

US-1: Full AI Nationalization can cause Misaligned Economic Incentives

Ryan Tovcimak and Nikolay Radev

Description

The escalating geostrategic importance of frontier AI development increases the likelihood of nationalization. While no explicit plans have emerged in the United States, such action would likely be swift and comprehensive. A government seizure of critical AI infrastructure would fundamentally transform the sector's economic foundation – shifting funding from traditional private sources to the American tax base, thereby repositioning AI as a public good. The objectives driving development would similarly pivot from user engagement to national security imperatives. Given the history of American adversaries pursuing intellectual property theft, this transition would likely establish a more restrictive diffusion model that prioritizes security over openness. By tightly controlling crucial elements of the AI stack, that approach risks diminishing the broader societal benefits that might otherwise emerge from AI advancement.

Introduction

Along many dimensions, AI progress has accelerated at a blinding pace. Even in the brief period since ChatGPT's initial research release, advancements continue at an extraordinary rate. Interestingly, generative AI is one of the few innovations with little to no trace of the US governments' fingerprints. Unlike nuclear energy, semiconductors, computing, GPS, and the internet, AI has been primarily funded and developed by private players – largely hyperscalers and AI labs. As for the latter, the American landscape is dominated by OpenAI, Anthropic, Google DeepMind, Meta, and xAI. To date, the US government has not sat within this loop.

Yet, this rapid evolution has caused the gravitational pull of the AI sector to increase immensely, commanding the full attention of many leaders in Washington. Simultaneously, AI development has become the focal point of US-China competition. In many ways, the international balance of power rests on the outcome of this technological race. In an attempt to guarantee victory on this front, the US government could decide to nationalize AI development.

While various experts disagree on the likelihood of such a prospect, the advantages of a national AI project could be tempting to US leaders. Still, [many recent high-profile essays](#) have offered the contention that this future could be realized. Such an undertaking would not be unprecedented. In the past, the US has seized key industries or launched national scale projects internally. Given the tightly contested competition with China, a nation devoted to strengthening its approach to [Military Civil Fusion](#), the United States could be forced to pursue similar tactics. There are tremendous national security risks associated with both frontier model development and falling behind in this domain. The US government may eventually feel that consolidating control over crucial layers of the AI stack is the best way to quell these threats.

In the present day, there are many emerging motivations that could fuel the US' eventual decision to nationalize AI development. While many of these rationales may emerge later, the most immediately obvious vectors that could pressure such a shift are listed below:

1.) Many US agencies (DOD, NSA, CIA) maintain some of the most rigorous security controls in the world, supported by expertise from organizations like CISA. Comparatively, the security of leading labs is porous. The inadequacy of their respective defensive postures could enable damaging forms of economic espionage perpetuated by key rivals. Above all else, exfiltration remains the largest risk, where an attacker could potentially steal the crown jewel of AI development: the model weights and architecture. Already, [attempts have been](#) made to steal essential algorithmic secrets by those based in adversarial nations. Given the [scale](#) of the [cyber operations](#) deployed by the CCP (largely carried out by the MSS and PLA SSF), prepositioning across each level of the AI infrastructure stack could be well underway. Ceding significant algorithmic efficiency gains or even the end-product of multi-billion dollar training runs works to ensure American defeat. This could potentially force the hand of the US government.

2.) US AI development is running up upon many material bottlenecks. Chiefly, power availability could constrain the attempts of AI labs to reach training runs of sufficient scale. The US' attempts to expand its grid are handcuffed by labor shortages, supply chain woes, lengthy permitting processes, and investment shortfalls. Without sufficient powered land, AI progress could come to a grinding halt. Conversely, it is worth noting that these challenges [are not likely](#) to limit Chinese AI projects. Going further, many other constraints exist in the US, namely those related to training data, potable water, construction capacity, and data center operator talent. Additionally, if Taiwan were taken off the table, the chokepoint could immediately shift to fabrication and packaging capacity. In total, if any of these constraints were to cause significant friction, disaggregation only worsens the issue. Spreading resources thinly across five labs could impede current paradigms by limiting the allocation of compute to different phases of the AI lifecycle. Nationalization then offers an economies of scale approach. If such a project effectively pooled compute, talent, and AI R&D assistant fleets, it could potentially accelerate the advancements of US AI capabilities beyond those achieved by one lab.

3.) As model development intensifies, there have been [many observations](#) of concerning behavior patterns. Deceptive tendencies and scheming have both been widely cited. Meanwhile, although there have been improvements to the adversarial robustness of AI systems, universal jailbreaks continue to be [unturned](#). These episodes further reinforce the consensus that AI alignment remains an increasingly complex undertaking. This becomes a pressing issue given the progress that has been made on potentially harmful fronts (e.g. [virology](#) and [cybersecurity](#)). Today, these trends are precautionary. While many [reported](#) incidents of harm have unfolded, major AI-enabled catastrophes have not yet occurred. Should the consequences of misuse or misalignment worsen, government intervention could be demanded by the public. Nationalizing the development and deployment of powerful AI systems would likely cause the formation of more stringent safety controls. Today, moral hazard dilemmas could cloud the risk management protocols of AI labs. Therefore, it is possible that AI development might need to be offloaded to the government, given its commitment to public welfare.

4.) While some labs have formed partnerships with the US government, including key elements of the defense community, this does not seem to be a focal point of their operations. At the same time, the CCP is placing a major priority on accelerating the integration of championed AI systems like DeepSeek R1 throughout their reach. Given the track record of the CCP on similar endeavors, one cannot underestimate the depth of future points of Military Civil Fusion. Nationalization in the US could eventually be seen as a counterweight to ensure the US does not fall behind in leveraging this technology across critical national security apparatuses. Here, the US' motivations could be heightened by escalating geopolitical tensions or the outbreak of a war.

5.) A financial crisis could impede AI labs from raising the capital needed to achieve next-generation scale. This quandary could be intensified by rising interest rates, a loss of confidence in AI progress, or struggles around commercial viability. In the past, the government has bailed out industrial giants too significant to fail. While labs do not seem to be facing imminent threats to funding, the global economy is delicate. Should turmoil arise, the US government could then be a last line of defense to inject the needed capital. To this point, only some public funding has flowed into the semiconductor industry through different [CHIPS Act vehicles](#). Allocations have not yet been directed towards AI labs (e.g. Stargate does not receive contributions from the US government). Yet, China's recently announced one trillion yuan "[state venture capital guidance fund](#)" could also push America to open its warchest. If aggressive investments are made, ownership stakes could be transferred to the US government, as was the [case with GM](#).

Precedent

In the AI sector, nationalization would be a shocking reversal. Government involvement has been minimal thus far, highlighted by an increasingly deregulatory stance promoted by the current administration. Yet, given the motivating pressures that could arise, such an outcome should not be ruled out. Likewise, there have also been numerous examples of the US government seizing critical industries or launching national projects internally. Dating back to the beginning of the last century, several instances of nationalization have transpired that could mirror nationalization efforts which might occur in the AI sector:

USRA

During World War I, President Woodrow Wilson nationalized the American railway system by establishing the United States Railroad Administration (USRA) through the Federal Possession and Control Act. This unprecedented government takeover was a response to critical transportation congestion that threatened to hamstring America's war effort. At the time, private railroads proved ineffective at coordinating the clash of military and civilian demands. The USRA standardized equipment designs, consolidated facilities, eliminated competitive practices, and coordinated operations across the industry. After the war ended, the railroads were returned to private ownership, though the USRA's standardization left a lasting legacy on the sector.

Manhattan Project

The Manhattan Project represents one of the most significant government-directed research and development initiatives in history. It exemplifies the government's ability to rapidly organize and fund transformative technological development when facing urgent national security imperatives. Rather than nationalizing existing infrastructure, the government created entirely new research and production facilities while coordinating contributions from universities and the private sector. In this pursuit, the project yielded substantial economic impacts. It accelerated technological development in nuclear physics, materials science, and computing with significant post-war civilian applications. This worked to create entirely new industries and capabilities that market forces alone would not have produced in such a compressed time frame. The program also pioneered novel approaches to managing large-scale scientific enterprises that influenced subsequent government R&D initiatives.

While its top-down, mission-oriented approach proved effective for addressing a specific technological challenge with clear parameters and defined outcome, this model may not transfer well to AI development where the requirements and end goals are more unclear.

Apollo Program

NASA and its Apollo program represent another landmark case of government-directed technological development. Established in 1958 following the Soviet Union's Sputnik launch, NASA exemplifies how geopolitical rivalry can drive massive public R&D investments.

The research in [Kantor and Whalley \(2023\)](#) reveals several important insights about NASA's economic impact:

1. Scale and concentration: NASA's budget peaked at over 0.7% of GDP in the mid-1960s, employing over 400,000 workers at the height of the Space Race. Spending was highly concentrated in transportation equipment and electronics sectors.
2. Contractor model: Rather than directly nationalizing industries, NASA utilized an extensive contracting network with private companies. This approach enabled the agency to leverage existing industrial capabilities while directing them toward mission-specific goals.
3. Local economic effects: Regions with pre-existing capabilities in space-relevant technologies experienced significant growth in manufacturing output, employment, and capital accumulation in space-related sectors during and after the Space Race.
4. Fiscal multiplier effects: Kantor and Whalley estimate that NASA contractor spending had a local fiscal multiplier of approximately 0.3-0.4, which is on the lower end compared to typical government spending (0.6-0.8) or military spending (0.3-1.2).
5. Limited technological spillovers: Despite popular narratives about NASA spinoffs, their analysis found limited evidence of substantial local productivity enhancements or technological spillovers beyond directly targeted industries.

Importantly, the research suggests that the "Moonshot" model of government R&D had "only modest economic effects" despite its cultural significance and technological achievements. The space program's impacts were largely concentrated in directly related sectors rather than

generating broader economic growth. This however seems less likely to apply in the AI case, because the result of the project is in itself expected to be a significant economic driver.

While other examples have unfolded in the U.S., other more relevant templates have taken shape abroad:

Japan's Ministry of International Trade and Industry

Japan's Ministry of International Trade and Industry (MITI) represents one of the most studied cases of government coordination of industrial development. Active primarily from the 1950s through the 1980s, MITI helped guide Japan's transformation from a war-devastated economy to the world's second-largest economic power.

Rather than direct nationalization, MITI employed various policy tools to influence industrial development:

- Strategic sector targeting: MITI identified and promoted industries deemed critical for long-term growth, including steel, shipbuilding, automobiles, consumer electronics, and semiconductors.
- Coordinated investment: Through "administrative guidance" and control over foreign exchange, MITI influenced private investment decisions, reduced duplicate investments, and facilitated technology transfer.
- Protection and promotion: MITI employed temporary import restrictions, export promotion, subsidized credit through the Japan Development Bank, and tax incentives to nurture nascent industries.
- Research coordination: MITI sponsored collaborative research projects between companies that might otherwise be competitors, helping to spread R&D costs and accelerate technology adoption.

The economic impacts of MITI's policies remain debated among economists. While Japan experienced extraordinary growth during this period, skeptics argue that MITI may have simply picked industries that would have succeeded anyway. However, Japan's approach offers several lessons. First, the effectiveness of industrial policy may depend on broader institutional factors, including a competent bureaucracy insulated from political pressures. Next, policies that worked during catch-up industrialization may be less effective for frontier innovation. By the 1990s, Japan had significantly reduced MITI's role as the economy matured, suggesting that different approaches may be appropriate at different stages of economic development.

South Korea's Heavy and Chemical Industry (HCI) Drive

South Korea's transformation from an industrial laggard in the 1960s to a manufacturing powerhouse represents another successful case of state-directed economic development. The HCI drive was launched in 1973 in response to security concerns following U.S. policy shifts under the Nixon Doctrine and North Korean military threats. Key aspects included:

- Sector targeting: The government targeted six strategic sectors: steel, nonferrous metals, shipbuilding, machinery, electronics, and petrochemicals based on their upstream position in supply chains and potential for military applications.

- Policy instruments: The government used directed credit through subsidized "policy loans" and tax incentives rather than direct nationalization. By 1973, approximately half of all domestic credit consisted of these policy loans, which had interest rates around five percentage points lower than benchmark loans.
- Economic outcomes: [Lane \(2024\)](#) shows that targeted heavy-chemical industries expanded their output by more than 100% relative to non-treated manufacturing sectors, with labor productivity increasing by over 15%. These effects persisted long after the policy ended.

The Korean case demonstrates that temporary, targeted industrial policy can successfully catalyze long-term industrial development and shift a country's comparative advantage, particularly when security concerns create political consensus for intervention. However, it's important to note that these policies operated within a broader export-oriented strategy and were eventually phased out.

These historical precedents offer several insights relevant to potential nationalization or strategic direction of AI research:

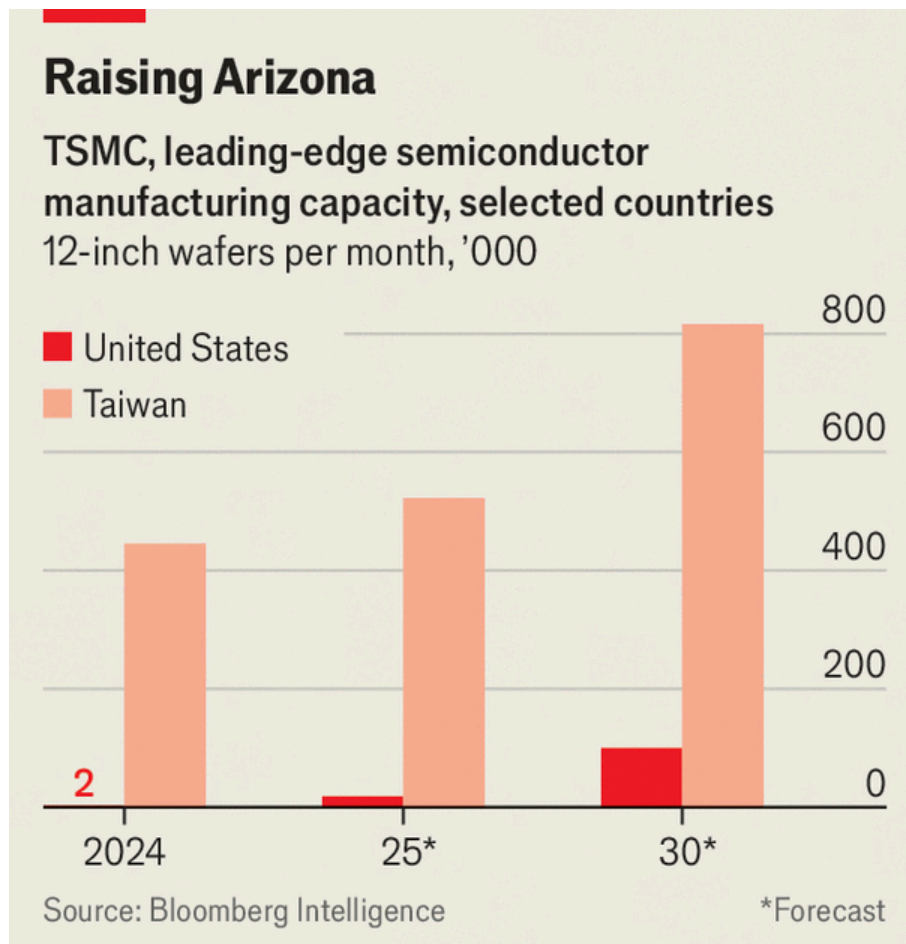
- Implementation matters: The institutional design, governance structures, and specific policy instruments significantly influence outcomes of government intervention.
- Mixed ownership models: Alternatives to full nationalization, such as public-private partnerships, strategic investments, or directed credit may achieve policy goals while maintaining market incentives.
- Time-limited interventions: The most successful cases of industrial policy, like South Korea's HCI drive, were implemented as temporary measures with clear phase-out strategies.
- Security rationales: Successful large-scale government interventions have typically been justified by clear national security imperatives that create political consensus, as seen in the Manhattan Project and South Korea's HCI drive.
- Complementary policies: Successful technology development typically requires complementary investments in human capital, infrastructure, and supportive regulatory frameworks.

Talent

While governmental control can vary in scope, many scenarios imagine full nationalization. In such a scenario, AI development and deployment would be provided by a single, well-coordinated government entity. This organization would have an arsenal of tools to marshal AI infrastructure as needed. AI data centers could be commandeered through DPA Title I authority, where one department might issue DPA priority contracts to secure a dominant share of the US' AI cloud capacity. The government could also take possession of physical data centers through a Declaration of Taking, which relies on the power of eminent domain.

Such an entity may also choose to construct its own data centers, in order to fulfill certain security objectives that may be impossible to accomplish through [retrofitting](#). Once again, the

government can redirect the supply of critical IT infrastructure, cooling solutions, materials, and power infrastructure through DPA authorities. The resulting clusters might be designed with key controls in mind that mitigate exfiltration risks. Securing AI infrastructure might also entail more incentives that push providers to onshore manufacturing in ways that mitigate supply chain attacks. Recent examples of industrial policy, which offer subsidies, tax exemptions, and grants, have proven [moderately effective](#) in guiding manufacturing capacity back to the US. Other punitive levers, like tariffs, are also being explored by the current administration.



Overall, many mechanisms exist that might provide a fully nationalized AI effort the ability to gain more control over elements of the AI value chain. However, levers do not exist that would enable such a program to absorb the AI talent base. A national AI project could only attract willing AI personnel. There are no known legal pathways that the US government could use to mandate key AI personnel to leave the private sector.

Nationalizing frontier AI development in the United States would face significant challenges reshaping the AI labor market. While human capital remains a critical input for AI advancement, federal entities struggle to compete with private-sector compensation packages. Nevertheless, the U.S. government has historically overcome this obstacle, attracting distinguished talent to public service at various stages of their career. Here, evidence has long sat within DARPA, national laboratories, intelligence agencies, and military research programs.

The US government has a deep toolbox that could be used to recruit top AI researchers and engineers. Full nationalization could also lead to more advanced approaches. Below are a list of advantages that might make the US government well-positioned to attract and retain leading AI talent during such a scenario:

Job Security and Stable Access to Funding and Compute

A national AI project could offer stability that might appeal to researchers focused on long-term, high-risk projects that may not yield immediate commercial applications. Since many AI researchers prioritize high-impact exploration over products, this type of program is likely to align well. Additionally, government employment often offers greater job security than private sector positions, which could be attractive during times of heightened turmoil. Finally, joining a national AI program would guarantee researchers access to the largest allocations of compute following any commandeering of cloud capacity. Similarly, joining this effort would provide stable and predictable funding streams detached from venture capital cycles or commercial pressures.

Mission-Driven Work and Public Purpose

Many researchers are motivated by the potential social impact of their work. A nationalized AI initiative could align research with strategic national priorities, potentially attracting talent drawn to working on projects with clear public benefit and national importance. As with previous large-scale government scientific projects like the Manhattan Project or Apollo program, a national initiative could inspire researchers with a sense of purpose and historical significance.

Coordinated National Research Agenda

Under a nationalized approach, the government could implement a unified national research agenda potentially reducing duplication of efforts across multiple private labs and enabling more efficient resource allocation. For researchers frustrated by the fragmented nature of current AI research, this coordination might prove attractive.

Despite these upsides, a national AI program would introduce significant complexities. Public backlash would almost certainly ensue if the government consolidated AI development. Beyond optics, the nature of such an undertaking could pose challenges around talent attraction:

Immigration and Visa Restrictions

Strong security concerns would almost certainly entail stricter immigration controls and security clearance requirements. This approach could severely limit the international talent pool. The history of other nationalized technology sectors suggests that citizenship requirements and security clearances would become standard, potentially excluding many international researchers. This is particularly concerning given [Peterson et al.'s \(2021\)](#) observation that the US has traditionally benefited from "attracting and retaining elite talent, including Chinese nationals," with their research showing that "91% of top Chinese students with US AI doctorates are still in the United States five years after graduating." Today, the breakdown of AI labs is heavily skewed towards foreign born talent, including a large proportion of talent from adversarial nations.

Compensation Disparities

Government salary structures typically cannot match the compensation packages offered by

leading technology companies. The research in [Rock \(2021\)](#) suggests that AI skills command significant premiums in the private sector, with strong correlations between AI talent and market value. Under nationalization, researchers might face substantial income reductions compared to private sector alternatives. While government employment might offer non-monetary benefits like loan forgiveness programs or retirement security, these will be unlikely to offset the compensation gap. Furthermore, public positions are often subject to strict salary tiers that do not reward high performers. This is a sharp break from the bonus packages awarded to top talent at leading labs. These challenges would be particularly acute for senior researchers and engineers who could command seven-figure salaries in the private sector. Consequently, nationalization could spark a significant talent outflow to other technical fields.

Research Autonomy and Culture

A nationalized, security-focused AI initiative would likely impose significant constraints on research autonomy. In other technology areas with national security implications, researchers often face restrictions on publication, international collaboration, and selection of research problems. For many, such constraints could represent a strong deterrent. The culture of government research agencies also differs markedly from both academia and industry. Industry settings typically offer more resources and flexibility than academic environments, while maintaining fewer bureaucratic constraints than government institutions. A nationalized AI research environment might struggle to replicate the innovative cultures that have made leading AI labs attractive to top talent.

International Collaboration Limitations

Security-driven nationalization would likely impose constraints on international research partnerships. [Aaronson \(2024\)](#) cautions about the broader implications of nationalistic AI policies, noting they risk alienating potential allies and could trigger retaliatory policies from other nations. Such dynamics might significantly impede the cross-border exchange of ideas and talent that has historically accelerated scientific advancement. Many AI researchers value international collaborative networks as essential to their work. The implementation of security protocols that limit these global connections could substantially reduce the attractiveness of positions within nationalized AI research environments, potentially causing talented researchers to seek opportunities elsewhere.

To address these challenges, a nationalized AI approach might implement several mitigating strategies:

Special Compensation Authorities

The government could create special compensation authorities for AI researchers, similar to those used at DARPA or for medical professionals in government service. These might include retention bonuses, student loan repayment, or performance-based incentives that help narrow the gap with private sector compensation.

Public-Private Partnerships

Rather than complete nationalization, a "soft nationalization" approach as described by [Cheng and Katzke \(2024\)](#) might better preserve talent attraction capabilities. Their research suggests that the government can increase control over frontier AI labs through an evolving application of US government control without fully nationalizing them. This could include joint ventures,

FFRDC (Federally Funded Research and Development Center) models, or government-commissioned research within private labs.

Dual-Use Research Programs

Creating pathways for researchers to work on dual-use technologies (technologies with both civilian and national security applications) could help maintain a balance between security concerns and research openness. Such programs could allow for classified applications of research while still permitting some publication of fundamental scientific advances.

Moratorium on Frontier AI Development

In a more extreme case, the US government could ban all frontier AI research taking place outside the national AI program. Although coercive, this would funnel all personnel who might be tied to this career path.

Economic Reverberations

A restructuring of this nature would position frontier AI systems as a public good distributed by the US government. This shift would effectively rewire the incentive structures underpinning the modern AI sector. Based on differing priorities between the private sector and US government, this could limit the diffusion of AI capabilities to the broader public.

First, funding sources would be shifted in monumental ways. Today, AI labs predominantly fund their operations through investments from hyperscalers, VC firms, PE firms, and even sovereign wealth funds. AI labs also generate growing pools of revenue through subscriptions, API usage, and other products. Yet, due to the high burn rates of leading AI labs, most of the capital deployed is not sourced from their own earnings.

Full nationalization would largely move this financial burden onto the American tax base. This could add greater fragility to the future financing of AI development. A long-horizon program would likely run into appropriation risks as this shift creates exposure to potential issues related to budget caps, sequestration, or government shutdowns. While AI development presents a strategic priority, political willingness can fade over time. Current investment streams are backed by a steady optimism in this technology, whereas Congress could potentially introduce fickleness, especially as the US pushes towards greater austerity.

At today's scale, funding such a project is not unprecedented. During the peak of [both](#) the Manhattan Project and Apollo Project, the US directed roughly 0.7% GDP to each respective program. Translating this into 2024 terms, this breakdown would result in an allocation of approximately \$203B. Currently, this budget is far larger than the combined outlays of the five major AI labs. However, the capital intensity of AI development continues to rise, driven by cluster scaling, data acquisition, and people costs. The leading labs' expenditures follow very steep curves that significantly outpace GDP growth. For this reason, government appropriations could eventually become a limiting factor if the project was to rely purely on tax revenue.

Leading AI systems are already strong engines of economic growth. If the government nationalized AI development, its distribution efforts could secure additional layers of funding through usage fees. Unlike many previous national projects, serving public usage demand

would need to be a fundamental aspect of the program. In this scenario, AI infrastructure would have to be provided in a way that enabled the national project to approach self-sufficiency.

Yet, self-financed public providers have a history of under-performing. There are a number of theories that can be applied in ways that suggest a national AI deployment strategy could lead to suboptimal performance:

- Budget-Maximizing Bureaucracy (Niskanen, 1971): This theory proposed the concept that bureaucrats act as “budget-maximizers.” These actors instinctually attempt to expand their agency’s budget to gain prestige and power. Under this framework, scenarios revolving around full nationalization could lead to major cost-overruns. Armed with the rhetoric of “AI Dominance,” national AI leaders might pressure Congress for gargantuan compute budgets. Without profit constraints, such a project might also inefficiently allocate compute across model lifecycles. In turn, the US could deprioritize the development of models that can be served in economically meaningful ways.
- X-Inefficiency (Liebenstein, 1966): This theory posits that workers supply less effort in the absence of strong competitive pressure. By removing commercial pressures at the national level, leading AI researchers and engineers could grow complacent. While competition with China is quite potent today, US labs seem more concerned with one upping one another. Their model release schedules and product enhancements all show signs of this pressure. Defusing this form of competition could have severely negative effects on the momentum backing iteration velocity.
- Soft-Budget Constraint (SBC) (Kornai, 1979): In this context, SBC is a concept applied to agencies that operate with the expectation of government bailouts. Oftentimes, risk tolerance can be distorted by leaders' confidence in the government's willingness to cover over budgeted costs. Tying this together, a national AI project might suffer from poor cost-discipline which could test the patience of Congress. While current US labs are not widely heralded for their strong cost controls, these issues could spiral following nationalization. By continuously rationalizing cost overruns, national AI leaders might only accelerate appropriation risks.

More generally, a nationalized program would drastically change the landscape underpinning AI development in the US. AI staff that once prioritized rapid ship cycles would likely instead be pressured to focus solely on absolute capability gains. Progress on chatbot or agent arenas would mean little compared to evaluations related to national security objectives (e.g. cyberdefense, logistic planning, persuasion campaigns, autonomous weapon systems). Directives to iterate in the product space would be replaced with plans to accelerate AI integration into key agencies and other aspects of the defense community.

Conversely, the current cast of American AI labs are well aligned with diffusing this technology as broadly as possible. To do so, the existing landscape largely promotes differentiation over two dimensions. First, price wars have kicked off, intensified by the rise of key Chinese AI offerings, from firms like DeepSeek, Baidu, Alibaba, Tencent, Moonshot, and ByteDance. This encourages research directions that uncover new algorithmic efficiency gains and other

practices that make models more lightweight such as overtraining, distillation, and quantization. Second, firms are gravitating closer to differentiated product offerings, oftentimes split between chatbot platforms and agentic tracks. While many labs may be pursuing several different streams, distinct focuses have emerged within each organization. AI labs are still motivated to maximize overall model capabilities, but may hill-climb on proxy metrics that are more closely related to user engagement (as was the case with Llama 4 and LMArena). Overall, this environment has led to a surge in global adoption rates, with exceptionally strong traction in the US.

Departing from this norm could delay the productivity gains that eventually come from AI progress. Current AI organizations predominantly attempt to maximize the pool of active users and overall token demand. Yet, national AI projects would be instead preoccupied with more geostrategic aspects, especially in achieving advantages over the rival bloc of China, Russia, North Korea, and Iran. This could have a dramatic effect on global adoption curves, especially given concerns surrounding [illicit acts of distillation](#). If leading AI talent is redirected away from improving user, developer, and enterprise experiences, traction in these spaces could slow.

Recommendations

Nationalization should be seen as a last resort. While many onlookers overestimate the US' lead over China, there is still a buffer between the two ecosystems. Due to the drawbacks presented by full nationalization, it is unclear whether such a strategy would extend or evaporate the US' lead. As a consequence of this uncertainty, leaders should take a measured caution when considering this option.

To stave off futures that could warrant nationalization, the US should be taking immediate action on two main fronts. First, the US should be seeking international cooperation that works to limit economic espionage. Similarly, the federal government should push extensive cooperation between key agencies (CISA, CIA, NSA) and the leading labs to harden the defensive posture of these organizations. Moving the major AI labs up the [security levels](#) should be seen as a preeminent national security priority.

Next, the US government should begin injecting considerable funding into Frontier AI Safety research. Following the direction of the CHIPS Act and IRA, the US must incentivize more work on this front. By supercharging research in areas like alignment, interpretability, and control, while simultaneously attracting more third-party safety evaluators, the US can help minimize the probability of AI-enabled catastrophes. Given the potential reach of this technology, AI testing will be a massive market that should be guided by more measurable technical standards. The US should also emphasize investment into vehicles like NIST and the AISI network to ensure proper frameworks exist for third-party entrants to exercise.

Yet, should the US reach a state where nationalization becomes unavoidable, an emphasis should be placed on [softer](#) approaches that do not sacrifice AI diffusion rates. Drawing from the work of [Jeffrey Ding](#), general-purpose technologies reshape global power balances not through

initial breakthroughs, but through the gradual, widespread adoption across the entire economy. If nationalization unfolds, leaders must stress a balance between developing specific strategic capabilities and more general capabilities that boost widespread usage. If America wishes to secure continued prosperity for its people, it must strive to maintain leadership across the AI technology stack that underpins emerging global ecosystems of products, applications, agents, and tools. The economic potential of these developing markets is too significant to neglect.