

Who Will Make Money on AI?

A Discussion Paper on Aligning Commercial Incentives with National Security Interests

Geoffrey Gertz and Emily Kilcrease



Center for a
New American
Security

Center for a New American Security

1701 Pennsylvania Ave NW, Suite 700
Washington, DC 20006

T: 202.457.9400 | F: 202.457.9401 | CNAS.org | [@CNASdc](https://twitter.com/CNASdc)

About the Authors

Geoffrey Gertz is a senior fellow with the Energy, Economics, and Security Program at the Center for a New American Security (CNAS). His work focuses on economic tools for protecting and promoting critical technologies, digital policy and data governance, and geoeconomic competition. Previously, he served as director for international economics for the National Security Council and the National Economic Council, as well as a senior advisor at the State Department's Bureau of Cyberspace and Digital Policy.

Emily Kilcrease is a senior fellow and director of the Energy, Economics, and Security Program at CNAS. Her research focuses on the U.S.-China economic relationship, the alignment of national security objectives and economic policy, and the use of coercive economic statecraft. Previously, she served as the deputy assistant U.S. trade representative for investment and held positions related to economics and national security at the National Security Council and the Department of Commerce.

About the Energy, Economics, and Security Program

The CNAS Energy, Economics, and Security Program explores the changing global marketplace and implications for U.S. national security and foreign policy. In a highly interconnected global financial and trade system, leaders must increasingly leverage economic and financial assets to defend and promote U.S. national interests. The program develops practical strategies to help decision-makers understand, anticipate, and respond to these developments.

Acknowledgments

Liam Epstein provided excellent research support for this project, including support for the research trips to the San Francisco Bay Area and the United Arab Emirates. The authors are grateful to Sam Winter-Levy for peer review of this paper, as well as to Paul Scharre, Janet Egan, Daniel Remler, Pablo Chavez, and Maura McCarthy for helpful comments on earlier drafts, and to the experts who agreed to be interviewed for this research and those who participated in a Washington, D.C., workshop to discuss the paper's findings. Thank you, as always, to the CNAS publications and communications teams. This paper was made possible with the generous support of Coefficient Giving.

As a research and policy institution committed to the highest standards of organizational, intellectual, and personal integrity, CNAS maintains strict intellectual independence and sole editorial direction and control over its ideas, projects, publications, events, and other research activities. CNAS does not take institutional positions on policy issues, and the content of CNAS publications reflects the views of their authors alone. In keeping with its mission and values, CNAS does not engage in lobbying activity and complies fully with all applicable federal, state, and local laws. CNAS will not engage in any representational activities or advocacy on behalf of any entities or interests and, to the extent that the Center accepts funding from non-U.S. sources, its activities will be limited to bona fide scholastic, academic, and research-related activities, consistent with applicable federal law. The Center publicly acknowledges on its [website](#) annually all donors who contribute.

Table of Contents

| | |
|---|-----------|
| EXECUTIVE SUMMARY | 4 |
| INTRODUCTION | 7 |
| COMMERCIALIZATION STRATEGIES OF AI COMPANIES | 12 |
| MARKET STRUCTURE ACROSS THE AI SECTOR | 20 |
| A POLICY FRAMEWORK TO ALIGN COMMERCIAL INCENTIVES AND NATIONAL SECURITY INTERESTS | 27 |
| CONCLUSION | 32 |

Executive Summary

The private sector is playing a leading role in advancing the frontier of artificial intelligence (AI). As a result, commercial incentives are likely to have a significant impact on how AI capabilities develop and diffuse across markets. Firms' commercial incentives will influence U.S. national security interests associated with the emergence of powerful AI systems. These interests include enabling beneficial uses of AI while limiting security risks associated with AI misuse, ensuring reliable and controllable AI system behavior in deployment, and maintaining strategic geopolitical advantage in the development and global diffusion of AI.

Yet to date, stakeholders focused on AI national security interests have paid only limited attention to AI companies' commercialization strategies and market dynamics across the AI stack. This paper seeks to bridge this gap, identifying potential scenarios for the future shape of AI markets and exploring the implications of these scenarios for U.S. national security. Rather than attempting to resolve core debates on the commercialization of AI, the paper seeks to prompt consideration in both the private and public sectors, and among economics and national security expert communities, of how commercial incentives can better align with U.S. national security interests.

Across the AI technology stack (including the infrastructure layer of chips and data centers, the foundation model layer, and the application layer), firms are exploring how to achieve profitability. Variations in companies' commercialization strategies may result in either more or less demand for AI safety and security from product end users, which may impact how these strategies do or do not reinforce U.S. national security interests. At the sector level, the concentration of market power across layers of the AI stack has implications for the government's ability to effectively guide the private sector-led development of powerful AI. Mapping out how such commercial dynamics may unfold will help prepare policymakers and other stakeholders to identify possible market interventions to better align commercial incentives with U.S. national security interests.

Key Takeaways

The following key takeaways emerged from the research:

Commercialization Strategies of AI Companies

- AI model and application companies are increasingly focused on generating revenues, targeting four potential end markets: enterprise use cases, consumer use cases, government use cases, or leveraging AI capabilities for internal use cases to indirectly earn revenues. The relative size and appeal of these four markets will shape the trajectory of the AI ecosystem.
- Different end markets have varying levels of demand for AI safety and security based on customer incentives. Government and enterprise markets will tend to be more risk averse, which may better align with U.S. national security interests. Consumer end markets may feature significant information asymmetries and risk externalities. Internal use cases may hinder transparency and create misuse and misalignment risks that are difficult for governments to manage.
- Compute policy—both restricting trade in advanced semiconductors through export control programs and supporting such trade through export promotion programs—will continue to play an important role in advancing U.S. national security interests on AI. U.S. strategy will need to balance the objectives of maintaining a preponderance of advanced compute in the United States with encouraging foreign counterparts to build and deploy their AI systems on top of U.S. infrastructure.

Market Structure Across the AI Sector

- The AI market has meaningful concentration at the infrastructure and model layers of the AI stack, which may be exacerbated by vertical integration and circular financing deals among key players.
- There is significant uncertainty on how the foundation models market will evolve, including how market demand may be segmented between open-weight and closed-weight models and the extent to which “good enough” commoditized AI will capture a substantial share of the total market.
- Shifts toward a monopolistic market structure at any layer of the AI stack may erode the government’s power to protect U.S. national security interests if any one AI provider becomes too big to fail. At the same time, a more decentralized market may present different challenges associated with difficulties in controlling access to powerful AI and coordinating on best practices for AI safety.

Recommendations

Policymakers face a formidable challenge in trying to shape AI markets to align commercial incentives and national security interests. Given the rapid pace of technological development, rigid rules are unlikely to be effective. Instead, a flexible and layered policy framework that prioritizes shaping market incentives so that economic actors internalize the costs and benefits of AI safety and security may be a more durable approach. Such a framework should:

1. Leverage U.S. infrastructure and compute capabilities to shape the geographic distribution of AI.

U.S. leadership in AI infrastructure and compute provides the United States with substantial ability to shape the geographic distribution of AI capabilities and adoption, as well as an opportunity to make access to U.S. technologies contingent upon alignment with U.S. national security interests. U.S. policymakers should continue to facilitate the build-out of AI infrastructure in the United States, use export controls and export promotion policies to limit China’s ability to develop and deploy frontier AI while encouraging other foreign counterparts to buy U.S. AI products where feasible, and embed basic AI safety and security conditions into AI infrastructure approvals.

2. Use transparency, liability, insurance, and procurement policies to align market demand for AI safety and security with commercial incentives.

Policymakers should prioritize strengthening the transparency of AI companies, which will better enable investor and customer demand for AI safety to influence AI companies’ behaviors. Policymakers should also mandate disclosures to the government for certain high-severity misuse and misalignment risks, particularly for companies’ internal use of models and applications which may not otherwise garner public scrutiny. Clarifying the legal liability for AI misuse and misalignment incidents up front through legislation, rather than litigating through the courts, could also help ensure that the costs of misuse and misalignment do not become national security externalities in AI markets. Supporting the emergence of AI insurance markets can also incentivize more responsible behavior. Finally, governments should leverage their role as customers of AI products to send strong demand signals for safe and secure AI.

3. Encourage safe and competitive AI markets in the United States.

Policymakers should provide antitrust guidance reflecting the economic and national security importance of AI markets, defining acceptable levels of concentration across the AI stack, and clarifying

under what conditions AI companies can coordinate and share information to advance their shared AI safety priorities. The government should also assess national security risks associated with the use of Chinese models in the U.S. AI market and adopt risk-based approaches to limit the access of Chinese firms where justified. Recognizing considerable demand for open-weight models and the downsides of ceding this market to Chinese competitors, policymakers should work with researchers and the private sector to foster a healthy ecosystem of safe and secure U.S. open-weight models.



Introduction

The development of breakthrough dual-use technologies has historically involved cooperation and interaction between academia, the government, and the private sector. Researchers in academia often provide the initial seeds for new innovations, governments can marshal resources and direct research programs toward national security priorities, and as technologies mature and develop, the private sector can scale new capabilities, operating within regulatory frameworks that account for dual-use risks identified by the government. For example, the initial ideas for nuclear energy emerged from European academics. The U.S. government then launched the massive Manhattan Project to develop nuclear weapons capabilities in the context of the Second World War. A decade later, the Atomic Energy Act of 1954 paved the way for the private sector to deploy civilian nuclear power plants.¹

The development of artificial intelligence (AI), however, has followed a distinct trajectory. The private sector is playing a leading role in advancing the frontier of this new dual-use technology even in its early stages, as novel capabilities and associated risks are still being discovered. Because AI is developing within commercial firms rather than government programs, market forces are likely to play a significant role in shaping how these capabilities develop (as companies direct research resources to align with commercial incentives) and diffuse (as companies choose which markets to prioritize). There is an urgent need to better understand how market structures and commercial pressures will influence the evolution of AI technologies, and what these dynamics mean for AI risks and U.S. national security interests.

Today, there are several wide-ranging debates on future scenarios for AI commercialization and market development. The unique characteristics of AI as a product, including significant uncertainties on what product fits will ultimately be most important for a powerful general-purpose technology, make it difficult to draw clear historical lessons from other industries. Will companies developing frontier AI models identify sustainable business models? Will enterprise AI users find productive use cases to justify rapid uptake of AI tools? What layers of the AI stack will be most profitable, and what layers may become commoditized? Will today's dominant incumbent companies across the AI stack maintain their positions, or will new upstarts encroach on their markets?

The starting point of this paper is that the answers to these and related questions will have important implications for U.S. national security interests. To date, however, experts focused on AI and national security have tended to pay more attention to technical AI developments, such as the trajectories and timelines to reach different AI capability benchmarks, than to companies' commercialization strategies and market incentives. To the extent that commercial dynamics are considered, it has been primarily around how competition to release new models may give an incentive to cut corners on safety.² While important, this is only one channel through which AI commercialization may impact national security, leading to an important gap in understanding. This paper does not seek to provide definitive answers to the big outstanding questions on the evolution of the market for AI, but rather to better understand the contours of such debates to identify what is at stake for U.S. national security and assess how policymakers and private actors might better align commercial incentives with national security interests.

The AI Economy

AI has evolved from an esoteric research topic to an emergent general-purpose technology that could have transformative effects on the U.S. and global economies. Researchers are beginning to identify the impact of AI in the latest aggregate economic data on productivity and employment in the United States, although such analyses remain contested.³ Expectations of the future economic impact of powerful AI vary widely. Some leaders at frontier labs have suggested that AI will drive an economic revolution in the near term, with growth in gross domestic product accelerating to 10–20 percent a year.⁴ Economists studying AI tend to be more cautious, highlighting that even as AI capabilities improve exponentially, economic growth will still be limited by physical constraints in the real world—including the ability to

scale chips and data centers—and that historical precedents suggest technology diffusion and adoption is a slow and staggered process.⁵ Some economists and AI experts warn powerful AI could spark mass unemployment, as AI becomes a cheaper and more capable substitute for nearly all knowledge workers.⁶ Other economists, however, are more sanguine, again often pointing to the history of earlier labor-saving innovations, which served as a complement rather than a substitute for labor and induced greater aggregate demand.⁷

What Is Powerful AI?

Several terms are commonly used to describe advanced AI systems. **Frontier AI** refers to the most capable general-purpose AI models available at a given time. **Artificial general intelligence (AGI)** refers to an AI system that can match or exceed human-level performance across a wide range of cognitive tasks. **Artificial superintelligence (ASI)** describes a system that far surpasses human capabilities in virtually all domains.⁸ **Transformative AI** refers to AI that could precipitate societal change on the scale of the Agricultural or Industrial Revolution, emphasizing impact rather than capability.⁹

AGI and ASI are defined by what a system can do; transformative AI is defined by what a system causes. An AI system could be transformative without reaching AGI, for instance, by automating enough economically valuable tasks to reshape labor markets. Similarly, an AI system could reach AGI-level capabilities without proving transformative if bottlenecks in deployment or adoption limit its real-world impact.¹⁰

This paper uses the term **powerful AI** to refer to AI systems with advanced cognitive capabilities, including the ability to reason, act autonomously, and perform tasks across a range of domains. Unlike frontier AI, which describes a relative position (the most capable systems at any given time), powerful AI describes a level of capability that carries national security implications regardless of whether a system is at the frontier. Like AGI and ASI, powerful AI is defined by capabilities without assuming specific impacts of those capabilities, which is most relevant for the analysis in this paper given its focus on interrogating the extent to which powerful AI may, or may not, lead to economic transformation.

This paper is not primarily concerned with debates on timelines for reaching powerful AI (or other AI thresholds) or understanding the technical advances needed to reach such capabilities. Instead, it adopts the simplifying assumption that leading U.S. AI systems will be capable of producing powerful AI in the medium term (within the next 3–5 years), which is well within the expectations of experts at the leading labs themselves.¹¹ The national security interests examined in this paper do not depend on resolving these debates, and these interests are implicated across a range of capability levels, including those available today.

Meanwhile, the American public is increasingly anxious about AI. One recent poll found that 57 percent of Americans believed the risks of AI outweigh its benefits, with only 34 percent holding the opposite opinion.¹² Seventy percent of Americans think advancements in AI will lead to a decrease in job opportunities, while only 7 percent expect an increase.¹³ Such perceptions are sparking a backlash against the prospect of transformative AI, epitomized in the fights state and local governments are launching against new data center construction.¹⁴ Whether these public pressures ultimately translate into policy shifts to restrict AI development and deployment, however, remains to be seen; Americans have also had predominantly negative views of social media companies for several years, but this has not resulted in substantial regulatory changes to date.¹⁵

Against this backdrop, companies across the AI stack are seeking to establish or cement profitable business strategies in rapidly evolving commercial AI markets. Individual companies' commercial decisions are shaped by broader macroeconomic trends in the AI market, and likewise the aggregate economic impact of AI will be determined by the interaction of individual companies' decisions. Certain segments of the AI tech stack, such as semiconductor companies, have more clarity on commercialization and have already reaped substantial revenues from the AI boom. However, fundamental questions remain about the path to commercialization and profitability for AI models and applications. A central question will be to what extent AI companies are incentivized to design their products for enterprise, consumer, government, or internal (e.g., indirectly monetizing AI capabilities by better targeting ads) uses. How companies balance between these market segments, and which grow faster than others, will drive commercial strategy decisions that will shape incentives for investing in AI safety, potentially leading to knock-on effects for U.S. national security interests.

Companies' commercialization strategies will also impact market structure and concentration across the key layers of the AI stack: AI infrastructure, including chips and data centers; foundation models; and AI applications. Throughout the AI stack, companies are seeking to establish market positions that provide a moat to shield them from competition while also worrying about overconcentration among both their suppliers and their customers. These evolving market dynamics will influence what levers governments might use to shape AI development and deployment to ensure its consistency with U.S. national security interests.

The Layers of the AI Stack

This paper analyzes AI commercialization and market structure across three layers of the AI technology stack: infrastructure, foundation models, and applications.

The **infrastructure** layer encompasses the physical and computational resources underpinning AI development, including the advanced semiconductors used to train and run AI systems, the data centers that house and operate them, and the cloud computing platforms that make computing power available to customers.

The **foundation model** layer refers to large-scale AI models that are trained on broad datasets and can be adapted for a wide range of tasks.

The **application** layer comprises the products and services built on top of foundation models that deliver AI capabilities to end users.

U.S. National Security Interests

Powerful AI capabilities may bring significant benefits for, but may also have serious repercussions on, U.S. national security interests. This paper considers three categories of national security interests: (1) how to enable beneficial AI uses while minimizing risks associated with the misuse of powerful AI, (2) how to ensure reliable and controllable AI system behavior, aligned with the intentions and values of their developers and users, and (3) how to maintain the United States' strategic geopolitical advantage in the development and global diffusion of AI.

Realizing the benefits of powerful AI will require broad adoption and diffusion of AI capabilities. But expanding the use of AI may also increase security risks associated with the *misuse* of AI, namely the ability of malicious actors to deliberately employ powerful AI for harmful purposes. As AI systems become more capable, they can lower the barriers to conducting cyberattacks, developing biological or chemical weapons, generating disinformation at scale, or supporting adversary military operations. The April 2026 announcement by Anthropic that its Mythos model had discovered a dangerous number of

zero-day cybersecurity exploits, and its decision to withhold a public release of the model, underscores that these risks are no longer hypothetical.¹⁶ AI developers build safeguards into their models to prevent misuse, but these guardrails are imperfect, susceptible to circumvention through jailbreaking, and require continual investment to maintain as model capabilities advance. They are also discretionary, left up to individual firm decisions. Moreover, the growing availability of open-weight models, where model parameters are publicly released, means that safety guardrails can be removed by anyone with sufficient technical capability.

The United States also has a national security interest in ensuring AI systems behave in ways intended by their developers or users and avoiding AI misalignment. At one end of the spectrum, misalignment includes relatively mundane technical failures, such as an AI system that misinterprets instructions, produces unreliable outputs, or acts unpredictably in novel situations. At the other end, it encompasses more severe scenarios in which highly autonomous AI systems pursue objectives that diverge from human intentions in ways that are difficult to detect or correct. At an extreme, misalignment of powerful AI systems may create large-scale risks to human safety. AI developers use approaches such as reinforcement learning from human feedback and related techniques to train models to better align with what users want, but “interpretability”—understanding why a model acts the way it does—remains a significant challenge, particularly as systems grow more capable and are deployed with greater autonomy.¹⁷

Finally, the United States also has a set of geopolitical national security interests associated with how AI development and diffusion may shift the balance of power between states. The United States has a strategic interest in maintaining a leading position in AI relative to competitors, particularly China, given AI’s potential military and geoeconomic implications. AI systems are already being deployed in U.S. military operations, and a dominant AI capability could substantially alter the balance of power between the United States and its adversaries. Beyond direct military applications, if U.S. competitors gain a large share of the global AI market, this could create systemic dependencies and provide them with a powerful tool for economic leverage over other countries.

Importantly, these three interests are not independent and may at times be in tension. For example, the geopolitical interest in maintaining more advanced AI capabilities than strategic competitors may lead to dangerous AI race dynamics, where concerns about misuse and misalignment are de-prioritized to develop powerful AI capabilities as quickly as possible. The reverse is also true: Pausing U.S. AI progress until misuse and misalignment risks are fully resolved could compromise the country’s competitive position relative to adversaries who do not apply the same constraints. The overarching national security interest for the United States is advancing each of these three objectives in tandem, without sacrificing any one for the sake of another.

About This Paper

This paper does not seek to resolve any particular debate about how AI markets will evolve, nor to make direct causal arguments that certain market structures will be better or worse for U.S. national security interests. In fact, in most scenarios, there will likely be a mix of competing incentives as firms experiment with different ways to achieve profitability. By charting a range of plausible commercial scenarios and then mapping out how these scenarios may implicate U.S. national security interests in a variety of ways, the goal of the paper is to inform decision-makers in the private and public sector as they seek to align commercial incentives with national security interests. A richer understanding of these market dynamics will help U.S. policymakers and private sector actors seeking to shape these markets to appropriately calibrate their interventions to avoid unintended consequences.

The analysis and recommendations in this paper are informed by extensive desk research by the authors, as well as dozens of interviews conducted with industry, civil society, and government

representatives in the San Francisco Bay Area, Washington, D.C., and the United Arab Emirates. Interview subjects included companies making AI chips, building AI data centers, and developing foundation models and AI applications; enterprises integrating AI into their business operations; venture capital firms and other firms investing in AI; AI safety experts; academic economists; and government officials regulating AI and developing strategies to integrate AI into their economies. Interviews were conducted confidentially to allow participants to share candid views.

The next section examines potential variations in AI companies' commercialization strategies and assesses to what extent these commercial incentives align with U.S. national security interests. The subsequent section zooms out from the company level to the sector level, assessing how market structure and concentration may evolve across layers of the AI stack and the national security implications of these trends. The final section outlines a high-level framework for strengthening policy interventions to better align commercial incentives with national security interests.

Commercialization Strategies of AI Companies

Several years into the AI era, how AI companies will make money is only now coming into focus. Companies' corporate structures initially provided some insulation from such market pressures: Both Anthropic and OpenAI are organized as public benefit corporations with missions beyond profit maximization, while other key players in the AI market, such as Google and Meta, have massive non-AI revenue streams which could fund their AI development.¹⁸ In 2019, OpenAI's CEO Sam Altman famously remarked in an interview that the company had "no current plans to make revenue [and] no idea how we may one day generate revenue. ... Once we've built this sort of generally intelligent system, basically, we will ask it to figure out a way to generate an investment return."¹⁹

As of April 2026, however, as OpenAI and Anthropic both prepare for initial public offerings (IPOs) later this year which are anticipated to be among the world's largest public listings of all time, the question of where sustainable profits will come from has gained more urgency.²⁰ Though semiconductor companies are already getting rich off the AI boom, the long-term commercial sustainability of the broader AI market will ultimately depend on realizing profits from the final end users of AI, which will require determining the market fit for AI models and applications.

Key Takeaways

- AI model and application companies are increasingly focused on generating revenues, targeting four potential end markets: enterprise use cases, consumer use cases, government use cases, or leveraging AI capabilities for internal use cases to indirectly earn revenues. The relative size and appeal of these four markets will shape the trajectory of the AI ecosystem.
- Different end markets have varying levels of demand for AI safety and security based on customer incentives. Government and enterprise markets will tend to be more risk averse, which may better align with U.S. national security interests. Consumer end markets may feature significant information asymmetries and risk externalities. Internal use cases may hinder transparency and create misuse and misalignment risks that are difficult for governments to manage.
- Compute policy—both restricting trade in advanced semiconductors through export control programs and supporting such trade through export promotion programs—will continue to play an important role in advancing U.S. national security interests on AI. U.S. strategy will need to balance the objectives of maintaining a preponderance of advanced compute in the United States with encouraging foreign counterparts to build and deploy their AI systems on top of U.S. infrastructure.

Commercialization Across the AI Stack

At the infrastructure level of the AI stack, commercialization strategies have been clear for some time. Semiconductor companies have made billions of dollars amid surging demand for AI chips, particularly as the most sophisticated AI systems rely on ever-increasing amounts of compute to achieve technological advances. A handful of U.S. hyperscalers account for the majority of advanced AI chip purchases, some of which are then rented to frontier labs for model training and inference.²¹ Data center company valuations have jumped as their capacity is booked for years in advance.²² High expectations for the transformative potential of powerful AI have supported the commercial rationale for massive sales of AI chips and data center build-out.

One important commercial consideration for the build-out of AI infrastructure is geographic location. Several factors, on both the supply and demand sides, influence the location of AI data centers. Data centers require massive energy to run, and thus data center companies look to build where energy is cheap.²³ The regulatory environment also matters, both general permissiveness for permitting and

building AI data centers as well as copyright restrictions relevant for training models on datasets of published works. Several governments have prioritized developing local AI infrastructure as part of sovereign AI initiatives, allowing greater control over AI development and deployment and reducing foreign dependencies.²⁴ There are also technical considerations: Data centers that have strong submarine cable connections or proximity to end users have lower latency levels and thus faster inference capabilities. Collectively, these factors are likely to encourage at least some geographic dispersion of AI data centers, though at least in the near term, it appears likely that only a handful of jurisdictions may host the biggest data centers used for training frontier models.

While money is currently pouring into chips and data centers, the long-term sustainability of their business plans is still open for debate. One key question relates to how technological breakthroughs allowing for less compute-intensive AI development will ultimately impact demand for AI infrastructure. In early 2025, when the Chinese company DeepSeek released a new, highly efficient model suggesting future generations of frontier AI models might require significantly less compute, NVIDIA's shares initially plunged on the expectation that demand for AI chips might slow in the future. However, the stock subsequently rebounded, as market sentiment shifted to suggest that more compute-efficient AI could actually increase, not decrease, demand for chips, as this would open new use cases and increase adoption, an economic phenomenon known as Jevons paradox.²⁵ Whether the Jevons paradox dynamic holds over the long term for demand for AI chips, however, is unclear. Another question relates to how quickly the useful shelf-life of an AI chip depreciates, which *The Economist* has called "the \$4 trillion accounting puzzle at the heart of the AI cloud."²⁶ Since leading chip designers release new cutting-edge semiconductors every year or so, chips installed in data centers last year are quickly no longer state-of-the-art (though they may still be useful for various inference tasks); failing to account for this rapid depreciation may lead to a significant overvaluation of data center assets. More generally, worries of an "AI bubble" persist, and if demand for AI collapses, then demand for AI infrastructure would as well. Stock markets punished the major U.S. hyperscalers earlier this year after they reported a larger-than-expected \$660 billion in planned capital expenditures for AI data centers, perhaps a warning sign that AI infrastructure commitments are outpacing commercial realities.²⁷

AI commercialization strategies at the model and application layers are less clear than at the infrastructure layer. At the model layer, there are significant questions about how to make the basic economics of developing frontier AI models work. Recent analyses suggest that current foundation models are likely not profitable once one accounts for both the research and development (R&D) costs associated with developing and training the model as well as the inference costs associated with running the model.²⁸ While revenues at frontier AI labs have been growing rapidly, the companies are incurring massive outlays for R&D and compute necessary to train new models, which are often proprietary or closed-weight models (i.e., the training parameters are not disclosed). Once the model is trained, its monetizable lifespan is typically short: Within a few months of the release of a new best-in-class model, there will likely be a newer, more capable closed-weight model release that further advances the frontier, as well as the release of a freely available open-weight model with comparable performance as the original closed-weight model.²⁹ Both dynamics will limit the durability of pricing margins for any given closed-weight model; each generation of new models builds on the previous one, but also erodes its predecessor's market appeal. Recent research demonstrates the nominal price users pay for foundation model inference has held relatively steady in recent years, despite a dramatic increase in model quality, implying that the quality-adjusted price of AI inference has dropped precipitously.³⁰

Moving to the application layer, the biggest uncertainties on commercialization strategies relate to identifying product fit. Notably, the success of some recent AI applications that achieved broad market penetration, including OpenAI's ChatGPT chatbot and Anthropic's Claude Code, apparently came as a surprise to their companies.³¹ Other AI products that initially became viral hits, like OpenAI's Sora, an app allowing consumers to create their own AI-generated videos, ultimately flopped on the market as the cost of compute needed to run these services outpaced their ability to earn revenue.³² Such

uncertainty reflects a common feature of technology start-ups, where product fit may not be obvious *ex ante* and is only identified through iterative processes of testing new products on the market.

Of course, a current lack of profitability does not imply AI models or applications will never be profitable. Investor documents prepared for upcoming Anthropic and OpenAI IPOs project the companies to become profitable in 2028 and 2030, respectively.³³ Many new technologies initially operate at a loss before ultimately maturing into stable, profitable businesses. For now, U.S. and global capital markets, including venture capital, private equity, and sovereign wealth funds, have demonstrated a willingness to continue deploying billions of dollars to AI companies in the expectation that such bets will one day generate substantial returns. These companies' ability to ultimately develop sustainable business models will depend, in part, on how they prioritize among different types of customers and end markets.

End Markets for AI Foundation Models and Applications

At a general level, there are four categories of business use cases that AI companies might choose to focus on: providing products for enterprises, providing products for consumers, providing products for governments, or leveraging AI capabilities for internal use (which includes both companies using AI to further advance their next generation of AI capabilities, as well as companies using internal AI capabilities to indirectly make money in other markets through better targeting ads, picking stocks, designing pharmaceuticals, or other business lines). These four pathways to commercialization are not mutually exclusive: To date, leading labs have pursued versions of all four, while companies focused primarily on the application layer are more likely to tailor products to specific markets.

The relative size and appeal of these four markets will shape the trajectory of the AI ecosystem. AI companies will shift their practices and technological development plans to align with the demands and preferences of the customers they intend to serve. Moreover, the companies tailoring their products to markets that are ultimately more lucrative will subsequently gain larger revenues, build their market share, and thereby play a larger role in shaping the trajectory of the field.

Enterprise

AI insiders generally expect the enterprise market will become the largest revenue-generator for AI models and applications.³⁴ While data is limited, investor documents show Anthropic earns a substantial majority of its revenue from enterprise users and expects to continue to do so in the coming years.³⁵ OpenAI's current revenue base is more tilted toward consumers, but enterprise revenues are projected to grow more rapidly and account for around 40 percent of total revenues by 2030.³⁶ Current enterprise AI use cases focus mostly on using AI tools to improve software engineering productivity as well as automating back-office functions such as information technology (IT) and service operations. Looking forward, the biggest productivity gains and most transformative impacts could come from incorporating AI into manufacturing processes or other production tasks. Business cases built around enterprise users will tend to focus on higher price points and more tailored product offerings, earning revenues through high annual recurring subscriptions. Enterprise users will look for AI that can improve their productivity and/or lower their costs (including labor costs, where tasks previously carried out by humans can now be done by AI agents). AI products might conceivably replace many of today's business-to-business software-as-a-service (SaaS) offerings.³⁷

Consumer

The first AI product to break through into the broad public consciousness was the consumer-focused chatbot ChatGPT from OpenAI. Products for consumer markets will typically earn revenues through targeted advertising, small subscription fees, affiliate shopping links, or selling data. Many of today's largest tech and social media companies have attained large valuations on similar business cases, often

providing products to end users at a zero price point in order to build massive user bases that can then be monetized.³⁸ Notably, several companies that have large existing consumer platforms, including Google, Meta, and X, are also developing AI models and applications. The ability to leverage these existing customer bases may be a valuable asset in growing consumer-focused AI products.

Products targeted toward consumers will focus on maximizing users and engagement, as this is what will drive revenues. Consumer-focused AI products can be designed for broad general audiences (like ChatGPT) or for more defined niche audiences, such as tools specifically designed for health and wellness, personal finance, shopping, or education. In addition to tools to accomplish specific tasks, the consumer market also includes AI companions, which offer users an ongoing emotional and social relationship. Pornography and erotica could also emerge as a highly lucrative consumer market, based on historical examples in the development of the home video market and internet.³⁹ OpenAI reportedly developed plans to launch an “adult mode” for ChatGPT, but put such plans on hold indefinitely in March 2026.⁴⁰

Government

Governments (at the local, state, and federal levels) procure AI for a wide variety of use cases, typically related to improving the productivity of their own internal processes or enhancing citizen services. The size of the procurement market will vary country-by-country, based on the extent of the government’s involvement in the economy as well as how technology-forward governments are. Government use cases are generally similar to enterprise uses, but with two important differences.

First, the government is fundamentally a distinct actor in that it possesses a monopoly on the use of force, including military and law enforcement powers. When AI is deployed in support of these powers, the stakes are inherently higher, as the costs of any AI system mistakes could be lethal. Additionally, powerful AI systems used for military and law enforcement purposes may enable government intrusion into citizen privacy in more pervasive ways than ever before, with potentially serious implications for human rights and democratic institutions.

Second, many foreign government markets are also particularly shaped by strong demand for “sovereign AI” capabilities, which prioritize building local AI infrastructure and having full control over foundation models and applications, preferably ones tailored for the local market. A key driver for sovereign AI is the concern in certain foreign markets about a continued reliance on the dominant U.S. AI providers, along with an overall perception of rising levels of geopolitical risk. Sovereign AI can also reflect an industrial policy strategy to capture a greater share of the AI value chain and to build domestic capabilities that will have positive spillovers on other industries. There is considerable variation in sovereign AI approaches, including more narrow efforts to protect certain sensitive data while still relying on foreign compute and models, to complete full-stack sovereignty, though the latter is exceptionally rare in practice.⁴¹

Internal

Finally, companies creating frontier AI models might choose to keep these capabilities internal to the firm rather than sell them externally. Most obviously, frontier labs are already using their own models to write the code for their next models (though again, how these newest models will ultimately generate revenues remains an open question). Companies could also apply their own AI models to other ancillary revenue generation activities—Meta and Google, for example, could use proprietary AI models to better target their advertising, while Microsoft could incorporate AI capabilities into its existing suite of software applications. And Altman’s original idea—to create a superintelligent system and then ask it how to make profits and do that—could also be implemented as a closed business strategy; a frontier AI lab that identified an AI model for successfully picking stocks might reasonably decide the best

commercialization strategy is to carry out the financial trades the model suggests, rather than sell access to the model. DeepSeek, the Chinese company that shocked the AI industry when it demonstrated capabilities close behind those of the leading U.S. labs, was initially developed as a side project for a quantitative trading firm using AI for algorithmic trading.⁴² It is plausible that companies in biotech, robotics, or other sectors may also seek to develop their own tailored, niche, powerful AI systems for their own internal uses in developing new products.

How Commercial Incentives Align with National Security Interests

Companies' commercialization strategies across the AI tech stack will implicate U.S. national security interests and present both risks and opportunities for U.S. government interventions to shape markets.

Infrastructure Layer

At the chip and infrastructure level, commercial incentives are driven by rapidly growing global demand, including demand outside of the United States. Companies at this layer are seeking to expand their global footprint and break into new markets to meet this demand. This commercial drive for global expansion of AI infrastructure can cut against the U.S. interest in retaining AI global leadership, as infrastructure outside of the United States can enable emerging AI ecosystems to develop competitive capabilities to challenge the U.S. position. Among other markets, the lucrative Chinese market will appeal to chip companies, given its massive size and its government's commitment to advancing AI capabilities. U.S. export control policies have reflected this geopolitical interest, restricting China's access to advanced AI chips developed in the United States.

Yet a policy to retain maximal AI infrastructure in the United States would be in tension with another objective: As third countries (other than China) build out their own AI infrastructure, it will benefit the United States if they do so by partnering with U.S. and other like-minded companies, rather than with emerging Chinese competitors. If the United States hoards its technology too strictly to meet the objective of keeping advanced capabilities within the United States, it may push other countries to build on Chinese AI infrastructure, repeating the geopolitical dilemmas of the global 5G build-out, where Huawei established a dominant global position. This would not only mean lost sales for U.S. companies but would also mean U.S. policymakers would lose a critical opportunity to promote stronger AI safety and security standards globally by conditioning access to U.S. infrastructure capabilities on certain requirements. This could include restrictions on Chinese access to AI infrastructure, but also requirements to meet certain AI safety and security standards that would contribute to limiting misuse and misalignment. When done thoughtfully, a proactive position of promoting U.S. AI infrastructure globally can therefore advance all three national security interests contemplated in this paper.

At present, however, there is uncertainty about the future direction of U.S. compute policy. The Trump administration has relaxed controls on certain AI chip exports to China, although in practice has not yet approved sales in significant numbers. It has also signed agreements to partner with the United Arab Emirates (UAE) and Saudi Arabia on massive new AI data centers, but again, the implementation of these programs has been slow and inconsistent. The administration also announced it would not be enforcing the global regime for AI export controls released in the waning days of the Biden administration but has not clarified what, if anything, may replace it. Meanwhile, the administration is launching a new effort to facilitate U.S. AI exports to foreign countries, which could be an important program for both supporting U.S. company sales and establishing AI safety and security requirements, though as of early April 2026, details remain scarce.

Model and Application Layers

At the foundation model and application layers, there is significantly more uncertainty about the alignment or tension between commercial incentives and national security interests, given the less mature commercialization strategies at these layers. AI companies will design products that are responsive to the demands of their customers, which could be enterprises, consumers, governments, or internal users (or a mix of all of these). Thus, a key issue for U.S. national security is which customer markets will be most lucrative, and whether those customers' preferences will reflect a demand for safe and secure AI in a way that aligns normal market forces with national security interests. This, in turn, will depend largely on the extent to which such customers internalize costs associated with various AI misuse and misalignment risks.

Starting with the enterprise market, established businesses will tend to be risk averse in their uptake of AI tools, reflecting more cautious corporate cultures and the need to ensure compliance with various regulatory obligations. They will want total assurance that they understand how AI models work before they entrust them with any decision-making authority, and AI use and implementation will tend to be monitored and overseen by IT and cybersecurity departments that could help identify anomalous activities. This is particularly true for highly regulated industries, such as critical infrastructure companies, which would face significant reputational risk and potentially legal liability if an AI system deployed on their systems disrupts sensitive or lifesaving functions. Like governments pursuing sovereign AI initiatives, large enterprise users may also have a preference to bring more capabilities in-house to exert greater control over sensitive data and avoid depending on outside providers for mission critical capabilities, even if doing so means they are giving up some of the most advanced frontier capabilities.

The risk aversion of enterprise users means they will tend to have a stronger demand for safe and secure AI practices than individual consumers. The expectation that the enterprise market will ultimately be the most lucrative path to commercialization is a positive sign and could spark a race to the top as AI companies compete to demonstrate the interpretability and reliability of their models. This would be rewarded by enterprise customers and would have imperfect but nonetheless positive spillover benefits for national security interests related to use and alignment.

However, an enterprise-dominated AI market is still speculative, and corporate caution could end up being a double-edged sword by slowing overall adoption and constraining the viable market for enterprise end uses. According to surveys from the U.S. Census Bureau, only 18 percent of U.S. businesses currently use AI for any business functions, and only a marginally larger share expect to use AI in the next six months.⁴³ Enterprise AI adoption may be gated by legal and cybersecurity teams who will prohibit AI use or limit it to only certain prescribed, low-risk use cases.⁴⁴ A lack of technical understanding of both AI's capabilities and risks might contribute to low enterprise adoption. AI companies have taken a number of steps to share information and increase transparency of their products, such as the practice of publishing model cards to accompany the release of each new foundation model. Yet model cards are not standardized across different companies, and the primary audience for AI model cards appears to be the technical AI developer community rather than the broader ecosystem of AI users and stakeholders.⁴⁵ In practice, it is difficult for an enterprise choosing between different AI applications (which may be based on the same or different foundation models) to meaningfully use the information in AI model cards to weigh the risks of misuse or misalignment for its particular AI deployment. Similarly, while a number of independent organizations, including Epoch and Model Evaluation and Threat Research (METR), provide external assessments of model capabilities, these are aimed primarily at research and AI safety audiences rather than broader private sector AI customers.⁴⁶ The result is that enterprises looking to purchase safe and vetted powerful AI systems may struggle to identify these options in the market, leading them to either unwittingly purchase AI that is riskier than they would like or to forgo purchasing it all together.

Additionally, some major insurers are beginning to write AI exceptions into their general liability insurance offerings because they have struggled to precisely identify and quantify the associated risks.⁴⁷ This might further chill AI adoption among enterprises worried about liabilities for costs associated with AI mishaps. And though some companies are beginning to offer AI-specific insurance policies, to date this market is underdeveloped.⁴⁸

Turning to the consumer end market, customer demand for safe and secure AI practices will likely be more limited. Consumers are unlikely to internalize some of the broader societal risks, including national security risks, that AI poses. For instance, individual consumers may worry about cybersecurity breaches that target their personal information but be less concerned if the AI models and applications they are using contribute to more systemic cybersecurity risks.⁴⁹ Moreover, information asymmetries between end users and AI companies create a mismatch between customer demand and company incentives. The history of other consumer-focused online products suggests consumers place little value on accurately understanding how their data is being used and processed or in parsing the legal details of terms and conditions contracts.⁵⁰

In other contexts, imposing legal liability on the private sector is one mechanism to address information asymmetries and protect consumers. However, how liability applies in the case of AI misuse and misalignment remains highly uncertain. Recent legal scholarship suggests the liability protection of Section 230 of the Communications Decency Act of 1996, which shields social media companies and other internet sites from responsibility for user-generated content posted on their platforms, likely does not protect generative AI products, since the models and applications are developing content, not merely hosting content developed by others.⁵¹ Greater clarity on the bounds of legal liability for both underlying models and specific applications could help align incentive structures to appropriately motivate AI companies to develop adequate guardrails while not overly deterring innovation, though specific details on how to strike that balance are likely to be highly nuanced and contested. Ideally, such clarity would be provided in advance by new legal regulations adopted by Congress, but perhaps more realistically it may come from litigation. Indeed, a number of court cases assessing the scope of liability for consumer-facing AI products are already underway, arising out of tragic cases where individuals interacted intensively with AI-powered chatbots before committing suicides and murders.⁵²

The government end market for AI models and applications raises a complex web of risks and opportunities. Some features of the government market will tend to alleviate national security risks associated with powerful AI. Governments will likely be even more risk averse than enterprise customers, particularly for military and national security use cases, where there will be strong demand for reliability and zero tolerance for AI systems that act in unpredictable ways.⁵³ Moreover, relative to any other potential customers, government end users are more likely to internalize the costs of a broader range of catastrophic AI safety risks. Governments can also use their market power to shape AI markets to encourage safety and security, such as using procurement policies to define safety requirements that become industry standards.

Yet the government's inherent powers can also create risks of misuse that would not be present in the private sector. For example, the government may push AI companies to relax or remove certain guardrails against misuse, including in the context of military and law enforcement applications. The ongoing Anthropic-Pentagon contract dispute is over exactly these issues.⁵⁴ Law enforcement and surveillance applications raise similar concerns, as increasingly more powerful AI systems deployed for these purposes may enable government intrusion into citizen privacy in more pervasive ways than ever before, with potentially serious implications for human rights and democratic institutions. Different governments will have different definitions of what they consider misuse of AI in controlling domestic populations or gaining military advantage, risks that will rise when the government in question is not democratic.

Finally, models and applications designed for internal use cases pose a particular set of risks. Since they will never be brought to market for external customers, they will not be subject to the same levels of transparency and outside scrutiny as other AI products, and thus there may be greater uncertainty on their true capabilities or risks of model misalignment. Companies designing powerful AI systems for internal use cases may perceive that they do not need to invest in the same level of guardrails and safety protections since they assume it will never be released to a broader public. However, the model could later be leaked or otherwise made more widely available in a way that harms U.S. national security interests. Consider, for instance, a hypothetical biotech company developing its own tailored AI models and using them to automate biological research. If the company is incentivized to carry out this work in secret, shielded from public scrutiny and oversight, there may be a heightened risk that such processes could accidentally lead to the creation of dangerous pathogens. At the same time, AI models and applications for internal use cases would, by definition, restrict access to a smaller set of approved end users, which could alleviate risks of their misuse for malicious purposes. Developing an AI model capable of creating biological pathogens in secret and without oversight creates one set of challenges but developing it publicly and offering it for sale to any willing buyer would also come with its own set of risks. Differentiated transparency requirements for disclosing the risks of internal models to the government, investors, or the public (including through Security and Exchange Commission filings for public companies) may be important as a tripwire to ensure that external actors have sufficient visibility to guard against emerging misuse and misalignment risks.

Market Structure Across the AI Sector

Throughout the AI stack, companies are seeking to define market positions that provide a moat shielding them from competition while also worrying about overconcentration among both their suppliers and their customers. How these market dynamics develop, at each individual level of the stack as well as holistically across the stack, will shape opportunities and constraints to enable beneficial AI uses while limiting risks associated with misuse, misalignment, and threats to the U.S. geopolitical position.

Key Takeaways

- The AI market has meaningful concentration at the infrastructure and model layers of the AI stack, which may be exacerbated by vertical integration and circular financing deals among key players.
- There is significant uncertainty on how the foundation models market will evolve, including how market demand may be segmented between open-weight and closed-weight models and the extent to which “good enough” commoditized AI will capture a substantial share of the total market.
- Shifts toward a monopolistic market structure at any layer of the AI stack may erode the government’s power to protect U.S. national security interests if any one AI provider becomes too big to fail. At the same time, a more decentralized market may present different challenges associated with difficulties in controlling access to powerful AI and coordinating on best practices for AI safety.

Concentration in the AI Tech Stack

The degree of market concentration varies across different layers of the AI stack. The majority of advanced AI chips are designed by one company (NVIDIA), and most would-be challengers are also U.S. companies, with the notable exception of China’s Huawei. Advanced AI chip fabrication is concentrated at the Taiwan Semiconductor Manufacturing Company’s (TSMC’s) facilities in Taiwan, and while there is a push to increase and diversify production capacity, this will likely remain an important bottleneck on AI build-out. The data center market is dominated by a small handful of U.S. companies. At the model layer, three U.S. AI labs compete for the frontier, followed closely by other labs and increasingly competitive open-weight models. The market at the application layer remains fluid as companies determine product fit and often compete with established players, such as incumbent SaaS companies.

AI Infrastructure

The AI infrastructure layer consists of the companies designing and fabricating advanced AI semiconductors, as well as the data center and cloud companies that install these chips into compute clusters for AI training and inference. The largest share of realized revenue in the AI boom to date has flowed to the companies making the advanced semiconductors necessary to train AI models. NVIDIA earned \$130 billion dollars in revenues in FY2025, more than twice its income from a year earlier.⁵⁵ The company currently has a dominant position in designing semiconductors, and thus as demand for such chips has soared, the company has been able to exploit its market position to reap outsized profit margins. In FY2025, its gross profit margin reached a remarkable 75 percent; even with huge outlays on R&D, its net profit margin of 56 percent is still about five times the rate of the S&P500 average.⁵⁶ This combination of strong revenues, astronomical profit margins, and expectations of high future demand has helped make NVIDIA the most valuable company in the world.

Looking forward, however, there are some signals that NVIDIA’s position could be challenged in the future. A handful of other companies are actively working to develop and gain market share with

competing AI chips, including established companies like AMD, newer chip companies like Cerebras, as well as large tech companies now looking to develop their own chips, including Google, Amazon, and Tesla. The research institute Epoch AI tracks data on AI chip sales using financial reports and company disclosures. Its data suggests that though NVIDIA clearly remains the dominant player, alternative chip designers made up about one-third of the market by the end of 2025. Notably, this includes some limited production of Huawei chips, although research suggests Huawei remains far behind leading-edge U.S. companies.⁵⁷

Greater competition among AI chip designers could, in principle, cut into the extreme profit margins of NVIDIA, potentially leading to less expensive compute (which is, at present, one of the most important bottlenecks for training new frontier AI models). This is perhaps one of the reasons leading labs are actively pursuing compute strategies that focus on diversification; for instance, Anthropic publicly notes that it trains its models on a multiplatform combination of Google's tensor processing units (TPUs), Amazon's Trainium, and NVIDIA's graphics processing units (GPUs).⁵⁸

Yet in the short term, the price of AI chips reflects not only NVIDIA's dominant market position in chip design but also severe constraints in the production of AI chips. There are only a limited number of labs able to produce the most advanced chips needed for training frontier models; most of these are in Taiwan and operated by TSMC. Chip producers Intel and Rapidus are looking to bring more advanced production capacity online, and TSMC is also expanding its geographic footprint in the United States and Japan.⁵⁹ Such efforts will increase the overall capacity for AI chip production and strengthen the resilience of chip supply chains by diversifying away from Taiwan. But bringing these new facilities online will take time, so in the near term, AI chip demand will likely continue to outstrip supply, and compute will remain a key bottleneck in expanding AI capabilities. Meanwhile, just as Huawei is seeking to emerge as a Chinese alternative for chip design, Semiconductor Manufacturing International Corporation (SMIC) hopes to become a globally competitive chip manufacturer but has struggled to produce advanced chips at scale.

Turning to data centers and cloud computing, historically three American companies—Google, AWS, and Microsoft—have accounted for about two-thirds of the global cloud market generally, including both traditional cloud compute services and newer AI cloud capabilities.⁶⁰ Today, there is a scramble to build new data centers and cloud compute capacity to meet projected AI demand, with substantial investments by established hyperscalers and foundation model companies, as well as “neocloud” compute providers, such as CoreWeave and Nscale, that focus almost exclusively on AI cloud compute infrastructure. These entities are spending massive sums to acquire AI chips, which they intend to recoup by leasing compute to companies training and/or running large AI models (or by using them for their own model development). While multiple companies are racing to build new data centers, they are facing various constraints in the real world: delays in the permitting and approval process for new buildings, strained power grids struggling to provide sufficient energy for massive data centers, and capacity constraints on the chips to fill them.

Data centers provide a mostly undifferentiated product: access to computing power.⁶¹ Yet the economies of scale arising from operating a global network of data centers, tied to the significant financing necessary to acquire chips, have functioned as a barrier to entry, allowing cloud providers to earn substantial profits in recent years. The economics of the current AI build-out, however, are more speculative and thus potentially riskier than traditional cloud offerings. Uncertainty over the ultimate demand for AI training and inference translates into uncertainty on demand for the new crop of data centers, and it remains to be seen if today's investments in AI data center capacity will ultimately be profitable.

Meanwhile, as is true across the AI stack, Chinese companies are developing their own offerings both for the domestic China market and to compete overseas. Alibaba, Tencent, and Huawei all offer data

center and cloud services in emerging markets, typically priced below the rates of Western alternatives. These companies have pre-established market footholds in information and communications technology infrastructure in many of these countries, which may give them a leg up. However, they continue to be constrained by limited access to AI chips, both because of market constraints as well as the restrictions imposed by U.S. export controls.

Foundation Models

Three firms—OpenAI, Anthropic, and Google—are generally seen as leading the race to develop the most capable foundation models, with a few others—including Meta, xAI, Alibaba, and DeepSeek—also in contention alongside a longer tail of smaller companies. The leading companies are in sharp competition with each other, releasing new models within weeks or days of each other and measuring themselves against common benchmarks. Meanwhile, high barriers to entry, namely the need for access to large quantities of compute to train new models, limit new entrants. Access to top-tier AI talent also serves as a barrier to entry, as only a select number of engineers globally are capable of developing frontier AI models. Recent efforts by AI labs to recruit AI engineers with salaries running to the hundreds of millions of dollars, and in some cases higher, is just one example of the fierce competition for AI talent.⁶²

Yet it is plausible to imagine alternative market structures. At one extreme, the market could trend toward a monopoly, where a single AI firm captures the market.⁶³ This might occur, for instance, if one of the current leading labs achieves a technological breakthrough that allows for recursive self-improvement, such that a small technological lead grows exponentially over time as the AI system is increasingly able to train itself without human intervention. In such a scenario, foundation models could become a winner-take-all market where one of the labs' capabilities so far exceeds that of its rivals that it fully dominates the market.

At the other extreme, the market could become less concentrated due to competition from new entrants. Some “fast-follower” AI companies, such as DeepSeek, which train models at a fraction of the cost of the large leading labs, have achieved impressive capabilities that could translate into significant market share. If fast-follower companies are able to quickly erode any competitive advantages of the frontier labs, this could ultimately stall the rate of technological progress: the labs pushing forward the frontier will see little commercial incentive to invest massive sums necessary to develop the next generation of models if their less well-capitalized competitors can quickly copy their work at much lower cost. Yet it is unclear to what extent these fast-follower models depend on illicit distillation, the practice of training a model based on a competitor's model output in violation of the competitor's terms of service. OpenAI, Anthropic, and Google have all publicly noted that they have identified such distillation from Chinese AI companies and are implementing technical measures to limit it.⁶⁴ On April 23, 2026, the White House announced a new initiative to counter adversarial distillation of American AI models.⁶⁵ If such measures are successful, they may prevent fast-follower models from keeping pace with the leading labs.

Competition at the frontier might also come from so-called neolabs, a group of smaller companies and research institutes seeking to disrupt the grip current leading labs hold on the market.⁶⁶ Some of these neolabs are attempting to develop less compute-intensive methods for training advanced models. This could lower the currently high barriers to entry and lead to a more open market. At the extreme, leading AI models could effectively be commoditized.

Whether the market for foundation models trends toward concentration or dispersion, there is also a related question of how the market may ultimately be segmented along two (related) dimensions of closed-weight versus open-weight models and exquisite, frontier models versus “good enough” AI. Most AI experts expect there will be persistent market demand for both open-weight and closed-weight models: Certain cost-sensitive customers, including tech start-ups, may opt for open source, but many

customers will prefer closed models either because they are more capable or simply more user-friendly and convenient.⁶⁷ One recent analysis suggests open-weight models make up about 30 percent of the market for foundation models, while the price of using an open-weight model is approximately 90 percent less than closed-weight model equivalents.⁶⁸ Open-weight options are likely to promote competition and choice in the market for foundation models, but might have market distorting effects if companies seek to use open-weight models to lock customers in to tech ecosystems and/or purposefully undercut rivals' ability to sustainably sell closed-weight models. Today, the leading open models are primarily Chinese, although some U.S. companies, including Meta and OpenAI, have released open models and more might do so in the future based on shifting market dynamics.

Similarly, even if the market for frontier models with the most advanced capabilities does not become commoditized, it could still be the case that market segmentation results in a smaller market for the top frontier models, used for advanced scientific research for instance, alongside a much larger, mostly commoditized market of "good enough" AI for most consumer and basic enterprise services. The "good enough" models could be either open-weight or closed-weight models offered at low price points, offering products that achieve ~80 percent of frontier performance at only ~10 percent of the price. Such a scenario would mean a smaller total addressable market for models at the frontier, which could disincentivize investments in advancing the frontier. At the same time, however, if the leading labs determine the biggest market opportunities are in lower-cost, lower-capability models, they could adjust their product offerings to compete at both the high and low end of the market, rather than simply ceding the low end to open-weight and/or "fast-following" competitors.

Applications

AI applications are built on top of foundation models, although they may involve significant fine-tuning to adapt the underlying model to a particular use case. It is difficult to precisely characterize the market structure of the AI application layer, as it is segmented across a wide variety of end uses and because the major models also offer their own application interfaces (such as ChatGPT for OpenAI and Claude for Anthropic). For example, a number of companies compete against each other to offer AI applications focused on the legal industry, a separate set of companies offer rival products applying AI to the robotics industry, etc., while the major companies developing foundation AI models also offer access to their proprietary models via their own applications. Companies developing AI applications compete not only against each other but also against incumbent SaaS companies. In early 2026, several leading SaaS companies experienced sharp market falls as investors worried new AI products would eat into profits.⁶⁹

While foundation models compete on benchmarks of model performance, AI applications compete on specified domain expertise (often linked to access to proprietary technical data) and appealing user experience and workflow integration. A key question for the future of the AI application market structure is how switching costs develop between the model layer and the application layer. That is, one could imagine a market structure where customers easily switch between different AI applications but prefer to stick with the same underlying foundation model, or where customers are locked into a particular application, but that application can easily switch between competing underlying models. This will also vary by market segment: established enterprise companies may find it difficult to switch between applications after they have built workflows around them, while start-ups or consumers may be more open to switching. Such dynamics will shape how much competition or concentration there is at both the model and application layers.

While companies focused on the application layer have not garnered the same sky-high valuations or popular attention as companies working on foundation models, the application of AI capabilities will be crucial in translating technological AI advances into concrete economic, scientific, and social progress. And though U.S. companies lead in developing the most capable foundation models, China's position at the AI application layer is arguably more competitive.⁷⁰ This is in part because a wide range of

application companies have defined market niches by building on top of freely available Chinese open-source models and fine-tuning them for particular use cases. The ability to apply AI in a variety of real-world use cases will likely have strong implications for public opinion on AI.

Vertical Integration Across the AI Stack

In practice, many AI companies operate at multiple layers across the AI stack. Google, for example, can make its own chips, run its own data centers, has developed its own model (Gemini), and has a built-in suite of consumer and business applications (e.g., search, email) into which it can plug its growing AI capabilities. This type of vertical integration will likely be as important in shaping the overall level of competition in the AI stack as the horizontal concentration at separate layers. Vertically integrated AI companies may have commercial incentives to restrict, or in some cases subsidize, other market players' access to one layer of the stack in an effort to maximize their profits at another layer. For instance, a company with popular AI applications might force customers to use its own proprietary model, or a company with a dominant position in the cloud and data center market could provide compute to train its own models at a much cheaper price than it charges to other companies developing competing frontier models. These incentives could distort markets and lead to less competition in the AI stack.

These vertical integration challenges may also be compounded by various complex investment and partnership relationships across layers of the AI stack. NVIDIA, for example, has invested in companies focused on data centers (Crusoe, Nscale), foundation models (OpenAI, Anthropic, xAI, Mistral, ReflectionAI, Thinking Machines Lab), and applications (Cursor, Perplexity).⁷¹ Such deals will often include linked customer contracts: NVIDIA's investments in OpenAI, for instance, were tied to commitments from OpenAI to buy NVIDIA chips. Similar interdependent arrangements have become common throughout the AI stack.⁷² Thus, even where vertical integration within companies may be limited, AI companies may still have commercial incentives to see particular companies at other layers of the stack succeed, again distorting normal market behavior and potentially limiting competition. Circular financing partnership deals can also obscure true levels of market demand and give inflated impressions of future cash flows, thereby driving artificially high valuations. The complex web of financial relationships across the AI sector could mean that if one central node faces market pressure, it could quickly spread to other major AI companies.

How Commercial Incentives Align with National Security Interests

The evolution of market structures in the AI stack, and particularly the degree of concentration at each layer of the stack as well as holistically across the stack, will have important implications for AI national security interests. Many AI experts have identified the degree of centralization of AI development as a critical factor in understanding how advanced AI will impact national security, as this will fundamentally shape how governments control AI development and its consequences.⁷³ Understanding centralization dynamics across the AI stack can inform policy responses, as policymakers may have more leverage at one layer of the stack than at others (and such leverage may shift over time and across jurisdictions).

There are three main cross-cutting effects to consider. First, higher market concentration could allow one or a small handful of companies to acquire so much power that they could have undue influence over the development of AI and pose significant threats to other actors. Second, lower market concentration could make it more difficult for both public and private stakeholders to impose some degree of control and restrictions on who has access to powerful AI technologies. And third, particular features of the AI market, namely the need for massive capital expenditures to be globally competitive, mean that a substantial scale is necessary for companies to maintain a viable market position. Taken together, these three effects suggest a complex, non-linear relationship between market concentration and U.S. national security interests in AI.

Consider a scenario where a single entity emerges as a monopoly with complete control over one or several layers of the AI stack. Such an outcome might reduce some national security risks associated with AI. Because access to the technologies could be more tightly controlled, it could be easier to prevent adversaries from acquiring substantial compute that could be used for military AI, or to prevent terrorist groups from using advanced AI models to plan biological or chemical weapons attacks. If the most powerful AI is controlled by one firm, it may be easier to adopt a responsible approach to technological development and ensure safety and interpretability research keeps pace with advances in AI capabilities. Anthropic's decision to selectively release its Mythos model only to approved users for the purpose of strengthening cybersecurity defenses is an example of how tight control over frontier capabilities can be leveraged to lessen AI misuse risks.

Yet a single company with a monopoly on powerful AI would be the epitome of too big to fail, as society would be incredibly dependent on this single company. Concentrating so much power within a single entity would require great trust, akin to hoping for a benign autocrat rather than a despot. A private sector company with a monopoly on powerful AI would occupy such a critical role in society it would be difficult for government to regulate it as a normal entity, as the company would have leverage to push back against any unwanted constraints. The variety of ways in which a single entity (either a firm or government) with a monopoly on powerful AI could abuse this power suggests at least some dispersion of AI capabilities among multiple actors might be beneficial so that "good AI" can serve as a check and counterweight to "bad AI."⁷⁴ There may be particular risks with monopolization at the frontier model layer, given that this could imply deep societal dependencies effectively on a single product, and vulnerabilities or failures with the product could have catastrophic effects.

The other end of the spectrum, a world where there is complete decentralization at all layers of the stack, would present a different set of national security opportunities and threats. Widely available and affordable AI could allow for broader adoption of beneficial AI and would ensure no single entity could use its dominance in AI as a leverage point for greater political or geopolitical control. But it would also become much more difficult for governments to prevent bad actors from acquiring dangerous capabilities, particularly if markets are broadly decentralized across countries. Export controls will be ineffective if substitute products are generally available in foreign markets. It would be more difficult for companies to coordinate around AI safety best practices. Companies may be less willing to pay an "alignment tax"—costs associated with AI safety, red teaming, etc.—if they perceive that they are competing in a wide-open market where their competitors are not similarly bearing such costs.⁷⁵ For instance, reports suggest the cost of including safety guardrails in inference of leading AI models runs is 5 percent to 16 percent of total inference compute costs.⁷⁶

Between these two extremes of monopolization and widely decentralized markets lies what is perhaps the most likely scenario, where significant barriers to entry prevent AI markets from becoming open and commoditized, but where no single entity is able to gain an entrenched dominant position.⁷⁷ Such market dynamics should help alleviate some of the national security risks associated with high concentration of power: If any single entity with access to powerful AI capabilities sought to abuse its position, then other market actors could respond by turning to an alternative supplier. Whether such market structures support or undermine national security interests associated with responsible use and alignment policies will depend in part on whether the companies involved can avoid an arms race dynamic where competitive pressures push them to cut corners on AI safety.⁷⁸ A market structure with a handful of key players may also make it easier for policymakers to shape market behavior through more nimble forms of soft regulation, such as voluntary commitments and codes of practice: It may be more feasible for the government (or other stakeholders seeking to advance AI safety) to coordinate a small group rather than negotiate with one behemoth company or deal with a flock of smaller companies.

All of this suggests that antitrust and competition policy will be important for advancing national security interests in the AI market. Policymakers, however, will need to take an agile and nuanced approach. They will need to carefully look for signs that any layer of the AI stack is trending toward monopoly and consider how vertical integration and cross-stack financing deals and partnerships may contribute to a consolidation of market power. But they should also recognize that there may be some downsides associated with an overly zealous approach to competition policy. For instance, in scenarios where there are a handful of key market players, allowing these companies to coordinate on shared safety practices could help dampen arms race dynamics and encourage responsible AI development.⁷⁹ Yet such coordination efforts might otherwise look like collusion to competition policymakers, who are generally very wary any time market competitors seek to coordinate business activities.

Finally, it is worth considering the national security implications of Chinese-headquartered companies across the AI stack gaining global market share. While Chinese companies still lag significantly behind their U.S. counterparts when it comes to producing AI chips and developing the best foundation models, they may be positioning themselves to win a significant share of the market for “good enough” models and applications.⁸⁰ This could pose a range of threats to U.S. national security interests.⁸¹ While data is limited, analyses suggest Chinese-developed models tend to have less well-developed transparency and safety standards than leading U.S. models, and thus likely pose greater risks associated with AI misuse and misalignment.⁸² Moreover, because most Chinese models are open weight, the safety filters that do exist could be removed by malicious actors through fine-tuning.⁸³ If Chinese AI companies achieve substantial global market shares, this could provide China with significant geopolitical leverage and allow China to steer the overall development of AI technologies.

While a comprehensive assessment on how to keep the United States ahead of China in the global AI race is beyond the scope of this paper, there is one issue directly tied to evolving commercial incentives that is worth highlighting. China’s ability to win significant global market share is closely tied to its companies’ decisions to offer open-weight models and sell AI applications at very low price points. Given the perception among AI market participants that there will be persistent demand for open-weight options, it would be a mistake for the United States to cede this market share to China. U.S. companies and policymakers should consider how to enhance U.S. offerings of open-weight models and related inexpensive AI applications.

A Policy Framework to Align Commercial Incentives and National Security Interests

As this paper has highlighted, there are several outstanding questions and uncertainties about how commercial markets for AI will develop. AI markets are evolving extremely rapidly, and commercial strategies may shift quickly. For example, even just during the months researching and writing this paper, AI agents emerged from a small niche tool to potentially revolutionizing all SaaS business models.⁸⁴ It is difficult to predict with any level of certainty what new market opportunities may open or be foreclosed in the months and years ahead.

As this paper has also noted, the commercial incentives of companies throughout the AI stack may not always align with the national security interests of the United States and its allies. And even when private sector actors have a collective interest in promoting responsible AI diffusion, their individual incentive structures may be at odds, for instance, when it comes to avoiding arms race dynamics and the pressure to release new models as quickly as possible, even when doing so entails heightened national security risks.

The challenge for policymakers, then, is how to develop policies amid this uncertainty to shape market behaviors to better align commercial incentives with national security interests. Their objective should be crafting policy frameworks that are robust and adaptable to a series of different market scenarios. This section outlines the elements such policy frameworks should include.

A Layered, Adaptable Approach

Over the long run, both the public and private sectors would be well-served by a regulatory regime for governing AI that provides clear technical guidance on what commercial activities are off limits to safeguard national security interests, and what targeted policy interventions are needed to correct market failures or otherwise align commercial incentives with public interests. For the moment, however, technological and commercial uncertainties and the incredibly fast pace of change mean any near-term effort to rely on top-down technocratic rules would likely backfire. The technology is moving so quickly that rigid rules would risk both shutting down potentially beneficial developments while also missing potentially threatening new vulnerabilities, as government rule-making processes would undoubtedly lag behind the technological development curve. In place of a comprehensive set of rigid rules, there are several layers of risk-based policies that public and private actors can jointly implement to advance national security interests while promoting responsible AI diffusion.

First Layer: Infrastructure and Compute Policy Remain the Bedrock

The United States and its close allies have a substantial lead in producing AI compute and infrastructure, and government policy should capitalize on this advantage to support the development of global AI markets in line with U.S. geopolitical interests. There are four key steps policymakers should pursue:

- **Address constraints to AI infrastructure build-out in the United States.** Demand for building AI data centers in the United States is strong, but various constraints and bottlenecks could slow development. U.S. policymakers should continue ongoing efforts to meet these challenges, which include permitting reforms, bringing new power sources online to meet the energy needs for AI infrastructure, and addressing bottlenecks on chip fabrication. The continued reliance on TSMC chips from Taiwan is a vulnerability to U.S. AI dominance, and U.S. policymakers should support more diverse and resilient chip supply chains. This should include ensuring CHIPS Act-supported plants currently under development come to fruition; capitalizing on investment promises attached to recent trade deals with Taiwan, Japan, and South Korea to build further

domestic fabrication capacity; and leveraging the government's 10 percent ownership stake in Intel to accelerate the company's revival as a domestic chip-making champion.

- **Develop a coherent global vision for governing AI compute that uses both export controls and export promotion to advance U.S. geopolitical interests.** Confusion about the direction of U.S. compute policy makes it difficult for both U.S. companies and allied governments to develop long-term plans. The U.S. strategy should balance the U.S. objectives of keeping the preponderance of compute in the United States and its close allies while also encouraging foreign governments and companies to build on the U.S. AI stack where feasible. The strategy should include strict limits on China's access to advanced chips and semiconductor manufacturing equipment, including with robust enforcement. This limits China's own infrastructure build-out and hampers Chinese companies' efforts to gain market share in third countries, as their limited productive capacity will be devoted to China's domestic market instead of breaking into foreign markets. The government should also properly resource the Commerce Department's AI Exports program, reflecting the fact that U.S. global leadership in AI infrastructure is a key geopolitical advantage.
- **Leverage U.S. advantages on compute to embed basic AI safety and security conditions into AI infrastructure approvals.** U.S. dominance at the infrastructure layer provides important leverage to encourage foreign companies developing and deploying AI on top of U.S. infrastructure to invest in practices to prevent AI misuse and misalignment. This not only provides an opportunity to spread U.S. regulatory standards around the world, but also crucially will ensure U.S. AI exports do not inadvertently create an opening for regulatory arbitrage that could undermine U.S. AI competitiveness. Demanding that foreign customers acquiring U.S. AI infrastructure adopt a common baseline regulatory approach in developing and deploying AI will promote a level playing field and prevent foreign companies from seeking a competitive advantage by, for instance, cutting corners on data center security requirements. AI experts at the National Institute of Standards and Technology (NIST) and the Center for AI Standards and Innovation (CAISI) should design conditions the United States will require of foreign customers that want access to U.S. AI infrastructure, which should include data center security standards (both cybersecurity and physical access) as well as basic know-your-customer (KYC) rules and deployment auditing to limit access for harmful end uses or by malicious end users. To the extent practicable, these conditions should be standardized across various policy instruments, including export control licensing, requirements for participation in the AI Exports Program, and government-to-government frameworks such as those recently agreed with the UAE and Saudi Arabia, and harmonized across U.S. domestic and international markets.
- **Track both technological and market developments that might lessen the importance of compute and infrastructure dominance in shaping AI trajectories and begin thinking through such hypotheticals now.** For instance, should new chip companies emerge to rival NVIDIA outside the reach of U.S. export controls, then updates to an infrastructure and compute-based strategy would be needed. Algorithmic advances could also make it easier to train very capable models on significantly less compute, which would make it more difficult to prevent certain AI national security risks. In a market scenario with multiple powerful open-weight AI models that can be run efficiently on small compute clusters, it might be impossible to rely on compute controls to prevent bad actors from misusing AI for individual cyberattacks or similar threats.⁸⁵ This is not to say compute controls would lose all efficacy in such a scenario: KYC requirements for data centers could help catch bad actors, AI misuse at scale would still require large quantities of compute that would be difficult to acquire outside of approved channels, and adversaries, including China, would still need access to large quantities of the most advanced compute in order to widely deploy AI models and applications. But they will

need to be complemented with risk management strategies to improve resilience to AI misuse, particularly in the cyber and bio domains.⁸⁶

Second Layer: Use Transparency, Liability, Insurance, and Procurement Policies to Align Market Demand for AI Safety and Security with Commercial Incentives

For foundation models and AI applications, in many instances commercial incentives could in principle support U.S. national security interests in AI, yet the market regulatory framework is incomplete or undefined, resulting in broken feedback loops and misaligned incentives. End market customers for AI products may want to purchase products that are safe and secure, but such demand signals may not reach the companies developing AI models and applications if customers have insufficient information to assess relative safety and security or are unclear on who will bear the liability for AI-related mishaps. There are four key steps policymakers should take to reinforce market mechanisms that could reward companies for adopting behaviors that will limit national security risks associated with AI misuse and misalignment:

- **Increase transparency and information-sharing requirements related to AI safety and security risks.** AI companies have already gone further than most private sector entities by voluntarily disclosing information related to risks associated with their products and business strategies, including in company-wide responsible scaling policies and AI model cards. While these informal best practices provide a useful baseline for AI transparency, there are sharp limits to such approaches that rely predominantly on voluntary commitments and self-assessments. This will require action from a patchwork of government agencies. NIST and CAISI should further expand their model evaluation program, working with industry to develop standardized best practices that can raise the floor for external model evaluations. The Federal Trade Commission (FTC) should assess consumer-facing AI applications and ensure they provide sufficient information to customers to enable users to understand relevant risks.⁸⁷ The Securities and Exchange Commission should issue guidance for enterprises on when potential risks associated with AI misuse and misalignment rise to the level of materially relevant financial information that must be disclosed to investors, including for companies deploying AI for internal use cases.⁸⁸ All such disclosures should focus on providing digestible and actionable information for a wide range of relevant AI users, to enable such users to make informed decisions on the AI products they purchase, invest in, or build their applications on top of. Additionally, the government should also consider mandating disclosure to CAISI for certain defined, high-severity misuse and misalignment risks. This could include, for example, mandatory disclosures should a U.S. model developer, AI application company, or cloud provider become aware that malicious actors are using U.S. AI capabilities to launch AI-enabled cyberattacks on U.S. assets or critical infrastructure, when an AI developer crosses certain technical thresholds related to weapons development, or significant instances of model misalignment. Such information should initially be provided confidentially to prevent the dissemination of new vulnerabilities to malicious actors, but the government should consider publicly releasing information when doing so is in U.S. national security interests (including by sending signals to market participants). Critically, these requirements should also apply to internal use cases of models, not only to models designed to be made commercially available.
- **Clarify legal liabilities to align incentives for companies to be accountable for AI safety and security and avoid liability gaps.** The market for AI safety would benefit from clarifying the legal liabilities arising from various potential risks associated with the misuse and misalignment of powerful AI.⁸⁹ AI mishaps will result from the interaction of decisions made by foundation model developers, companies building applications on top of foundation models, and end users deploying AI tools in particular contexts, but at present there is little guidance for what factors should be taken into account to apportion liability between these three sets of actors. Legislation

on liability for AI products would provide the clearest guidance, but even in the absence of any changes to law, policymakers and other stakeholders can still provide courts guidance on how current liability and tort laws should be applied to AI. Part of the problem is a lack of consensus on the best practices for AI testing and safety, which can make it difficult to determine whether an entity acted recklessly or negligently in developing and/or deploying an AI product. Government policymakers, industry, and other stakeholders should further define and promulgate best practices on AI safety and security, which courts can then point to in determining liability for any market actors not following such practices. This should also include standards on the due diligence required to understand AI model risks (both for those developing and deploying models) as well as how to effectively communicate these risks to end users. While clear liability standards should provide guidance on what actions AI model and application companies need to take to lessen their legal exposure associated with AI misuse and misalignment risks, policymakers should avoid broad blanket exemptions from liability along the lines of those provided to social media companies and other internet platforms under Section 230. This would lead to critical liability gaps, particularly as powerful AI systems increasingly act autonomously. In certain cases, increasing liability standards may be necessary, which deserves particular consideration for AI applications geared toward the consumer end market.

- **Support the development of AI insurance markets.** Clarifying liability standards will also help foster a more robust market ecosystem for AI insurance. A healthy insurance market for enterprise AI could create market incentives for safer practices among companies developing and deploying AI, as companies demonstrably addressing risks of misuse or misalignment could pay lower premiums. Yet to date, pricing such insurance has proven to be difficult, given uncertainty about the likelihood of AI mishaps and the ultimate liability for such events. Policymakers should work with industry and other stakeholders to support independent certification bodies that can provide assurance to enterprises that the AI models and applications they are purchasing and incorporating into their workflows have been appropriately vetted, which will serve as a key instrument for underwriting AI insurance.⁹⁰
- **Use procurement processes to send demand signals for AI safety and security and promote competition.** When governments act as commercial buyers of powerful AI systems, they should leverage this position to set high safety and security standards. Government buyers should prioritize developing capabilities and technical expertise to evaluate AI systems for misuse and misalignment risks. The procurement requirements governments develop can then also be adopted and adapted by private enterprises procuring their own AI products to become industry best practices. AI systems that are trusted by the U.S. government can use this as a marketing point to appeal to risk averse enterprise customers.

Third Layer: Targeted Market Interventions Can Encourage Safe and Competitive AI Markets in the United States

There are some instances where collective market outcomes may not align with U.S. national security interests. Targeted government interventions can shape and steer markets to encourage safe and secure AI development and deployment. These targeted market interventions can encourage the expansion of safe and competitive AI markets:

- **Provide antitrust guidance reflecting the economic and national security importance of AI markets.** Over the last 50 years, U.S. antitrust policy has primarily adopted the consumer welfare standard as the key criteria in assessing antitrust and competition policy, viewing market concentration through the lens of its impact on the prices and choices facing consumers. More recently, however, some antitrust advocates have advocated for a more expansive approach to antitrust, weighing not only the impact on consumers but also a broader set of public interests,

including potential harms to workers, producers, and national security.⁹¹ Putting aside the merits of this position as an overall approach to antitrust policy, it is clear that AI is one area where U.S. interests expand beyond a narrow consumer welfare standard, and competition policy for the AI sector should consider broader national security and geopolitical stakes of AI development.⁹² The FTC and the Antitrust Division of the Department of Justice should publish guidance on competition policy for the AI sector that (a) highlights that competition throughout the AI stack is a national priority, given risks associated with monopolization of AI; (b) defines acceptable levels of concentration across the AI stack, recognizing that full decentralization of AI markets may raise its own security risks; and (c) explicitly provides a safe harbor for firms to coordinate certain safety practices and share information that, in other contexts, could appear like collusion. The latter is particularly important in the context of sharing nonpublic information on how frontier AI companies might defend against distillation attacks, because this could be a commercial secret and aligning on common approaches could appear like an anticompetitive action.⁹³ In a promising step, the recent White House memorandum on adversarial distillation stated that the Trump administration intends to enable the private sector to better coordinate against such attacks, though details have not yet been released.⁹⁴ An appropriate level of flexibility in antitrust policy to reflect some of the benefits associated with private sector coordination on AI safety, and clarity from the government on what forms of coordination are acceptable and which are forbidden, could ultimately contribute to promoting safe and competitive AI markets. While companies too often play the national security card to push for softer regulation, it is nonetheless important for competition policy to account for U.S. geopolitical interests in maintaining global AI leadership.

- **Assess risks associated with the deployment of Chinese models and applications in the U.S. market and implement risk-based restrictions as necessary.** U.S. AI and national security officials should carefully consider what role, if any, Chinese AI companies and products should play in the development of safe, competitive, stable U.S. AI markets. There are already multiple reports of U.S. companies—primarily small start-ups, but also some larger enterprises such as Airbnb—incorporating Chinese foundation models into their AI applications.⁹⁵ At a minimum, the government should commit to conducting national security risk assessments of leading Chinese models, issue cybersecurity alerts and advisories related to Chinese AI systems where necessary, and establish security testing guidelines for U.S.-based entities that host, distribute, or provide managed access to Chinese AI models and applications.⁹⁶ Based on these risk assessments, should policymakers determine there is a need to restrict Chinese companies' access to the U.S. AI market, they should develop a coherent strategy and implement it through existing relevant policy channels including the Commerce Department's Information and Communications Technologies and Services supply chain authorities, the Committee on Foreign Investment in the United States authorities to review mergers and acquisitions, as well as various entity-based restrictions maintained by the Commerce Department, Treasury Department, and Department of Defense.
- **Encourage development of safe and secure U.S. open-weight models to challenge Chinese competitors.** The market penetration of Chinese open-weight models, both in the United States and in third markets around the world, underlines that there is substantial demand for open-weight approaches. It is critical that the United States does not cede this market to Chinese firms. U.S. policymakers, academic researchers, frontier labs, start-ups, and AI funders should work together to foster a rich ecosystem of U.S. open-weight models, while prioritizing safety and security. The Trump administration's AI Action Plan correctly highlights the need to support open-weight models, including by leveraging the National AI Research Resource pilot to make compute available to researchers. The Commerce Department's AI Exports program should also prioritize supporting the export of U.S. open-weight models to prevent Chinese competitors from cornering global markets.⁹⁷

Conclusion

The prominent role of the private sector in advancing the frontier of AI technologies has resulted in several benefits for U.S. national security. For instance, private capital markets have demonstrated a remarkable ability to funnel resources to the AI sector, allowing U.S. AI labs and technology companies to devote billions of dollars to the compute and talent necessary for R&D. This financing has been critical in allowing the United States to build its lead in AI.

Commercial incentives in the private sector, however, may not always align with U.S. national security interests. Yet to date, AI national security stakeholders have paid only limited attention to the potential linkages between AI market dynamics and U.S. national security. This paper has sought to contribute to closing this gap by mapping several key outstanding questions on the commercialization of AI and assessing their implications for three key U.S. national security interests: (1) enabling beneficial uses of AI while limiting security risks associated with AI misuse; (2) ensuring reliable and controllable AI system behavior in deployment; and (3) maintaining strategic geopolitical advantage in the development and global diffusion of AI. Policymakers may have limited power to fully determine how AI market dynamics will develop, yet there are still a number of actions they can take to place guardrails on commercial activities. The analyses in this paper are designed to provide an initial framework for public and private sector actors to shape market incentives in the AI sector to better align with U.S. national security interests.

Looking forward, the challenge for policymakers is how to prepare today for an uncertain tomorrow, recognizing that the pace of technological development is so rapid that officials cannot wait until markets shift to begin determining appropriate policy responses. Building on the analysis in this paper, researchers, policymakers, and AI companies should cooperate to use novel research methods to explore future market trajectories and their impacts on commercial strategies and national security outcomes. Tabletop wargaming exercises, where players representing AI companies and public officials react to hypothetical market scenarios, could elucidate some of the tensions and unexpected consequences of government efforts to shape markets to advance national security interests, and allow public and private stakeholders to prepare for a range of plausible future economic developments.

Finally, one of the fundamental takeaways from this research is the need for greater exchange and cross-pollination between the research community studying the economics of AI and the research community studying AI and national security. Both fields are currently exploring a wide range of fascinating, complex, policy-relevant questions, but with seemingly little engagement with one another. There would be significant policy and research gains from greater collaboration between these two expert communities.

1. David Fischer, *History of the International Atomic Energy Agency: The First Forty Years* (International Atomic Energy Agency, 1997), https://www-pub.iaea.org/MTCD/Publications/PDF/Pub1032_web.pdf; U.S. Department of Energy, *The History of Nuclear Energy* (U.S. Department of Energy, accessed April 9, 2026), <https://www.energy.gov/ne/articles/history-nuclear-energy>.
2. See, for instance, Yoshua Bengio et al., *International AI Safety Report 2026* (AI Security Institute, 2026), 97, <https://internationalaisafetyreport.org/publication/international-ai-safety-report-2026>, which notes that “Due to competitive pressures, AI companies may face trade-offs between faster product releases and investments in risk reduction efforts.”
3. For an overview of the debate on productivity, see Alex Imas, “What Is the Impact of AI on Productivity?” Ghosts of Electricity (Substack), January 29, 2026, <https://aleximas.substack.com/p/what-is-the-impact-of-ai-on-productivity> and Erik Brynjolfsson, “The AI Productivity Take-off is Finally Visible,” February 15, 2026, *Financial Times*, <https://www.ft.com/content/4b51d0b4-bbfe-4f05-b50a-1d485d419dc5>. For an overview of the debate on labor market impacts, see Jed Kolko, “Research on AI and the Labor Market is Still in the First Inning,” Peterson Institute for International Economics Realtime Economics (blog), March 20, 2026, <https://www.piie.com/blogs/realtime-economics/2026/research-ai-and-labor-market-still-first-inning>.
4. Dario Amodi, “Machines of Loving Grace: How AI Could Transform the World for the Better,” *Dario Amodi* (blog), October 2024, <https://darioamodei.com/essay/machines-of-loving-grace>.
5. Ezra Karger et al., “Forecasting the Economic Effects of AI,” Working paper (Forecasting Research Institute, March 2026), <https://forecastingresearch.org/economic-effects-of-ai>; Thomas Cunningham, “Forecasts of AI & Economic Growth,” (blog), November 9, 2025, <https://tecunningham.github.io/posts/2025-10-19-forecasts-of-AI-growth.html>; and Charles I. Jones, “A.I. and Our Economic Future,” Working Paper No. 34779 (NBER, January 2026), <https://www.nber.org/papers/w34779>.
6. Anton Korinek and Donghyun Suh, “Scenarios for the Transition to AGI,” Working Paper No. 32255 (NBER, March 2024), <https://www.nber.org/papers/w32255>; Matthew Barnett, “AGI Could Drive Wages Below Subsistence Level,” Epoch AI Gradient Updates (blog), January 24, 2025, <https://epoch.ai/gradient-updates/agi-could-drive-wages-below-subsistence-level>.
7. Erik Brynjolfsson, Danielle Li, and Lindsey Raymond, “Generative AI at Work,” *The Quarterly Journal of Economics* 140, no. 2 (2025): 889–942, <https://doi.org/10.1093/qje/qiae044>; Martin Neil Baily, Erik Brynjolfsson, and Anton Korinek, “Machines of Mind: The Case for an AI-Powered Productivity Boom,” Brookings Institution, May 10, 2023, <https://www.brookings.edu/articles/machines-of-mind-the-case-for-an-ai-powered-productivity-boom/>.
8. Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford University Press, 2014).
9. Holden Karnofsky, “Some Background on Our Views Regarding Advanced Artificial Intelligence,” Coefficient Giving, May 6, 2016, <https://coefficientgiving.org/research/some-background-on-our-views-regarding-advanced-artificial-intelligence/>.
10. Arvind Narayanan and Sayash Kapoor, “AGI is Not a Milestone,” AI as Normal Technology, May 1, 2025, <https://www.normaltech.ai/p/agi-is-not-a-milestone>.
11. Multiple interviews with employees at AI labs, January 2026. Interviews were conducted in confidentiality, and the names of the interviewees are withheld by mutual agreement. It is worth noting, however, that there is a range of views within the labs and no single consensus on timelines. See also Amodi, “Machines of Loving Grace”; Sam Altman, “The Gentle Singularity,” Sam Altman (blog), June 10, 2025, <https://blog.samaltman.com/the-gentle-singularity>.
12. Allan Smith, “Poll: Majority of Voters Say Risks of AI Outweigh Benefits,” NBC News, March 10, 2026, <https://www.nbcnews.com/politics/politics-news/poll-majority-voters-say-risks-ai-outweigh-benefits-rcna262196>.
13. Quinnipiac University Poll, “The Age of Artificial Intelligence: Americans’ AI Use Increases While Views on It Sour, Quinnipiac University Poll on AI Finds; 7 in 10 Think AI Will Cut Jobs with Gen Z the Most Pessimistic,” press release, March 30, 2026, <https://poll.qu.edu/poll-release?releaseid=3955>.
14. Lydia DePillis, “Local Opposition Is Slowing A.I. Data Centers: Wall Street Has Noticed.” *The New York Times*, March 26, 2026, <https://www.nytimes.com/2026/03/26/business/economy/ai-data-centers-construction-local-opposition.html>.
15. Monica Anderson, *Americans’ Views of Technology Companies* (Pew Research Center, April 29, 2024), <https://www.pewresearch.org/internet/2024/04/29/americans-views-of-technology-companies-2/>.
16. Huo Jingnan, “How AI is Getting Better at Finding Security Holes,” National Public Radio, April 11, 2026, <https://www.npr.org/2026/04/11/nx-s1-5778508/anthropic-project-glasswing-ai-cybersecurity-mythos-preview>.
17. Tim G. J. Rudner and Helen Toner, *Key Concepts in AI Safety: Interpretability in Machine Learning* (Center for Security and Emerging Technology, March 2021), <https://cset.georgetown.edu/publication/key-concepts-in-ai-safety-interpretability-in-machine-learning/>.
18. OpenAI was originally founded as a non-profit organization controlling a for-profit limited liability company, and in 2025 restructured as a public benefit corporation partially owned by a non-profit. OpenAI, “Evolving OpenAI’s Structure,” press release, May 5, 2025, <https://openai.com/index/evolving-our-structure/>.
19. Matt Levine, “OpenAI Has a Business Plan,” Bloomberg Opinion, Money Stuff newsletter, October 15, 2025, <https://www.bloomberg.com/opinion/newsletters/2025-10-15/openai-has-a-business-plan>.
20. Berber Jin, Corrie Driebusch, and Kate Clark, “OpenAI Plans Fourth-Quarter IPO in Race to Beat Anthropic to Market,” *The Wall Street Journal*, January 29, 2026, <https://www.wsj.com/tech/ai/openai-ipo-anthropic-race-69f06a42>.
21. “AI Chip Owners,” AI chip ownership over time, cumulative compute capacity, Epoch AI, accessed April 7, 2026, <https://epoch.ai/data/ai-chip-owners/>.
22. Berber Jin, “Oracle, OpenAI Sign Massive \$300 Billion Cloud Computing Deal,” *The Wall Street Journal*, September 10, 2025, <https://www.wsj.com/business/openai-oracle-sign-300-billion-computing-deal-among-biggest-in-history-ff27c8fe>.
23. Ben Cottier and Yafah Edelman, “What You Need to Know About AI Data Centers,” Epoch AI (blog), November 4, 2025, <https://epoch.ai/blog/what-you-need-to-know-about-ai-data-centers/>.
24. Pablo Chavez, “Sovereign AI in a Hybrid World: National Strategies and Policy Responses,” *Lawfare*, November 7, 2024, <https://www.lawfaremedia.org/article/sovereign-ai-in-a-hybrid-world-national-strategies-and-policy-responses>.
25. Martin Chorzempa, “How the AI Boom Shrugged Off the DeepSeek Shock and Keeps Gaining Steam,” Peterson Institute for International Economics Realtime Economics (blog), February 5, 2026, <https://www.piie.com/blogs/realtime-economics/2026/how-ai-boom-shrugged-deepseek-shock-and-keeps-gaining-steam>.
26. “The \$4trn Accounting Puzzle at the Heart of the AI Cloud,” *The Economist*, September 18, 2025, <https://www.economist.com/business/2025/09/18/the-4trn-accounting-puzzle-at-the-heart-of-the-ai-cloud>.

27. Stephen Morris, Michael Acton, and Rafe Rosner-Uddin, “Big Tech’s ‘Breathtaking’ \$660bn Spending Spree Reignites AI Bubble Fears,” *Financial Times*, February 5, 2026, <https://www.ft.com/content/0e7f6374-3fd5-46ce-a538-e4b0b8b6e6cd>.
28. Jaime Sevilla, Hannah Petrovic, and Anson Ho, “Can AI Companies Become Profitable?” *Epoch AI*, January 28, 2026 (revised March 6, 2026), <https://epoch.ai/gradient-updates/can-ai-companies-become-profitable>.
29. “Epoch Capabilities Index,” score vs release date trends, *Epoch AI*, accessed April 7, 2026, <https://epoch.ai/benchmarks/eci>.
30. Mert Demirel et al., “The Emerging Market for Intelligence: Pricing, Supply, and Demand for LLMs,” Working paper, December 12, 2025, https://andreyfradkin.com/assets/LLM_Demand_12_12_2025.pdf.
31. On ChatGPT, see Sarah Jackson, “OpenAI Executives Say Releasing ChatGPT for Public Use Was a Last Resort After Running into Multiple Hurdles—and They’re Shocked by its Popularity,” *Business Insider*, January 25, 2023, <https://www.businessinsider.com/chatgpt-openai-executives-are-shocked-by-ai-chatbot-popularity-2023-1>. On Claude Code, Boris Cherney, the lead engineer on the project, later noted: “We released Claude Code like a year ago, and at the time, we weren’t sure if agentic coding was even going to be a thing. We had this hypothesis that maybe the model is ready for something like this. Almost immediately, it started to click.” Maxwell Zeff, “How Claude Code Is Reshaping Software—and Anthropic,” *Wired*, January 22, 2026, <https://www.wired.com/story/claude-code-success-anthropic-business-model/>.
32. Berber Jin and Jessica Toonkel, “The Sudden Fall of OpenAI’s Most Hyped Product Since ChatGPT,” *The Wall Street Journal*, March 29, 2026, <https://www.wsj.com/tech/ai/the-sudden-fall-of-openais-most-hyped-product-since-chatgpt-64c730c9>.
33. Berber Jin and Nate Rattner, “An Inside Look at OpenAI and Anthropic’s Finances Ahead of Their IPOs,” *The Wall Street Journal*, April 5, 2026, <https://www.wsj.com/tech/ai/openai-anthropic-ipo-finances-04b3c9b9>.
34. Multiple interviews with venture capital investors and employees working at AI labs, January 2026. Interviews were conducted in confidentiality, and the names of the interviewees are withheld by mutual agreement.
35. Jin and Rattner, “An Inside Look at OpenAI and Anthropic’s Finances.”
36. Jin and Rattner, “An Inside Look at OpenAI and Anthropic’s Finances.”
37. Fears that AI would destroy the SaaS market led to sharp market sell-off in early 2026; Lynn Doan and Carmen Reinicke, “What’s Behind the ‘SaaS Apocalypse’ Plunge in Software Stocks,” *Bloomberg*, February 4, 2026, <https://www.bloomberg.com/news/articles/2026-02-04/what-s-behind-the-saaspocalypse-plunge-in-software-stocks>.
38. For instance, roughly three-quarters of Alphabet/Google’s total revenues come from advertising. Alphabet, “Alphabet Announces Fourth Quarter and Fiscal Year 2025 Results,” press release, February 4, 2026, https://s206.g4cdn.com/479360582/files/doc_financials/2025/q4/2025q4-alphabet-earnings-release.pdf.
39. Tim Harford, “Does Pornography Still Drive the Internet?,” *BBC*, June 4, 2019, <https://www.bbc.com/news/business-48283409>.
40. Cristina Criddle and Stephen Morris, “OpenAI Puts Erotic Chatbot Plans on Hold ‘Indefinitely,’” *Financial Times*, March 26, 2026, <https://www.ft.com/content/de9bf0af-b241-424f-8229-5870b1c0d93d>.
41. See Brooke Tanner et al., *Is AI Sovereignty Possible? Balancing Autonomy and Interdependence* (Brookings Institution, February 17, 2026), <https://www.brookings.edu/articles/is-ai-sovereignty-possible-balancing-autonomy-and-interdependence/>.
42. Jason Ma, “Meet the Hedge Fund Manager Who Founded DeepSeek, the Chinese AI Startup that Began as a Hobby and is now Laying Waste to U.S. Stocks,” *Fortune*, January 27, 2025, <https://fortune.com/2025/01/27/deepseek-founder-liang-wenfeng-hedge-fund-manager-high-flyer-quant-trading/>.
43. “AI Use—National Average,” *Business Trends and Outlook Survey* data, U.S. Census Bureau, accessed March 27, 2026, <https://www.census.gov/hfp/btos/data>. Ramp, a fintech payments processing company, offers an alternative methodology for assessing AI uptake, which suggests significantly higher adoption among U.S. businesses of nearly 50 percent. Yet the Ramp sample is based on companies that have adopted the Ramp platform and thus is likely not representative of broader U.S. businesses. As Ramp’s own economist describes their data, “My preferred framing is not that Ramp businesses are reflective of the average American business, but that Ramp businesses are reflective of high-growth, fast-moving, tech-forward businesses.” AI adoption among tech-forward businesses will clearly tend to be higher than among a more representative sample of U.S. businesses. See Ara Kharazian, “How Ramp Data Works,” *Ramp*, April 3, 2026, <https://ramp.com/leading-indicators/how-ramp-data-works>.
44. Interviews with employees of companies in the energy and finance industries deploying AI, February 2026. Interviews were conducted in confidentiality, and the names of the interviewees are withheld by mutual agreement.
45. Anokhy Desai, “5 Things to Know About AI Model Cards,” *International Association of Privacy Professionals*, August 23, 2023, <https://iapp.org/news/a/5-things-to-know-about-ai-model-cards>.
46. A notable exception is Arthur, a platform for enterprise AI users to better monitor and evaluate their AI deployments. See <https://www.arthur.ai>.
47. Lee Harris and Christina Criddle, “Insurers Retreat from AI Cover as Risk of Multibillion-Dollar Claims Mounts,” *Financial Times*, November 23, 2025, <https://www.ft.com/content/abfe9741-f438-4ed6-a673-075ec177dc62>; Hunton Andrews Kurth LLP, “The Continued Proliferation of AI Exclusions,” *Hunton Insurance Recovery* (blog), accessed March 27, 2026, <https://www.hunton.com/hunton-insurance-recovery-blog/the-continued-proliferation-of-ai-exclusions>.
48. Companies beginning to offer insurance policies specifically for AI risks include the Artificial Intelligence Underwriting Company (<https://aiuc.com>) and Armilla (<https://www.armilla.ai>); see also discussion at Miranda Bogen, “What Will It Look Like to Insure Against AI Risks?” *Center for Democracy and Technology Insights* (blog), February 10, 2026, <https://cdt.org/insights/what-will-it-look-like-to-insure-against-ai-risks/>.
49. This is already a known problem in other consumer-facing connected devices, for instance, home internet routers, where compromised devices pose threats not only to individuals’ data but to broader network integrity, yet any individual consumer may have limited incentive to act on these broader risks. Michale Fagan et al., *Recommended Cybersecurity Requirements for Consumer-Grade Router Products* (U.S. National Institute for Standards and Technology, September 2024), <https://nvlpubs.nist.gov/nistpubs/ir/2024/NIST.IR.8425A.pdf>.
50. Alessandro Acquisti and Jens Grossklags, “Privacy and Rationality in Individual Decision Making,” *IEEE Security and Privacy* 3, no.1 (2005): 26–33, <https://dl.acm.org/doi/10.1109/MSP.2005.22>; Helen Nissenbaum, “A Contextual Approach to Privacy Online,” *Daedalus* 140, no. 4 (2011): 32–48, https://doi.org/10.1162/DAED_a_00113.
51. Matt Perault, “Section 230 Won’t Protect ChatGPT,” *Lawfare*, February 22, 2023, <https://www.lawfaremedia.org/article/section-230-wont-protect-chatgpt>.
52. Kashmir Hill, “Lawsuits Blame ChatGPT for Suicides and Harmful Delusions,” *The New York Times*, November 6, 2025, <https://www.nytimes.com/2025/11/06/technology/chatgpt-lawsuit-suicides-delusions.html>; Julie Jargon, “He Blames ChatGPT for the



- Murder-Suicide That Shattered His Family,” *The Wall Street Journal*, December 11, 2025, <https://www.wsj.com/tech/ai/chatgpt-murder-suicide-greenwich-openai-fd14fac2>.
53. Caleb Withers, Jay Kim, and Ethan Chiu, *Off Target: A Working Paper on AI Alignment Challenges for National Security* (Center for a New American Security, March 24, 2026), <https://www.cnas.org/publications/reports/off-target>.
54. “How Anthropic-Pentagon Dispute over AI Safeguards Escalated,” Reuters, March 11, 2026, <https://www.reuters.com/world/how-anthropic-pentagon-dispute-over-ai-safeguards-escalated-2026-03-11/>.
55. NVIDIA Corporation, “NVIDIA Announces Financial Results for Fourth Quarter and Fiscal 2025,” press release, February 26, 2025, <https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-fourth-quarter-and-fiscal-2025>.
56. The aggregate net profit margin for the S&P500 over the last decade is 11 percent; John Butters, “S&P 500 CY 2025 Earnings Preview: Analysts Expect Earnings Growth of 12.1%,” FactSet Insight, December 15, 2025, <https://insight.factset.com/sp-500-cy-2025-earnings-preview-analysts-expect-earnings-growth-of-12.1>.
57. Chris McGuire, *China's AI Chip Deficit: Why Huawei Can't Catch Nvidia and US Export Controls Should Remain* (Council on Foreign Relations, December 15, 2025), <https://www.cfr.org/articles/chinas-ai-chip-deficit-why-huawei-cant-catch-nvidia-and-us-export-controls-should-remain>.
58. Anthropic, “Expanding Our Use of Google Cloud TPUs and Services,” press release, October 23, 2025, <https://www.anthropic.com/news/expanding-our-use-of-google-cloud-tpus-and-services>.
59. Intel Corporation, “Intel Reports Fourth-Quarter and Full-Year 2025 Financial Results,” press release, January 26, 2026, <https://www.sec.gov/Archives/edgar/data/50863/000005086326000009/q425earningsrelease.htm>; Luke James, “Rapidus Targets Mass 2nm Chip Production in 2027, Quadruples Capacity Ramp up—Company Plans to Scale to 25,000 Wafer Starts per Month in Just One Year,” Tom's Hardware, <https://www.tomshardware.com/tech-industry/semiconductors/rapidus-targets-2nm-mass-production-in-2027-with-a-four-times-capacity-ramp>; and Reuters, “TSMC CEO Flags 3-Nanometre Chip Production in Japan, Investment Reported at \$17 Billion,” February 4, 2026, <https://www.reuters.com/world/asia-pacific/tsmc-plans-3-nanometre-chip-production-japan-with-17-billion-investment-yomiuri-2026-02-04/>.
60. “Worldwide Market Share of Leading Cloud Infrastructure Service Providers,” Statista, accessed March 27, 2026, <https://www.statista.com/chart/18819/worldwide-market-share-of-leading-cloud-infrastructure-service-providers/>.
61. There is, however, some differentiation between various cloud providers, particularly among companies tailoring their services for AI. For an overview ranking of different cloud providers on various quality dimensions. Jordan Nanos et al., “ClusterMAX™ 2.0: The Industry Standard GPU Cloud Rating System,” SemiAnalysis, November 6, 2025, <https://newsletter.semianalysis.com/p/clustermax-20-the-industry-standard>.
62. Mike Isaac, Eli Tan, and Cade Metz, “A.I. Researchers Are Negotiating \$250 Million Pay Packages: Just Like N.B.A. Stars,” *The New York Times*, July 31, 2025, <https://www.nytimes.com/2025/07/31/technology/ai-researchers-nba-stars.html>.
63. Jon Schmid, Tobias Systma, and Anton Shenk, *Evaluating Natural Monopoly Conditions in the AI Foundation Model Market* (RAND Corporation, September 12, 2024), https://www.rand.org/pubs/research_reports/RR3415-1.html.
64. OpenAI letter to the US House Select Committee, February 12, 2026, https://assets.bwbx.io/documents/users/iqjWHBFdfxIU/rRmqI_jCxb4/v0; Anthropic, “Detecting and Preventing Distillation Attacks,” press release, February 23, 2026, <https://www.anthropic.com/news/detecting-and-preventing-distillation-attacks>; and Google Threat Intelligence Group, “GTIG AI Threat Tracker: Distillation, Experimentation, and (Continued) Integration of AI for Adversarial Use,” press release, February 12, 2026, <https://cloud.google.com/blog/topics/threat-intelligence/distillation-experimentation-integration-ai-adversarial-use>.
65. Michael Kratsios, “Memorandum for the Heads of Executive Departments and Agencies: Adversarial Distillation of American AI Models,” memorandum, April 23, 2026, <https://whitehouse.gov/wp-content/uploads/2026/04/NSTM-4.pdf>.
66. Kate Clark, “These Billion-Dollar AI Startups Have No Products, No Revenue and Eager Investors,” *The Wall Street Journal*, January 27, 2026, <https://www.wsj.com/tech/ai/these-billion-dollar-ai-startups-have-no-products-no-revenue-and-eager-investors-97c0a9ba>.
67. Multiple interviews with venture capital investors and employees working at AI labs, January 2026. Interviews were conducted in confidentiality, and the names of the interviewees are withheld by mutual agreement.
68. Demirer et al., “The Emerging Market for Intelligence.” Note that this analysis is based on API usage data from OpenRouter and Microsoft Azure.
69. Richard Waters, “AI Upheaval Forces Software Industry to Ask if this is an ‘Adapt or Die’ Moment,” *Financial Times*, February 24, 2026, <https://www.ft.com/content/8867bff7-8632-40b7-8fde-7c95b3e624f9>.
70. Kyle Chan, “China is Running Multiple AI Races,” Brookings Institution Commentary, March 9, 2026, <https://www.brookings.edu/articles/china-is-running-multiple-ai-races/>.
71. Marina Temkin, “Nvidia's AI Empire: A Look at Its Top Startup Investments,” TechCrunch, January 2, 2026, <https://techcrunch.com/2026/01/02/nvidias-ai-empire-a-look-at-its-top-startup-investments/>.
72. Cedric Sam et al., “A Guide to the Circular Deals Underpinning the AI Boom,” Bloomberg, March 11, 2026, <https://www.bloomberg.com/graphics/2026-ai-circular-deals/>.
73. Oscar Delaney et al., *Strategic Visions in AI Governance* (Institute for AI Policy and Strategy, January 29, 2026), <https://www.iaps.ai/research/strategic-visions-in-ai-governance>; Barry Pavel et al., *How Artificial General Intelligence Could Affect the Rise and Fall of Nations: Visions for Potential AGI Futures* (RAND Corporation, July 2, 2025), https://www.rand.org/pubs/research_reports/RR3034-2.html.
74. See related discussion in Dario Amodei, “The Adolescence of Technology,” *Dario Amodei* (blog), January 2026, <https://www.darioamodei.com/essay/the-adolescence-of-technology>.
75. Jan Leike, “Distinguishing Three Alignment Taxes,” Musings on the Alignment Problem (Substack), December 19, 2022, <https://aligned.substack.com/p/three-alignment-taxes>; Amanda Askill, Miles Brundage, and Gillian Hadfield, “The Role of Cooperation in Responsible AI Development,” preprint, *arXiv:1907.04534*, July 2019, <https://arxiv.org/pdf/1907.04534>.
76. Anthropic has cited a figure of “close to 5% of total inference costs” (Amodei, “The Adolescence of Technology”). OpenAI has cited a figure of “as high as 16%” (OpenAI, “Introducing gpt-oss-safeguard,” OpenAI, October 29, 2025, <https://openai.com/index/introducing-gpt-oss-safeguard/>).
77. In interviews conducted for this paper (January–February 2026) both outside AI experts and market participants in venture capital and working at AI labs suggested oligopolistic competition was the most likely equilibrium outcome for frontier AI models. Interviews were conducted in confidentiality, and the names of the interviewees are withheld by mutual agreement.

78. Stuart Armstrong, Nick Bostrom, and Carl Shulman, "Racing to the Precipice: A Model of Artificial Intelligence Development," *AI & Society*, no. 31 (2016): 201–206, <https://link.springer.com/article/10.1007/s00146-015-0590-y>; Dan Hendrycks, Mantas Mazeika, and Thomas Woodside, "An Overview of Catastrophic AI Risks," preprint, *arXiv:2306.12001*, October 9, 2023, <https://arxiv.org/abs/2306.12001>.
79. Askill, "The Role of Cooperation in Responsible AI Development."
80. Ryan Fedasiuk, *China's Transition to Scalable Intelligence* (American Enterprise Institute, February 18, 2026), <https://www.aei.org/research-products/working-paper/chinas-transition-to-scalable-intelligence/>.
81. For an overview of these threats, see Daniel Remler, *Red Lines: Understanding the National Security Risks of China's Advanced AI* (Center for a New American Security, forthcoming).
82. NIST, "CAISI Evaluation of DeepSeek AI Models Finds Shortcomings and Risks," press release, September 30, 2025, <https://www.nist.gov/news-events/news/2025/09/caisi-evaluation-deepseek-ai-models-finds-shortcomings-and-risks>; Remler, *Red Lines*.
83. Elizabeth Seger et al., *Open-Sourcing Highly Capable Foundation Models: An Evaluation of Risks, Benefits, and Alternative Methods for Pursuing Open-Source Objectives* (Center for the Governance of AI, 2023), https://law-ai.org/wp-content/uploads/2023/10/Open-Sourcing_Highly_Capable_Foundation_Models_2023_GovAI-1.pdf.
84. Dan Gallagher, "AI Won't Kill the Software Business, Just Its Growth Story," *The Wall Street Journal*, February 4, 2026, <https://www.wsj.com/tech/ai/ai-wont-kill-the-software-business-just-its-growth-story-05673e07>.
85. Notably, an export controls-based strategy may not work as well applied at the model or application layer rather than the compute layer. Controls so far have been applied predominantly to tangible, physical goods (i.e., chips), which are inherently more amenable to controls than are pieces of software or intellectual property, which only need to be transferred once to nullify the impact of a control.
86. For more details on such a policy agenda, see Janet Egan, Spencer Michaels, and Caleb Withers, *Prepared, Not Paralyzed: Managing AI Risks to Drive American Leadership* (Center for a New American Security, November 20, 2025), <https://www.cnas.org/publications/reports/prepared-not-paralyzed>.
87. The Federal Trade Commission (FTC) previously launched an inquiry into AI chatbots acting as companions to better understand how such companies measure, test, and monitor potential negative impacts of their apps on children and teens. Federal Trade Commission, "FTC Launches Inquiry into AI Chatbots Acting as Companions," press release, September 11, 2025, <https://www.ftc.gov/news-events/news/press-releases/2025/09/ftc-launches-inquiry-ai-chatbots-acting-companions>.
88. In December 2025, the U.S. Securities and Exchange Commission (SEC) Investor Advisory Committee proposed a recommendation that the SEC issue further guidance to standardize AI-related disclosures for companies using AI; however, at present, the SEC appears unlikely to take up this proposal. "Recommendation of the SEC Investor Advisory Committee Regarding Disclosure of Artificial Intelligence's Impact on Operations", approved by the Investor Advisory Committee at the December 4, 2025 Meeting, <https://www.sec.gov/files/approved-artificial-intelligence-disclosure-recommendation-120425.pdf>.
89. Ketan Ramakrishnan, Gregory Smith, and Conor Downey, *U.S. Tort Liability for Large-Scale Artificial Intelligence Damages: A Primer for Developers and Policymakers* (RAND, August 21, 2024), https://www.rand.org/pubs/research_reports/RRA3084-1.html.
90. The former Trump administration AI official Dean Ball has previously called for governments to license private AI standards-setting and regulatory organizations to provide certifications for AI developers; Dean Ball, "Putting Private AI Governance into Action," *Hyperdimensional* (Substack), March 20, 2025, <https://www.hyperdimensional.co/p/putting-private-governance-into-action>.
91. Lina Khan, "The New Brandeis Movement: America's Antimonopoly Debate," *Journal of European Competition Law & Practice* 9, no. 3 (2018): 131–132, <https://academic.oup.com/jeclap/article/9/3/131/4915966>.
92. Tejas N. Narechania and Ganesh Sitaraman, "An Antimonopoly Approach to Governing Artificial Intelligence," *Yale Law & Policy Review* 43, no. 1 (2024), <https://yalelawandpolicy.org/antimonopoly-approach-governing-artificial-intelligence>.
93. Interview with an AI industry expert, February 2026. The interview was conducted in confidentiality, and the name of the interviewee is withheld by mutual agreement. Shirin Ghaffary and Maggie Eastland, "OpenAI, Anthropic, Google Unite to Combat Model Copying in China," *Bloomberg*, April 6, 2026, <https://www.bloomberg.com/news/articles/2026-04-06/openai-anthropic-google-unite-to-combat-model-copying-in-china>, which notes "For now, information sharing on distillation remains limited due to AI companies' uncertainty about what can be shared under existing antitrust guidance to counter the competitive threat from China, according to people familiar with the matter."
94. Kratsios, "Memorandum for the Heads of Executive Departments and Agencies."
95. "China Is Quietly Upstaging America with Its Open Models: How Worried Should OpenAI and Other Labs Be?," *The Economist*, August 21, 2025, <https://www.economist.com/business/2025/08/21/china-is-quietly-upstaging-america-with-its-open-models>.
96. For more on these and other related policy recommendations related to national security threats associated with Chinese AI models, see Remler, *Red Lines*.
97. There is some evidence that the U.S. Commerce Department intends to pursue such a strategy in its AI Exports Plan. Amrith Ramkumar, "Nvidia-Backed AI Startup to Spend Billions on Korea Data Center to Combat China," *The Wall Street Journal*, March 17, 2026, <https://www.wsj.com/tech/ai/nvidia-backed-ai-startup-to-spend-billions-on-korea-data-center-to-combat-china-f945a326>.