U.S. firms, despite their apparent dominance in many aspects of AI, may lack a durable competitive moat. If that proves true, then unilateral U.S. export controls may simply cause American firms to lose market share to foreign competitors. U.S. controls on commercial satellites, for example, may have fueled the growth of rival space industries, eroding America's global market share. Similar dynamics may apply in AI, where a strategy of control could undermine U.S. competitiveness.

American chip design, cloud services, model development, or AI software companies may lose billions of dollars if they are blocked from supplying a broad group of potential trading partners. As in the case of the space industry, loss of foreign sales could also accelerate other countries' development of independent, high-end chip supply chains, as they seek to build an alternative to an industry dominated by the United States. China is already investing billions of dollars in building domestic semiconductor and AI industries; it would gladly forge new business relationships with other states excluded from access to U.S. computing power. Saudi Arabia, for example, which has faced restrictions on its ability to buy advanced chips from the United States, is increasing its joint ventures with Chinese companies like Alibaba and SenseTime, and recently invested $400 million in the Chinese startup Zhipu AI.

Efforts to hoard or steer critical resources like compute within a small handful of wealthy, industrialized democracies may also come with diplomatic costs. Diffusion proponents worry that shutting other states out of key sections of a lucrative value chain could undermine international cooperation on a variety of AI issues. In the past year, for example, the United States has led or supported multilateral efforts on AI safety and security, including corporate commitments and public sector projects to spur research and policy coordination. Countries like the UAE and India are tentatively involved, but this cooperation could quickly evaporate if the United States tries to curtail their access to models and compute—products that they may see as central to their economic and developmental agendas.

Advocates of a strategy of diffusion argue that the United States should instead focus, in the words of analysts Matthew Mittelsteadt and Keegan McBride, on "embracing competition and openness, enabling effective market access, and supporting the diffusion of U.S. AI-enabled technology and governance standards." Washington and Silicon Valley should thus seek not to hoard compute and restrict the dissemination of algorithmic advances and model weights but to accelerate the export and entrenchment of U.S. technology systems in foreign markets as rapidly as possible.

Exploit First-Mover Advantages

The rapid export of AI systems is especially important, in this view, because of two features of the technology, one economic, the other political. From an economic standpoint, the rapid export of U.S. AI technology could prove especially important in markets, such as cloud computing and AI-enabled software applications, in which network effects and high switching costs may generate important first-mover advantages for whichever country enters a particular market first. According to industry sources, this rationale was one justification for Microsoft's partnership with the Emirati company G42, which will help spread U.S. technology in emerging markets where Chinese companies might otherwise entrench themselves.

The United States currently has a window of opportunity, thanks to its dominance of key links in the semiconductor supply chain, but this window could be narrow: While China may be years behind the United States and its allies on leading-edge chip development and

has yet to begin building out global data center infrastructure, major Chinese companies such as Huawei, Biren Technology, and SMIC will continue ramping up chip production. China's appetite for global infrastructure spending has fallen from its peak but may rise again due to sectoral and macroeconomic factors, such as soaring demand for data centers, any future decline in interest rates, or a rebound of the Chinese economy.

Consider the computing infrastructure required for AI inference, which occurs after a model has been trained and allows it to carry out tasks such as answering queries. In some scenarios, inference may not require as many cutting-edge GPUs as a training center for frontier models. China could present a competitive alternative in the near future: If the United States refuses to operate data centers locally in certain countries, China can provide those governments and companies with a more attractive level of service and sovereignty than they are likely to enjoy from remote U.S.-based cloud services. Under the strategy of diffusion, the United States should thus seize the opportunity to establish itself in global AI markets before it is too late.

From a political perspective, meanwhile, supporters of a diffusion-focused export policy, such as Ball, emphasize that "information technologies such as AI are embedded with cultural, political and philosophical values." As a result, the countries that most effectively export AI technologies worldwide will also export those values to billions of people. "Most nations will face a choice in the years ahead between American-built AI embodying concepts of personal privacy, free speech, and intellectual property rights," writes Jacob Helberg, Trump's pick for the State Department's top economic policy and trade official, "and Chinese AI built for surveillance, censorship, and intellectual property theft." The United States must thus disseminate and promote what he calls a "Free World alternative" to Chinese models that are likely to reflect very different perspectives on topics like Taiwan, human rights, and democracy.

Slower Timelines?

One key assumption that often divides advocates of a diffusion-based strategy and a control-based one is the anticipated timing and pace for AI's emergence as a truly transformative technology. Advocates of control often posit that astonishing breakthroughs are just around the corner and could happen quite abruptly, creating a sudden, major gap in technological and strategic power between the first movers and everyone else. If true, then any measure to hold back adversaries and competitors, however briefly, may be worth trying. But advocates of diffusion sometimes take the opposite view: that AI, as a general-purpose technology, will take decades to gradually spread within economies and militaries. In other words, AI is more like electricity than it is the atomic bomb. As the scholars Jeffrey Ding and Allan Dafoe have argued, the effects of general-purpose technologies on the balance of power have historically been "broad, delayed, and shaped by indirect productivity spillovers"—not by monopolizing foundational innovations.

What really matters, in this argument, is not which states get access to which cutting-edge innovation first—the primary focus of a strategy of control—but which states can embrace new technologies at scale, embedding them across a wide range of industries and institutions while broadening innovative capacities such as the available talent pool. Trying to hold back the diffusion of a general-purpose technology through export controls on high-end chips, semiconductor manufacturing equipment, or models is thus most likely futile and counterproductive. By comparison, British attempts to hoard technological secrets during the Industrial Revolution through export controls on textile machinery served primarily to limit the international market for British companies, ultimately handing other countries a competitive edge. So too, in this view, might an emphasis on export controls and mandatory secrecy today serve only to undermine U.S. competitiveness. As Meta CEO Mark Zuckerberg wrote in his manifesto for open-source AI, "constraining American innovation to closed development increases the chance that we don't lead at all."

Finally, advocates of a strategy of diffusion generally emphasize the uncertainties inherent in predicting the future of technological progress. The dominance of the current AI paradigm based on scaling laws, centralized training runs, and a U.S.-dominated semiconductor supply chain may not last (just as other, once-dominant AI paradigms have come and gone), potentially undercutting the value of stringent U.S. attempts to control the technology's diffusion right now. Export controls on compute and model weights may extend the United States' lead marginally in the short run: They can buy time. But critics ask: To what end, and at what cost? Any effect, they say, is likely to be expensive, temporary, and strategically irrelevant. A strategy of diffusion would thus double down on some of the factors that have historically underpinned U.S. technological competitiveness: free markets and open innovation.

A Strategy of Leverage

A third approach stakes out a middle ground. It seeks neither to hoard compute and algorithmic breakthroughs within a small club of U.S. allies, nor to delegate decisions over the trade of AI inputs to the private marketplace. Instead, it would use AI exports as leverage to negotiate specific, country-by-country arrangements that advance U.S. technological and political objectives. For example, Washington could condition the export of U.S. chips on a country's acceptance of robust AI monitoring protocols. Or U.S. officials could even demand that the country endorse some regional peacemaking initiative, new bilateral trade terms, or unrelated sanctions on China. Such bargaining would likely mean that U.S. AI exports wind up falling somewhere in between the tightly constrained strategy of control and the laissez-faire strategy of diffusion. But crucially, American policymakers would extract concrete foreign policy concessions along the way.

Incentivize AI Security Practices

Some argue that Washington should use its AI leverage solely to achieve AI-related outcomes. For example, Lennart Heim (an analyst at RAND), and Cullen O'Keefe (the director of research at the Institute for Law & AI) have advocated that Washington use export licenses and compute deals to help enshrine and enforce AI safety and security norms worldwide. If states wanted access to chips designed by U.S. companies or model weights generated by U.S. labs, they might have to pledge that they will report AI incidents and implement rigorous cyber and physical security measures to protect frontier model weights from exfiltration by hostile actors, state or nonstate. They might have to establish national AI safety institutes that engage with the existing network of AI safety institutes in Canada, Europe, and Japan. They might have to agree that their AI companies will adhere to commitments like those outlined in the G7 Hiroshima Process. They might have to introduce know-your-customer requirements for new data centers to help track potentially dangerous uses of compute and advanced AI chips registries to help ensure chips remain at their intended destinations.

U.S. officials have already begun to explore these opportunities, as indicated by Washington's blessing of the Microsoft–G42 deal, which reportedly includes a variety of security protocols. But the list above illustrates that many different potential conditions are possible. O'Keefe, Heim, and other analysts sympathetic to this approach generally focus primarily on the risks associated with sharing the largest-scale computing resources needed to develop and deploy the most advanced frontier AI systems; they tend to exclude small-scale AI compute and non-AI compute from their export control proposals. And they generally want to use access to U.S. compute as leverage to promote regulations that are, in O'Keefe's words, "narrowly tailored to prevent global catastrophic risks from frontier AI"—not to further U.S. strategic objectives across the board.

Trade Access to U.S. Compute for Concessions

But U.S. leverage could be directed toward more expansive purposes too: The United States could link access to U.S. compute (or other scarce AI inputs and systems) to a wider range of geopolitical concessions. In exchange for access to U.S. compute, for example, U.S. policymakers might ask countries to distance themselves from Huawei or break off joint military exercises with China; reduce bilateral trade imbalances with the United States; support particular U.S. positions in UN votes or other regional diplomatic initiatives; allow U.S. military access, basing, or overflight over sovereign territory; stop supporting armed proxies in civil wars; commit to clean energy and climate decarbonization goals; or crack down on transnational criminal organizations. Again, the list is potentially endless; the overarching theme is leveraging issue linkages between access to compute and other U.S. foreign policy priorities—and doing so now, before U.S. leverage wastes away as Chinese chip production ramps up.

As a middle ground between two other extremes, the strategy of leverage has obvious appeal for U.S. policymakers. In Sullivan's words, it holds out the promise of allowing the United States "to balance protecting cutting-edge AI technologies on the one hand, while also promoting AI technology adoption around the world." This approach also suits Trump's transactional approach to foreign policy and his willingness to draw unconventional linkages in pursuit of new deals. Already, Trump has threatened tariffs against Mexico and Canada unless those governments stop the fentanyl trade, and Vance has said that U.S. involvement in NATO should be conditioned on the European Union loosening what he described as speech-restrictive regulations on U.S. social media platforms.

But making these kinds of compute-for-concessions deals work in practice can be challenging. For one thing, good dealmaking depends on an accurate assessment of U.S. leverage. If Washington underplays its hand, then the United States might give away a major strategic technology in exchange for what will ultimately seem like minor concessions. Conversely, if U.S. officials overestimate the value and scarcity of U.S. AI technology, then the United States might lose out on deals, cede market share and influence to China, and alienate potential partners in the process.

Negotiation also involves delay, yet U.S. companies and many analysts believe that American AI firms are in an intensifying race to entrench themselves globally. Microsoft, for example, has repeatedly expressed its dismay at how long it took to obtain the licenses to ship components needed for its partnership with G42 while Washington and Abu Dhabi negotiated the terms of the deal, even after G42 had reportedly divested from Chinese firms and stripped out its Huawei technology. According to Emirati sources, the Chinese government, for its part, has seized upon these delays as evidence of the United States' arrogance and unreliability as a partner for the UAE. And from the rest of the world's perspective, bespoke deals or country caps that allocate specific numbers of GPUs to each state could become what one industry source called a "trustometer," or an explicit comparative metric of how much the United States trusts each state—a recipe for diplomatic offense.

To be sure, the Trump administration is unlikely to focus on how U.S. computing power can help globalize so-called AI safety norms: The international network of AI safety institutes supported by the Biden administration, for example, is seen as suspect by many (though not all) Republicans. But his administration may well incorporate U.S. technological exports in a broader transactional approach to international politics, potentially using U.S. chips as leverage in pursuit of economic and diplomatic concessions. In the Middle East, for example, if access to U.S. computing power can help Trump secure an Israeli-Saudi normalization deal, he will likely not hesitate for long. In fact, he may not even drive an especially hard deal in negotiations with the Gulf states, given his strong personal relationships with many of their leaders. Nonetheless, the broader approach envisaged by a strategy of leverage is very much in keeping with Trump's self-conception as an international dealmaker.

Conclusion

All three strategies have elements to recommend them; they can all call upon influential and informed defenders on both sides of the aisle. Moreover, champions of all three approaches have much they agree upon. They share a common belief in the strategic and economic importance of AI, and as a result, nearly all analysts and companies—though not necessarily all the relevant political actors—endorse a complementary set of domestic reforms to improve U.S. national AI competitiveness. These include high-skilled immigration reform, expanded STEM education, federal investments in R&D, and permitting reform to unlock clean energy infrastructure.

The Biden administration has seemingly gravitated toward a strategy of leverage, based on the reported contours of the emerging U.S. AI relationship with the UAE and the administration's broader consideration of country-specific compute caps. Under Biden, Washington is moving toward sharing more compute with foreign countries, but under strict conditions designed to promote U.S. conceptions of best AI practices, limit the risks of diversion to U.S. competitors, and roll back Chinese technological influence abroad.

But the debate is far from over, and under the Trump administration, versions of all three approaches will remain on the table for the foreseeable future. Both the executive branch and Congress will be staffed with proponents with varying views. Lobbyists for open-source software, American chip manufacturers, and Gulf authoritarians will confront traditional national security hawks in what are likely to be fluid and unstable contests for influence. The relative strength of advocates of each of these strategies will rise and fall with changes in the domestic political balance of power, as well as with changes in underlying technological possibilities. A stark demonstration of a new AI capability, for example, will strengthen the appeal of the strategy of control. The empowering of political factions around Trump that are dismissive of concerns about the risks of advanced AI and more attuned to the economic interests of U.S. technology companies, meanwhile, would strengthen advocates of a diffusion-focused strategy.

In this new era of AI diplomacy and trade, the question of U.S. priorities and tactics will have enduring relevance in the coming years. It will apply not just to the Gulf monarchies but to a broad range of emerging powers whose interests only partially align with those of the United States—countries like Brazil, India, Indonesia, Nigeria, and Türkiye—and to every layer of the AI stack, from semiconductors and data centers to algorithms and model weights. Export controls will play a major role, but so will capital restrictions, visa rules, and many other tools for shaping the flow of technology goods, services, and inputs. Ultimately, in AI as elsewhere, U.S. policymakers will need to figure out how to reconcile national security pressures with economic constraints; they will need to make decisions that are fundamentally political, involving complex and uncertain trade-offs, amid a fog of technological and geopolitical uncertainty.

About the Author

Sam Winter-Levy is a fellow in the Technology and International Affairs Program at the Carnegie Endowment for International Peace, where his research focuses on emerging technology and national security. Before joining Carnegie, he was a PhD candidate in politics at Princeton University and a Peace Scholar Fellow at the United States Institute of Peace, and he has worked as a staff editor at *Foreign Affairs* and a reporter at the *Economist*.

He has published academic research in the *Journal of Politics*, received Princeton's George Kateb Preceptor Award for teaching, and written for publications including *Foreign Affairs*, *Foreign Policy*, the *New Yorker*, *Lawfare*, *War on the Rocks*, the *Washington Post*, the *Boston Globe*, the *London Review of Books*, and *Scientific American*, among others. He received his undergraduate degree from the University of Oxford and was the Michael Von Clemm Fellow at Harvard in 2014–15.

Carnegie Endowment for International Peace

In a complex, changing, and increasingly contested world, the Carnegie Endowment generates strategic ideas, supports diplomacy, and trains the next generation of international scholar-practitioners to help countries and institutions take on the most difficult global problems and advance peace. With a global network of more than 170 scholars across twenty countries, Carnegie is renowned for its independent analysis of major global problems and understanding of regional contexts.

Technology and International Affairs Program

The Technology and International Affairs Program develops insights to address the governance challenges and large-scale risks of new technologies. Our experts identify actionable best practices and incentives for industry and government leaders on artificial intelligence, cyber threats, cloud security, countering influence operations, reducing the risk of biotechnologies, and ensuring global digital inclusion.

CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

CarnegieEndowment.org