BEIJING'S AI PLAYBOOK

A Joint Investigative Report

★ ★ ★ ★ ★

How AI Models Developed in the People's Republic of China Became Embedded in America's Technology Stack and What It Will Take to Avoid the Next Huawei

U.S. HOUSE OF REPRESENTATIVES
COMMITTEE ON HOMELAND SECURITY
CHAIRMAN ANDREW R. GARBARINO
AND

THE SELECT COMMITTEE ON THE STRATEGIC COMPETITION BETWEEN THE UNITED STATES AND THE CHINESE COMMUNIST PARTY
CHAIRMAN JOHN MOOLENAAR

TABLE OF CONTENTS

FOREWORD FROM THE CHAIRMEN 2

EXECUTIVE SUMMARY 6

THE INVESTIGATION 12

FINDINGS 15

Finding 1. PRC-developed open-weight models are already widely used across the American economy, but much of that use is hard to track. 15

Finding 2. Cost, customization, and near-frontier performance are driving adoption, and the United States lacks a competitive open-weight alternative at the price and performance the market demands. 16

Finding 3. Much of the PRC's apparent progress rests on capabilities extracted from American models, and American firms are helping supply the data that trains PRC models. 20

Finding 4. Open-weight distribution is a deliberate PRC strategy to expand global dependence, backed by state policy and capital, and the monetization phase has begun. 22

Finding 5. The risks are increasing as PRC models become more capable, while their behavior and security properties can differ in ways ordinary evaluations may not detect. 29

RECOMMENDATIONS 34

A. Build the American Alternative 35

B. Know What Is Running 36

C. Share Threat Information and Strengthen Defenses 39

D. Impose Costs on Illicit Extraction 44

E. Set Durable Rules and Work with Allies 45

CONCLUSION 46

ANNEX 47

Annex A. Zhipu and the PLA Cyberspace Force 47

Annex B. State Support for Chinese AI Developers 47

CLASSIFIED ANNEX 49

FOREWORD FROM THE CHAIRMEN

The United States has seen this movie before. A foreign technology enters the market at a price American competitors cannot match and with performance that is good enough. It spreads because it solves real problems for real customers. By the time Washington understands that the technology was developed under the authority of the Chinese Communist Party (CCP) and advances the strategic interests of the People's Republic of China (PRC), it is already deeply embedded in our networks, products, critical infrastructure, and public services across the country. Removing it then takes years and costs American taxpayers billions of dollars.

Huawei and ZTE followed that path. In 2012, the House Permanent Select Committee on Intelligence warned that their equipment posed a direct threat to American networks. Yet carriers, mostly small and rural, kept buying it because it was affordable and it worked. Years later, Congress had to fund a $5 billion program to "rip and replace" the equipment. DJI followed a similar course, flooding the U.S. market with cheap drones that became widely used by federal agencies, state and local police departments, farmers, construction companies, surveyors, and critical infrastructure operators managing power grids, water systems, and energy pipelines. By the time the federal government moved to restrict new foreign-made drone models and components in December 2025, DJI had become a staple across American industries. In each case, the warning came early. The response came far too late.

The Committee on Homeland Security and the Select Committee on the CCP (the Committees) launched this joint investigation in April 2026 because we believe the same pattern is emerging in artificial intelligence (AI), and this time the stakes are even higher. The technology at issue is not a radio on a cell tower or a camera on a drone that can be physically removed and replaced. An AI model can be downloaded once, adapted, renamed, and passed into other products, each carrying the original model's training and behavior without disclosing its origin. Tracing where those derivatives end up is far harder than locating a piece of equipment, and replacing them can take months of work.

Most Americans encounter AI every single day without knowing it: when a pharmacy texts that a prescription is ready, when a chatbot helps them rebook a trip, when an app estimates when their dinner will arrive, or when their bank texts them about an unusual purchase. Behind those interactions, AI models are summarizing customer records, routing support tickets, and executing automated tasks deep inside corporate networks. When a company brings an AI model into its operations, it also brings in the influence of those who developed and trained it, including their priorities, assumptions, and choices that shape how it processes information, generates responses, and carries out tasks. Increasingly, those developers answer to Beijing.

Over the past seven months, the Committees examined how and why American companies are adopting PRC-developed AI models. The investigation included document requests and briefings from Airbnb, the short-term rental platform; Anysphere, the developer of the Cursor AI coding platform; and DoorDash, the food delivery platform, along with review of company productions, written responses, and internal evaluation records. The Committees also received classified briefings, consulted with frontier AI developers and cybersecurity executives, and held closed-door roundtables with federal agencies, industry representatives, and academic researchers. The review drew on an expanding body of public and government research, including a September 2026 joint advisory issued by the National Security Agency (NSA), the Cybersecurity and Infrastructure Security Agency (CISA), and the Federal Bureau of Investigation (FBI).

The investigation uncovered risks that American businesses, policymakers, and national security officials can no longer afford to overlook or downplay. From Silicon Valley startups to Fortune 500 companies, American businesses are increasingly adopting PRC-developed open-weight models because they are inexpensive, customizable, and competitive with leading American systems. Federal investigations, however, have found that major PRC developers used proxy networks and fraudulent accounts to harvest outputs from American frontier models and replicate their capabilities. While America is investing billions and racing to build "Super Intelligence," Beijing is using stolen intelligence to close the gap.

Beijing is now seeking to turn those technological gains into global market share and long-term dependence. The PRC State Council has explicitly directed the construction of an open-source AI ecosystem with "international influence." Researchers at a state academy have described the approach as combining "low-cost technology exports" with "high-stickiness ecosystem binding" to create lasting dependence on PRC technology. Leading PRC developers are sustaining heavy losses with support from state investment funds, state-owned shareholders, and subsidized compute, while offering models at prices American developers struggle to match. As enterprise adoption expands, however, the commercial terms are shifting. New licenses require some large commercial users to negotiate separate agreements, and Beijing has reportedly considered restricting overseas access to its most advanced models. What begins as cheap to adopt can become costly to depend on and harder still to replace.

Commercial dependence is only part of the risk. PRC developers are making increasingly powerful cyber capabilities freely available. In September 2026, an American frontier AI developer reported that the newest open-weight model from Zhipu, a company on the Commerce Department's Entity List, could autonomously build complete cyber exploits at a level comparable to Mythos Preview, an American model deemed too dangerous for general release just five months earlier. The National Institute of Standards and Technology (NIST)

identified Zhipu's model as the most cyber capable open-weight model released to date. Because its weights are publicly available, safeguards can be stripped away for a few thousand dollars in compute. The assumption that PRC models remain safely six months behind the American frontier no longer holds in critical areas of national security, particularly cybersecurity.

South Korea offers an early warning. Over a matter of days in late September and early October, several of the country's largest banks and other financial institutions were breached. According to public reports, the attacker used an open-source AI penetration testing tool developed in the PRC, powered by a DeepSeek model and supplemented by others, including the same Zhipu model that NIST identified as the most cyber capable open-weight model released to date. If those reports are correct, these tools allowed a single attacker to carry out multiple intrusions in far less time than a human team would need. The same tools are available to anyone targeting American banks, hospitals, and utilities. A company does not have to adopt a PRC model to become a victim of someone using one.

These findings call for a response that distinguishes among developers, capabilities, and uses, with the strongest measures reserved for federal systems and critical infrastructure. The federal government has already begun to act. In October 2026, the Department of Homeland Security (DHS) made public an order, issued in July, that excludes DeepSeek products and services from federal civilian agency procurement and requires their removal from agency information systems, including third-party products built on DeepSeek models. That order is an important step, but it reaches only one developer. Other leading PRC developers are subject to the same national security laws, were identified in the same distillation campaigns, and have released models with comparable or greater offensive cyber capability. The federal government should extend the same restrictions to those developers.

Our position is straightforward. PRC-developed models should be removed from federal systems, and federal contractors should not use them to perform federal work or handle federal data. They should also be removed from high-risk functions of critical infrastructure, where a compromised or manipulated model could cause physical harm or disrupt essential services. Any exception should be temporary, limited to cases where no trusted substitute yet exists, and approved by DHS, acting through CISA, after an evaluation by the Center for Advancing Innovation and Standards for Super Intelligence (CAISSI). Running a model on American servers keeps data out of the PRC, but it does not change how the model was built, what it was trained to do, or how it behaves under conditions no one thought to test. Across the rest of the economy, the goal is to reduce dependence on PRC models as quickly as trusted American and allied alternatives can replace them. Most companies do not need the most advanced and expensive frontier systems for everyday work, and many turned to PRC models because they perform well at a fraction of the cost. Forcing those companies off PRC models before trusted open-weight alternatives exist would raise their costs and

concentrate the market in a small number of closed, proprietary providers. Building those alternatives is not a substitute for restrictions. It is what makes them work.

To achieve these goals, the report recommends building a competitive American open-weight ecosystem capable of becoming the global default, model provenance standards to identify what is running inside critical systems, and long-term certainty for the Cybersecurity Information Sharing Act of 2015 (CISA 2015) and the State and Local Cybersecurity Grant Program (SLCGP). It also proposes a dedicated DHS channel to share developer-specific threat information with industry, stronger measures to hold those who steal American AI technology accountable, a prohibition on PRC-developed models in federal systems and federal contract work that builds on the action already taken against DeepSeek, a phased prohibition on their use in high-risk functions of critical infrastructure, a bar on using federal cybersecurity grants to buy them, clear guidance that steers critical infrastructure operators toward trusted alternatives, and transition planning now, while the cost of switching remains manageable.

There are encouraging signs. American and allied developers are beginning to release more competitive open-weight models, including new models from Reflection AI in the United States and Mistral AI in France released in October 2026. Common sense policy can build on that momentum and help ensure American innovation continues to lead.

Congress has too often been behind the curve on emerging technology. As one of our colleagues observed after a briefing on these issues, AI is advancing so quickly that Congress risks being "light years behind." With this investigation, we intend to narrow that gap. We should not wait five years to discover that much of the world, and much of the United States, runs on AI shaped by the CCP, leaving policymakers with little choice but another costly rip and replace effort. The time to act is now, while we still have the opportunity and leverage to do so.

The Foreword's conclusions are supported by the investigative findings and recommendations that follow. A classified annex accompanies this report and is available to authorized personnel with appropriate security clearances.





Andrew R. Garbarino
Chairman
Committee on Homeland Security

John Moolenaar
Chairman
Select Committee on China

EXECUTIVE SUMMARY

In April 2026, the Committees opened a joint investigation into the national security, cybersecurity, and economic security implications of the growing use of PRC-developed AI models by American companies.[1] The investigation focused on why American companies are adopting these models, how PRC developers acquired the capabilities that made them competitive with leading American systems, the role of PRC government support in their development and distribution, and the risks associated with growing American dependence on them.[2] As part of that investigation, the Committees examined three U.S. companies (Airbnb, Anysphere, and DoorDash) as case studies of how and why PRC-developed models are already being incorporated into American products and operations.[3]

- **Airbnb** runs customized versions of Alibaba's open-weight Qwen models on U.S.-based service provider infrastructure and uses them in customer support, including for an AI assistant handling roughly 80,000 guest and host interactions per day and a fine-tuned model used to summarize roughly 250,000 support cases and calls each day.[4]
- **Anysphere** built Composer 2, the flagship proprietary model in its Cursor coding platform, on a Moonshot AI base model, a fact that became public only after an independent developer identified the underlying model.[5]
- **DoorDash** reported that its engineers had downloaded or evaluated PRC-developed models from at least six developers, that 25 PRC model families had been authorized and deployed in a non-trivial manner, which it defined as incurring at least $5,000 in computing and serving costs, and that replacing one PRC-developed model embedded in a larger system could take up to six months, although most of its other uses could be transitioned almost immediately.[6]

The experiences of Airbnb, Anysphere, and DoorDash are illustrative, not exhaustive. Their representations to the Committees show different ways PRC-

[1] Press Release, H. Comm. on Homeland Sec., *Chairmen Garbarino, Moolenaar Announce Joint Investigation into National Security Risks Posed by PRC AI Models* (Apr. 29, 2026) [hereinafter Launch Press Release].
[2] Press Release, H. Comm. on Homeland Sec., *Chairmen Garbarino, Moolenaar Continue Joint Investigation into Security Risks Posed by PRC Open-Weight AI Models* (July 31, 2026) [hereinafter DoorDash Press Release].
[3] Letter from Andrew R. Garbarino, Chairman, H. Comm. on Homeland Sec., & John Moolenaar, Chairman, H. Select Comm. on the Strategic Competition Between the U.S. & the Chinese Communist Party, to Brian Chesky, Chief Exec. Officer, Airbnb, Inc. (Apr. 29, 2026) [hereinafter Airbnb Letter]; Letter from Andrew R. Garbarino, Chairman, H. Comm. on Homeland Sec., & John Moolenaar, Chairman, H. Select Comm. on the Strategic Competition Between the U.S. & the Chinese Communist Party, to Michael Truell, Chief Exec. Officer, Anysphere, Inc. (Apr. 29, 2026) [hereinafter Anysphere Letter]; Letter from Andrew R. Garbarino, Chairman, H. Comm. on Homeland Sec., & John Moolenaar, Chairman, H. Select Comm. on the Strategic Competition Between the U.S. & the Chinese Communist Party, to Tony Xu, Co-Founder & Chief Exec. Officer, DoorDash, Inc. (July 31, 2026) [hereinafter DoorDash Letter].
[4] Airbnb, Inc., *Customer Support GenAI: Model Portfolio and Stage-Gated Optimization Strategy*, *in* ABNB_HSCCCP_2026_0000118–0000180 (produced to the Committees) [hereinafter Airbnb Customer Support Strategy]; Email from Airbnb, Inc., to Majority Staff, H. Comm. on Homeland Sec. (Nov. 13, 2026) [hereinafter Airbnb Nov. 13 Email].
[5] Cursor Research Team, *Composer 2 Technical Report* (Mar. 26, 2026), produced as ANYSP-HSC-00000021–00000043 [hereinafter Composer 2 Technical Report]; Mannat Dora, *Cursor Founder Clears Air on Kimi Model Use in Composer 2: Here's All You Need to Know*, Econ. Times (Mar. 21, 2026).
[6] Letter from DoorDash, Inc., to John Moolenaar, Chairman, H. Select Comm. on the Strategic Competition Between the U.S. & the Chinese Communist Party, & Andrew R. Garbarino, Chairman, H. Comm. on Homeland Sec. (Aug. 28, 2026) [hereinafter DoorDash Response]; DoorDash, Inc., Document Production in Response to the Committees' Letter of July 31, 2026, DD-HSCCCP-HHSC-00001–00044, attachs. A, C (produced to the Committees) [hereinafter DoorDash Production]; Briefing by DoorDash, Inc. to Majority Staff of the Committees (Sept. 28, 2026) [hereinafter DoorDash Briefing]; Email from DoorDash, Inc., to Majority Staff, H. Comm. on Homeland Sec. (Nov. 14, 2026) [hereinafter DoorDash Nov. 14 Email].

developed models are entering American products and operations, but the Committees' broader investigation found that adoption extends well beyond these three companies and is being driven by forces that are reshaping the wider AI market.[7] The investigation produced five main findings:

Finding 1. PRC-developed open-weight models are already widely used across the American economy, but much of that use is hard to track. Open-weight models now carry a majority of U.S.-originating traffic on at least one major model marketplace, and PRC models account for most of that volume.[8] PRC models also accounted for 41 percent of downloads on the leading open-model repository over the preceding year.[9] Because downloaded weights do not require a commercial relationship with the developer, public metrics capture only part of their actual use, and neither customers nor the federal government can reliably see where PRC models are deployed.[10]

Finding 2. Cost, customization, and near-frontier performance are driving adoption, and the United States lacks a competitive open-weight alternative at the price and performance the market demands. Leading PRC models cost a small fraction of leading American proprietary models at published rates.[11] In the most recent independent ranking of open-weight model capability reviewed by the Committees, the highest-ranked American open-weight model placed tenth, behind six PRC models.[12] Companies examined by the Committees cited the lack of competitive American open-weight options as a principal reason for their choices.[13]

Finding 3. Much of the PRC's apparent progress rests on capabilities extracted from American models, and American data-labeling firms are helping supply the data that trains PRC models. U.S. agencies assessed that six leading PRC developers conducted industrial-scale distillation campaigns against American frontier models, likely with PRC government awareness.[14] American data-labeling firms reportedly earn approximately $500 million a year selling American training data to PRC laboratories.[15] The Committees assess with high

---

[7] *See* Cailee Moberg, *In-Region Routing: Keep Your Data in the US or EU*, OpenRouter (Sept. 9, 2026), https://perma.cc/A38F-GU8T; Hugging Face, *State of Open Source on Hugging Face: Spring 2026* (Mar. 17, 2026), https://perma.cc/GVW5-CEWM [hereinafter Hugging Face Spring 2026 Report]; Dean Meyer & Konstantine Buhler, *America's Open-Model Paradox*, Sequoia Capital (July 24, 2026), https://perma.cc/L45D-P6BF.

[8] Moberg, *supra* note 7.

[9] Hugging Face Spring 2026 Report, *supra* note 7.

[10] *See DeepSeek Now Handles 17% of Vercel Gateway Tokens at 1% of Spend*, AI Insiders (June 11, 2026); Tim Tully, Joff Redfern, Deedy Das & Derek Xiao, *2025: The State of Generative AI in the Enterprise*, Menlo Ventures (Dec. 9, 2025), https://perma.cc/K3GD-MYXU; Amy Chang, *The "U.S. vs. China" AI Trap: An Incomplete Proxy for AI Security*, Cisco Blogs (Aug. 27, 2026), https://perma.cc/7UFX-WEWV; Ivica Nikolic, Teodora Baluta & Prateek Saxena, *Model Provenance Testing for Large Language Models*, Advances in Neural Info. Processing Sys. (2025), https://perma.cc/DY6Y-6333.

[11] *See* Alibaba Cloud, *qwen3-235b-a22b-instruct-2507 Model Info*, https://perma.cc/JDA8-9Z9G [hereinafter Alibaba Cloud Pricing]; DeepSeek, *DeepSeek-R1 Release* (Jan. 20, 2025), https://perma.cc/274F-JF4L [hereinafter DeepSeek-R1 Release]; Kimi, *Kimi API Pricing and Billing Guide*, https://perma.cc/S9VM-LDMY (last visited Oct. 7, 2026) [hereinafter Kimi Pricing Guide]; Anthropic, *Introducing Claude Opus 4.8* (May 28, 2026), https://perma.cc/6R7G-LCYE [hereinafter Claude Opus 4.8 Announcement]; OpenAI, *GPT-5.5 Model*, OpenAI API Documentation, https://perma.cc/G2XJ-2UTL (last visited Sept. 30, 2026) [hereinafter GPT-5.5 Model Page].

[12] Artificial Analysis, *Open Weights Models: Intelligence Index* (Sept. 2026), https://perma.cc/A5BS-BP86 [hereinafter Artificial Analysis Index].

[13] *See* DoorDash Response, *supra* note 6; Airbnb Customer Support Strategy, *supra* note 4.

[14] Nat'l Sec. Agency, Cybersecurity & Infrastructure Sec. Agency & Fed. Bureau of Investigation, *China-Based Artificial Intelligence Companies Conducting Industrial-Scale Distillation Campaigns Against U.S. AI Companies* 1–2 (Joint Cybersecurity Advisory, TLP:CLEAR, Sept. 8, 2026) [hereinafter Joint Distillation Advisory].

[15] Anna Tong, *Silicon Valley's Other China Problem: It's Training Their AI*, Forbes (Aug. 5, 2026, updated Aug. 19, 2026).

confidence that illicit distillation materially accelerated the development of several PRC model families now used by American companies.[16]

Finding 4. Beijing is using open-weight distribution, state support, and low prices to expand global dependence on PRC models, while the terms of that dependence are beginning to change. PRC policy documents call for open-source AI with "international influence" and describe the goal of "high-stickiness ecosystem binding."[17] Leading developers sustain heavy losses with support from state funds, state-owned investors, and subsidized compute.[18] Developers are now introducing commercial license terms and revenue demands, while Beijing has reportedly considered restricting overseas access to its most advanced models.[19]

Finding 5. The risks are increasing as PRC models become more capable, while their behavior and security properties can differ in ways ordinary evaluations may not detect. Independent researchers found that some PRC models produced less secure code when prompts referenced politically sensitive topics or a U.S. government user.[20] At the same time, NIST and an American frontier developer assessed that Zhipu's newest open-weight model can autonomously build end-to-end cyber exploits at a level comparable to a leading restricted American model and trails the American frontier by about four months on aggregate government cyber benchmarks.[21] Its safeguards were reduced from roughly 95 percent refusal rates to between 3 and 6 percent for an estimated $4,400.[22]

---

[16] *See* Joint Distillation Advisory, *supra* note 14, at 1–2; Anthropic, *Detecting and Preventing Distillation Attacks* (Feb. 23, 2026) [hereinafter Anthropic Distillation Report]; Memorandum from OpenAI to the H. Select Comm. on the Strategic Competition Between the U.S. & the Chinese Communist Party (Feb. 12, 2026) [hereinafter OpenAI Memorandum]. The Committees express their assessments using confidence levels consistent with the analytic standards of the Intelligence Community. *See* Off. of the Dir. of Nat'l Intel., Intelligence Community Directive 203, *Analytic Standards* (Jan. 2, 2015), https://www.dni.gov/files/documents/ICD/ICD-203.pdf [hereinafter ICD 203].

[17] State Council, *Opinions on Deeply Implementing the "Artificial Intelligence+" Action* [关于深入实施"人工智能+"行动的意见] § III(11) (Aug. 26, 2025), https://perma.cc/5FEZ-T7JN [hereinafter State Council AI+ Opinions]; Chen Xiaodong & Zhu Xiaofei, *Six Dimensions of DeepSeek's Technological Breakthrough Reshaping the Global Artificial Intelligence Industry Competitive Landscape: Paradigms and Implications* [DeepSeek 技术突围重构全球人工智能产业竞争格局的六个维度：范式与启示], Chinese Acad. of Soc. Scis., Inst. of Indus. Econ. (Aug. 6, 2025), https://perma.cc/6KJF-LAZ4.

[18] *See* Kinling Lo, *China's MiniMax, Zhipu AI Beat OpenAI to IPO*, Rest of World (Jan. 6, 2026), https://perma.cc/PSH8-RCYU; *DeepSeek Completes First Financing Round; Tencent, CATL, and National AI Industry Investment Fund All Invest* [DeepSeek首轮融资落地 腾讯、宁德时代、国家人工智能产业基金悉数入股], STAR Market Daily (July 15, 2026), https://perma.cc/GV25-KA4P; *MiniMax Lists on the Hong Kong Stock Exchange; Shanghai International Group and Shanghai State-Owned Capital Investment Help AGI Company Begin New Chapter* [MiniMax港交所挂牌上市，上海国际集团、上海国投公司助力AGI大模型企业开启新征程], Shanghai State-Owned Assets Supervision & Admin. Comm'n [上海市国有资产监督管理委员会] (Jan. 12, 2026), https://perma.cc/TQV5-K6LN; Kyle Chan et al., *Full Stack: China's Evolving Industrial Policy for AI*, RAND Corp. (June 26, 2025), https://perma.cc/R53B-TRFZ.

[19] Moonshot AI, *Kimi K3 License*, Hugging Face, https://perma.cc/4JVP-GL7Q [hereinafter Kimi K3 License]; Stephen Nellis & Eduardo Baptista, *Alibaba Plans to Charge Big Users of Its Next Open-Source AI Model, Sources Say*, Reuters (Aug. 7, 2026); Eduardo Baptista et al., *China's Moonshot in Talks with Microsoft, Amazon, Google over K3 Revenue Sharing, Sources Say*, Reuters (Aug. 26, 2026); Fanny Potkin, *Exclusive: Beijing Is Looking at Curbing Overseas Access to China's Top AI Models, Sources Say*, Reuters (July 7, 2026).

[20] Stefan Stein, *CrowdStrike Research: Security Flaws in DeepSeek-Generated Code Linked to Political Triggers*, CrowdStrike (Nov. 20, 2025), https://perma.cc/22TG-HZCP; Press Release, Booz Allen Hamilton, *New Booz Allen Analysis Reveals Risks in Using Chinese AI Models for America's Software Supply Chain* (June 5, 2026), https://perma.cc/LT4K-FM9A [hereinafter Booz Allen Press Release].

[21] Ctr. for AI Standards & Innovation, Nat'l Inst. of Standards & Tech., *Assessment of Z.ai's GLM-5.3 Cyber Capabilities* (Sept. 17, 2026), https://perma.cc/U6XS-DJCB [hereinafter NIST GLM-5.3 Assessment]; Anthropic, *GLM-5.3 and the Spread of Advanced Cyber Capabilities* (Sept. 29, 2026), https://perma.cc/M9TG-CA5F [hereinafter Anthropic GLM-5.3 Analysis].

[22] Anthropic GLM-5.3 Analysis, *supra* note 21.

The Committees' recommendations are intended to address those findings, build on the action the federal government has already taken against DeepSeek, and strengthen American innovation.[23] They are organized around five objectives:

Build the American alternative.

- **Recommendation 1.** Launch a national American open-weight AI initiative that pairs public compute, shared evaluation infrastructure, and lawful access to American frontier models for training with anchor demand from federal procurement.
- **Recommendation 2.** Make trusted American AI the global default through export promotion, development finance, and partnerships with allies building sovereign AI capacity.

Know which models are running in sensitive systems.

- **Recommendation 3.** Direct NIST, in coordination with DHS through CISA, to develop a model provenance standard and an AI bill of materials, required for federal contractors and high-risk critical infrastructure functions and voluntary elsewhere, with protections for small businesses and open-source developers.
- **Recommendation 4.** Build on the order excluding DeepSeek from federal civilian agency systems by prohibiting the use of PRC-developed AI models and their derivatives in federal information systems and in the performance of federal contracts, with a narrow and transparent waiver process for research, testing, and defensive security work, and modernize CISA's Continuous Diagnostics and Mitigation (CDM) program so agencies can discover and inventory the AI models and agents running on their networks. Prohibit PRC-developed models in high-risk functions of critical infrastructure, with removal schedules tied to the availability of trusted alternatives and temporary exceptions only where none yet exists.

Share threat information and strengthen defenses.

- **Recommendation 5.** Enact a long-term reauthorization of CISA 2015 through the Widespread Information Management for the Welfare of Infrastructure and Government (WIMWIG) Act,[24] which the Committee on Homeland Security approved unanimously and the House passed as part of the Fiscal Year 2027 National Defense Authorization Act (H.R. 8800).[25] Following enactment of a long-term reauthorization that extends and modernizes CISA 2015 authorities, Congress may consider whether additional narrow clarifications related to sharing information specifically about threats to AI systems are necessary, as those technologies continue

[23] *See* Notice of Order Regarding Exclusion and Removal of DeepSeek Products from Federal Civilian Executive Branch Information Technology Systems and Procurements, 91 Fed. Reg. 66,346 (Oct. 20, 2026); Jenny Lee & Kai Nicol-Schwarz, *U.S. Bans DeepSeek from Government Agency Systems amid AI Rift with China*, CNBC (Oct. 20, 2026).

[24] Widespread Information Management for the Welfare of Infrastructure and Government Act, H.R. 5079, 119th Cong. (2025), https://www.congress.gov/bill/119th-congress/house-bill/5079.

[25] National Defense Authorization Act for Fiscal Year 2027, H.R. 8800, 119th Cong. (as passed by House, July 22, 2026), https://www.congress.gov/bill/119th-congress/house-bill/8800; Martin Matishak, *Extension of CISA 2015 Info-Sharing Protections Passes as Part of House's Defense Bill*, The Record (July 22, 2026), https://perma.cc/K3HU-RW6P.

to evolve. Complete a long-term reauthorization of the SLCGP through the Protecting Information by Local Leaders for Agency Resilience (PILLAR) Act, which the House passed in November 2025,[26] and prohibit the use of program funds for PRC-developed AI models while allowing those funds to be used to replace them. Direct DHS, acting through CISA, to establish a channel for sharing developer-specific AI threat information with industry.

- **Recommendation 6.** Direct NIST's CAISSI to evaluate leading PRC models on a recurring basis, direct CISA to translate the results into guidance that steers critical infrastructure away from PRC models in sensitive functions, and codify DHS, acting through CISA, as the Sector Risk Management Agency for the Information Technology Sector, including AI models and AI compute infrastructure.
- **Recommendation 7.** Prepare DHS and CISA to defend federal networks and critical infrastructure at machine speed. Establish a DHS pilot program that brings private sector AI and cybersecurity expertise into the Department, with a dedicated allocation for CISA, and direct CISA to serve as the operational lead for securing AI agents on federal civilian networks.
- **Recommendation 8.** Make CISA the single point of submission for AI-related cyber incident reports by building on the Cyber Incident Reporting for Critical Infrastructure Act of 2022 (CIRCIA), which reaches the American entities best positioned to detect incidents involving PRC-developed models, rather than creating new, duplicative reporting requirements. Establish within CISA an independent, nonregulatory board to investigate the root causes of significant cyber incidents and cyber incidents involving AI, including incidents involving foreign adversary models and illicit capability extraction.

Impose costs on illicit extraction.

- **Recommendation 9.** Use sanctions, the Entity List, and information and communications technology supply chain authorities against firms engaged in industrial-scale illicit distillation and require identity verification standards that close the gray market of proxy "transfer stations."
- **Recommendation 10.** Examine the sale of American-produced frontier training data to PRC AI laboratories designated by the U.S. government and consider the imposition of disclosure requirements or export controls.

Set durable rules and work with allies.

- **Recommendation 11.** Ensure that federal AI security policy addresses both closed and open models, and establish a federal baseline for AI transparency and security that does not overburden startups, researchers, or nonprofits.

[26] Protecting Information by Local Leaders for Agency Resilience Act, H.R. 5078, 119th Cong. (as passed by House, Nov. 17, 2025), https://www.congress.gov/bill/119th-congress/house-bill/5078; Press Release, H. Comm. on Homeland Sec., *PILLAR Act Garners Broad Industry and Government Support, Passes House* (Nov. 18, 2025), https://perma.cc/HZT5-X5ND [hereinafter PILLAR Press Release].

- **Recommendation 12.** Expand joint testing and standards work with allies and direct the Intelligence Community, with DHS, to deliver an assessment of PRC state involvement in open-weight AI strategy.

THE INVESTIGATION

On April 29, 2026, the Chairmen sent letters to Airbnb, Inc. and Anysphere, Inc., the developer of the Cursor AI coding platform, requesting documents, information, and in-person briefings concerning their use of "**PRC-developed AI models**," meaning models whose original weights were trained by a company organized, headquartered, or principally operating in the PRC.[27] On July 31, 2026, the Chairmen expanded the investigation to DoorDash, Inc.[28] The requests focused principally on "**open-weight models**,"[29] whose trained parameters can be downloaded, run, and modified independently of the original developer, and sought information on the models companies used, how they evaluated and deployed them, and their interactions with PRC-based developers.[30]

Airbnb, Anysphere, and DoorDash responded through written submissions, document productions, and briefings to Committee staff. Airbnb produced materials on its approved generative AI models and how it selects, fine-tunes, and deploys them in customer support, and it later briefed Committee staff.[31] DoorDash provided a model inventory, guardrail descriptions, cost data, and internal AI policies, and later briefed Committee staff.[32] Anysphere produced a technical report describing the development of its proprietary coding model and also briefed staff.[33] Across the three companies, the Committees examined both how the models were selected and used and whether they were run on company-controlled infrastructure or accessed through services operated by the original developer.[34]

The companies' representations reflected different forms of adoption, from direct use in customer support systems to the use of PRC-developed models as components of proprietary products. **Table 1** summarizes the PRC-developed models identified in those productions and responses.[35]

[27] Airbnb Letter, *supra* note 3; Anysphere Letter, *supra* note 3.
[28] DoorDash Letter, *supra* note 3; DoorDash Press Release, *supra* note 2.
[29] PRC policy documents generally use the term "open-source" [开源], although many of the models discussed in this report are more precisely characterized as open-weight.
[30] Airbnb Letter, *supra* note 3; Anysphere Letter, *supra* note 3; DoorDash Letter, *supra* note 3.
[31] Letter from Airbnb, Inc., to John Moolenaar, Chairman, H. Select Comm. on the Strategic Competition Between the U.S. & the Chinese Communist Party, & Andrew R. Garbarino, Chairman, H. Comm. on Homeland Sec. (May 26, 2026) [hereinafter Airbnb Response]; Email from Airbnb, Inc., to Majority Staff, H. Comm. on Homeland Sec. (June 30, 2026) [hereinafter Airbnb June 30 Email]; Airbnb, Inc., *Generative AI Models Currently Approved for Production Use*, ABNB_HSCCCP_2026_0000118–0000123 (produced to the Committees) [hereinafter Airbnb Approved Models List]; Airbnb Customer Support Strategy, *supra* note 4; Briefing by Airbnb, Inc. to Majority Staff of the Committees (Aug. 27, 2026) [hereinafter Airbnb Briefing].
[32] DoorDash Response, *supra* note 6; DoorDash Production, *supra* note 6, attachs. A–F; DoorDash Briefing, *supra* note 6.
[33] Composer 2 Technical Report, *supra* note 5; Briefing by Anysphere, Inc. to Majority Staff of the Committees (2026) [hereinafter Anysphere Briefing].
[34] *See* Airbnb Approved Models List, *supra* note 31; DoorDash Production, *supra* note 6, attach. B; Anysphere Briefing, *supra* note 33.
[35] Airbnb Approved Models List, *supra* note 31; Airbnb Customer Support Strategy, *supra* note 4; Composer 2 Technical Report, *supra* note 5; DoorDash Response, *supra* note 6; DoorDash Production, *supra* note 6, attachs. A, C.

Table 1. PRC-Developed Models Identified in Company Records

Company	Developer	Models	Reported Use
Airbnb	Alibaba	Qwen 2.5 Qwen 3 Qwen 3.5	• Customer AI assistant (~80,000 interactions/day) • Agent case and call summary (~250,000/day, six languages)
	Zhipu	GLM-4.5 GLM-4.5-Air	• Approved for production use • GLM-4.5-Air shortlisted for AI assistant
Anysphere	Moonshot	Kimi K2.5	• Base for proprietary coding model (Mar. 2026)
DoorDash	Alibaba	Qwen family†	• Inventory*
	DeepSeek	DeepSeek V4 Pro	• Inventory*
	Moonshot	Kimi K2.6 Kimi K3	• K2.6 as first-pass code reviewer
	Zhipu	GLM-5 GLM-5.1 GLM-5.2	• Inventory*
	MiniMax	MiniMax M2.5	• Inventory*
	ByteDance	Anything V2‡	• AI research • Data labeling

† Includes Qwen2.5 0.5B; Qwen2.5-VL 3B and 32B; Qwen3 4B, 8B, 30B A3B, and 32B; Qwen3 Embedding and Reranker models; Qwen3-VL 4B, 8B, and 32B; Qwen3.5 397B A17B; and Qwen3.6 35B A3B.
‡ DoorDash deployed ByteDance's Depth Anything V2 Large, developed with the University of Hong Kong, for AI research and data labeling.
*Evaluated, tested, or deployed as reported in company inventory.

The Committees supplemented the company materials with public hearings, classified briefings, closed-door roundtables, and other research on the broader development and security of PRC AI models. The Subcommittee on Cybersecurity and Infrastructure Protection held hearings in December 2025, March 2026, and June 2026 on AI-enabled cyber threats, PRC-developed AI and robotics, and the security implications of frontier models, agentic AI, and AI coding tools.[36] Subcommittee staff then traveled to Japan, South Korea, and Taiwan in August 2026 to engage with representatives of host governments and private sector

[36] *The Quantum, AI, and Cloud Landscape: Examining Opportunities, Vulnerabilities, and the Future of Cybersecurity: Joint Hearing Before the Subcomm. on Cybersecurity & Infrastructure Prot. and the Subcomm. on Oversight, Investigations, & Accountability of the H. Comm. on Homeland Sec.*, 119th Cong. (Dec. 17, 2025) [hereinafter December 2025 Hearing]; Press Release, Subcomm. on Cybersecurity & Infrastructure Prot., H. Comm. on Homeland Sec., *Subcommittee Hearing on National Security Risks Posed by PRC Artificial Intelligence, Robotics, and Autonomous Technologies* (Mar. 15, 2026); *The AI Security Landscape: How Frontier Models, Agentic AI, and AI Coding Tools Are Reshaping Cybersecurity and Critical Infrastructure Resilience: Hearing Before the Subcomm. on Cybersecurity & Infrastructure Prot. of the H. Comm. on Homeland Sec.*, 119th Cong. (June 4, 2026).

leaders on AI adoption and governance, cybersecurity constraints, supply chain risks, and critical infrastructure resilience. Select Committee staff also traveled to Silicon Valley, Singapore, Malaysia, and Vietnam to examine the broader ecosystem supporting PRC AI development and overseas deployment, including access to advanced compute, cloud and data center infrastructure, and the channels through which those capabilities reach foreign markets. Staff from both Committees received classified briefings from leading American frontier AI developers and met with federal agencies, cybersecurity companies, AI laboratories, and researchers.[37]

The Committees also examined how PRC developers acquired the capabilities that made their models competitive with American alternatives. That inquiry included "**distillation**," a common training technique in which outputs from one model are used to train another, and alleged large-scale efforts by PRC developers to extract capabilities from American frontier systems.[38] In this report, "illicit distillation" refers to unauthorized, large-scale extraction through methods such as fraudulent accounts, proxy networks, or violations of a developer's terms of service.[39]

[37] *See* Sam Sabin & Ashley Gold, *Exclusive: OpenAI, Anthropic Meet with House Committee over Advanced Cyber Models*, Axios (Apr. 28, 2026), https://www.axios.com/2026/04/28/openai-anthropic-congress-cyber-briefings; David DiMolfetta, *House Homeland Panel Gets Briefing on Anthropic's Mythos*, Nextgov/FCW (May 14, 2026), https://www.nextgov.com/artificial-intelligence/2026/05/house-homeland-panel-gets-briefing-anthropics-mythos/413542/; Information provided to the Committees (on file with the Committees) [hereinafter Committee Records].

[38] Geoffrey Hinton, Oriol Vinyals & Jeffrey Dean, *Distilling the Knowledge in a Neural Network*, arXiv:1503.02531 (Mar. 9, 2015), https://perma.cc/PL89-CU8L; Anthropic Distillation Report, *supra* note 16; Memorandum from Michael J. Kratsios, Assistant to the President for Sci. & Tech. & Dir., Off. of Sci. & Tech. Pol'y, to the Heads of Exec. Dep'ts & Agencies, NSTM-4, *Adversarial Distillation of American AI Models* (Apr. 23, 2026), https://perma.cc/P6U3-SQMD [hereinafter NSTM-4].

[39] Anthropic Distillation Report, *supra* note 16; Joint Distillation Advisory, *supra* note 14; NSTM-4, *supra* note 38.

FINDINGS

Finding 1. PRC-developed open-weight models are already widely used across the American economy, but much of that use is hard to track.

Multiple independent data sources show that American use of PRC-developed open-weight models has grown rapidly. On OpenRouter, the share of tokens from U.S.-originating requests directed to open-weight models rose from 26 percent in September 2025 to 60 percent in August 2026, with models from PRC laboratories accounting for most of that volume.[40] A separate analysis found that PRC models accounted for more than 30 percent of weekly token consumption by American companies on the platform beginning in February 2026, peaking at 46 percent, up from 4.5 percent during the first half of 2025.[41] PRC-developed models also accounted for 41 percent of downloads on Hugging Face over the preceding year.[42] Alibaba's Qwen family alone had more than 151,000 derivative models on the platform, and its share of new open-model fine-tunes and adaptations rose from 1 percent in January 2024 to 69 percent by February 2026.[43] One leading venture capital firm observed that a majority of American AI startups appear to use PRC open weights somewhere in their technology stack.[44]

Spending data, however, can drastically understate how widely these models are used. In May 2026, DeepSeek models accounted for 17 percent of tokens processed through one major developer platform's AI gateway but only about 1 percent of spending.[45] A December 2025 survey of enterprise AI spending similarly found that PRC open models represented only about 1 percent of enterprise large language model application programming interface (API) usage.[46] The reason is simple: once weights are downloaded and run on private infrastructure, continued use may generate little or no commercial activity that can be measured.[47]

The companies examined by the Committees illustrate the point. Neither Airbnb nor DoorDash had paid a PRC developer or maintained a commercial agreement with one.[48] Each reported running PRC-developed models on infrastructure in the United States, substantially reducing the risk that user data is transmitted directly to the PRC, and DoorDash reported that it also blocks connections to PRC-based AI services from its corporate devices and model hosting environments.[49] But the absence of a commercial relationship does not mean the technology is no longer in use. Once the weights are downloaded, the

[40] Moberg, *supra* note 7.
[41] Metir AI, *Chinese AI Models Now Power Nearly Half of US Enterprise Tokens* (July 14, 2026), https://perma.cc/4AD8-PNZZ.
[42] Hugging Face Spring 2026 Report, *supra* note 7.
[43] Adina Yakefu et al., *State of Open Models: Summer 2026 Observations*, Hugging Face (Aug. 14, 2026), https://perma.cc/A93H-DKWU; Meyer & Buhler, *supra* note 7.
[44] Meyer & Buhler, *supra* note 7.
[45] *DeepSeek Now Handles 17% of Vercel Gateway Tokens*, *supra* note 10.
[46] Tully et al., *supra* note 10.
[47] *See, e.g.*, DoorDash Production, *supra* note 6, attach. C (reporting that costs associated with PRC-developed models are paid to U.S. computing providers that run the models on U.S. infrastructure).
[48] DoorDash Response, *supra* note 6; DoorDash Production, *supra* note 6, attach. C; Airbnb June 30 Email, *supra* note 31.
[49] Airbnb Response, *supra* note 31, at 3–4; DoorDash Response, *supra* note 6; DoorDash Production, *supra* note 6, attach. B; DoorDash Nov. 14 Email, *supra* note 6.

developer may disappear from the transaction while its model remains embedded in an American product. There may be no contract, invoice, ongoing vendor relationship, or other conventional record showing that dependence.

Model lineage, or the chain of models and training steps from which a model descends, can become even harder to see as weights are modified and reused. Fine-tuning, continued pretraining, merging, and distillation can progressively obscure a model's origin, allowing PRC-developed weights or capabilities to be modified and repackaged until it is no longer obvious that the resulting product was PRC-developed.[50] Anysphere, for example, built Cursor's flagship Composer 2 model on a PRC-developed base model, a fact that became public only after an outside developer identified the underlying model.[51] Researchers have shown that upstream model relationships can remain technically detectable after post-training and release under a different name, but identifying them requires the right tools and disclosure practices.[52] Today, however, no standard requires that lineage to be disclosed. An American-branded application, cloud service, or industrial system may therefore depend on a PRC-developed model several layers beneath what a customer sees.

That visibility gap becomes especially important when something goes wrong. If a serious vulnerability, hidden behavior, or disruption affected a widely used PRC model family, weak lineage records could make it difficult to determine which downstream products inherited the problem. The same problem would arise if future releases, security updates, or other parts of the model ecosystem became unavailable. Without reliable provenance records, neither customers nor the federal government can readily identify where PRC-developed models and their derivatives are operating or how far a problem could spread.

Finding 2. Cost, customization, and near-frontier performance are driving adoption, and the United States lacks a competitive open-weight alternative at the price and performance the market demands.

The companies examined by the Committees generally selected models based on performance, cost, latency, and flexibility. Airbnb Chief Executive Officer Brian Chesky publicly explained that the company relied on Alibaba's model because it was "fast and cheap."[53] Representations made to the Committees similarly showed that Airbnb, Anysphere, and DoorDash viewed PRC-developed open-weight models as competitive options in their own right, not substitutes of last resort.[54]

Price is a major part of that appeal. AI developers generally charge based on how much text a model reads and generates, measured in "tokens," or small pieces of words and characters. For the same standardized workload, one million input

[50] *See* Nikolic et al., *supra* note 10; Chang, *supra* note 10.
[51] Composer 2 Technical Report, *supra* note 5; Sumit Pandey, *A $29B Startup Got Caught. A Developer, an API Call, and 24 Hours.*, Medium (Mar. 26, 2026).
[52] Chang, *supra* note 10; Nikolic et al., *supra* note 10.
[53] Natalie Lung, *Chesky Says OpenAI Tools Not Ready for ChatGPT Tie-Up With Airbnb App*, Bloomberg (Oct. 21, 2025).
[54] *See* Airbnb Customer Support Strategy, *supra* note 4; DoorDash Response, *supra* note 6; Anysphere Briefing, *supra* note 33; *see also* Dora, *supra* note 5.

tokens and one million output tokens, Alibaba's Qwen3-235B-A22B-Instruct costs about $1.15 at published API rates, DeepSeek-R1 about $2.74, and Moonshot AI's Kimi K2.6 about $4.95.[55] The same workload costs approximately $30 using Anthropic's Claude Opus 4.8 and $35 using OpenAI's GPT-5.5.[56] In other words, depending on the model, an American company can pay roughly six to thirty times more for the same volume of model usage. These figures do not adjust for differences in capability or for caching, batch discounts, promotions, or self-hosting costs, but they illustrate the scale of the price gap that companies face when deciding which models to deploy.

Table 2. Published API Prices for Selected Models

Model (Developer)	Origin	Approx. cost per 1M input + 1M output tokens
Qwen3-235B-A22B-Instruct (Alibaba)	PRC	$1.15
DeepSeek-R1 (DeepSeek)	PRC	$2.74
Kimi K2.6 (Moonshot)	PRC	$4.95
Claude Opus 4.8 (Anthropic)	United States	$30.00
GPT-5.5 (OpenAI)	United States	$35.00

Notes: Figures are not adjusted for capability and do not reflect caching, batch discounts, promotions, or self-hosting costs.[57]

Open weights also give companies greater control over how a model is adapted and deployed. Companies can fine-tune models on their own data, adjust their behavior, run them on infrastructure they control, and reduce dependence on a single vendor that could change terms or cut off access.[58] Airbnb rated open-weight candidates as offering "full flexibility" for fine-tuning, compared with "limited flexibility" for proprietary alternatives.[59] DoorDash similarly told the Committees that open-weight models "reduce the risk a single developer could change or turn off access to a model without notice."[60]

That flexibility can matter in practice. In July 2026, after one American model under evaluation exploited a weakness in its testing environment and breached the network of Hugging Face, the company reported that American frontier models refused to assist with parts of the defensive response because their

[55] Alibaba Cloud Pricing, *supra* note 11; DeepSeek-R1 Release, *supra* note 11; Kimi Pricing Guide, *supra* note 11.
[56] Claude Opus 4.8 Announcement, *supra* note 11; GPT-5.5 Model Page, *supra* note 11.
[57] Alibaba Cloud Pricing, *supra* note 11; DeepSeek-R1 Release, *supra* note 11; Kimi Pricing Guide, *supra* note 11; Claude Opus 4.8 Announcement, *supra* note 11; GPT-5.5 Model Page, *supra* note 11.
[58] *See* Taylar Rajic & Lauryn Williams, *China's Open-Weight Challenge to U.S. AI Leadership*, Ctr. for Strategic & Int'l Stud. (Sept. 14, 2026), https://perma.cc/XT8S-3BBT; Nathan Lambert, *The Inevitable Need for an Open Model Consortium*, Interconnects (Apr. 11, 2026).
[59] Airbnb, Inc., *AI Assistant Overview*, ABNB_HSCCCP_2026_0000135–0000140 (produced to the Committees) [hereinafter Airbnb AI Assistant Overview].
[60] DoorDash Response, *supra* note 6.

safeguards could not distinguish a defender from an attacker.[61] Hugging Face instead used Zhipu's open-weight GLM-5.2 model to complete portions of the work.[62] The episode illustrates why companies and cybersecurity researchers may value models they can run and modify themselves, even when those models raise separate national security concerns.

PRC-developed models are also increasingly capable. DoorDash's internal benchmark, DashBench, measured how well paired AI models identified real problems in 105 historical code changes. A configuration pairing Moonshot's Kimi K2.6 with an American frontier model achieved the highest score of any configuration tested, outperforming pairings built on leading American open-weight models by more than 20 percentage points.[63] Table 3 summarizes those results.

Table 3. DoorDash DashBench Code Review Benchmark Results

Initial Reviewer (Origin, Type)	Second Reviewer	Weighted F1 Score
Kimi K2.6 (PRC, open-weight)	Claude Fable 5 (U.S., proprietary)	75.3%
Claude Sonnet 4.6 (U.S., proprietary)	Claude Opus 4.8 (U.S., proprietary)	66.3%
Kimi K2.6 (PRC, open-weight)	Claude Opus 4.8 (U.S., proprietary)	58.9%
NVIDIA Nemotron 3 Ultra (U.S., open-weight)	Claude Fable 5 (U.S., proprietary)	53.5%
Inkling (U.S., open-weight)	Claude Fable 5 (U.S., proprietary)	52.1%

Notes: Scores reflect a weighted measure of precision and recall on 105 historical code changes. The company stated that these evaluations do not determine which models it adopts.[64]

Market-wide benchmarks highlight this rapid convergence. The highest-ranked American open-weight model now sits at tenth place in the latest Artificial Analysis Intelligence Index, an independent ranking that measures overall capability across reasoning, knowledge, mathematics, coding, and agentic tasks, trailing six PRC models.[65] While PRC systems have lagged American frontier models by roughly seven months on average since 2023, rapid releases from

[61] OpenAI, *OpenAI and Hugging Face Partner to Address Security Incident During Model Evaluation* (July 21, 2026, updated Aug. 26, 2026), https://perma.cc/F7MS-9PJW [hereinafter OpenAI Incident Disclosure].
[62] *Id.*; Rebecca Bellan, *As US Weighs Response to Chinese AI, Industry Urges Against Broad Open-Weight Restrictions*, TechCrunch (July 24, 2026); Rajic & Williams, *supra* note 58.
[63] DoorDash Response, *supra* note 6.
[64] *Id.*; *see also* DoorDash Briefing, *supra* note 6.
[65] Artificial Analysis Index, *supra* note 12.

DeepSeek, Moonshot, Alibaba, and Zhipu have dramatically narrowed that lead across several core domains.[66]

That narrowing gap is changing how companies deploy AI. Rather than choosing between an American model and a PRC-developed model, companies increasingly use different models for different tasks. Airbnb described a portfolio approach in which "no single LLM can efficiently serve all" of its support use cases,[67] while DoorDash described routing some work to a PRC-developed model and more difficult tasks to an American model.[68] Companies also described controls that treat model outputs as untrusted and subject them to additional review.[69] Those controls can reduce the risk of a bad output causing harm, but they do not answer other questions, including where a model came from, how it may behave under different conditions, or what happens if access to it is disrupted. Using several models at once can also make it harder to know which model handled a particular task or had access to particular data.

For American companies, the larger problem is the limited supply of competitive domestic open-weight models near the frontier. DoorDash cited the lack of competitive American open-weight alternatives as a principal reason for using PRC-developed models and described American open-weight models as "a strategic asset for U.S. competitiveness and security."[70] In the September 2026 Artificial Analysis ranking, the highest-ranked American open-weight model, Thinking Machines Lab's Inkling, placed behind six PRC models.[71] Thinking Machines Lab also disclosed that Inkling's initial training used synthetic data generated by other open-weight models, including Moonshot's Kimi K2.5.[72] That means even some American efforts to build alternatives are still drawing on PRC-developed models.

One reason American alternatives remain limited is that U.S. developers are playing by a different set of rules. Training frontier models is expensive, and American companies generally rely on proprietary access to recover those costs rather than releasing their strongest models as open weights.[73] They also cannot lawfully replicate the distillation campaigns described in Finding 3, which used fraudulent accounts and proxy networks to extract capabilities from competing American models in violation of provider restrictions.[74] PRC developers, meanwhile, benefit from state investment and subsidized compute.[75] The result is an uneven playing field in which American companies bear the cost of developing

---

[66] Luke Emberson, *Chinese AI Models Have Lagged the US Frontier by 7 Months on Average Since 2023*, Epoch AI (Jan. 2, 2026), https://perma.cc/7TQ7-AFAF; Nikolaus Lang et al., *The Great Divide: How the US and China Are Splitting the AI World*, Bos. Consulting Grp. (June 30, 2026), https://perma.cc/7FZ4-BRST; Rajic & Williams, *supra* note 58; Stanford Inst. for Hum.-Centered AI, *2026 AI Index Report*, ch. 2 (2026), https://perma.cc/5YJA-689A [hereinafter Stanford AI Index].
[67] Airbnb Customer Support Strategy, *supra* note 4.
[68] DoorDash Response, *supra* note 6.
[69] *Id.*; DoorDash Production, *supra* note 6, attach. B; Airbnb Response, *supra* note 31, at 3–4.
[70] DoorDash Response, *supra* note 6.
[71] Artificial Analysis Index, *supra* note 12.
[72] Thinking Machines Lab, *Inkling: Our Open-Weights Model* (July 15, 2026), https://perma.cc/Z254-468S [hereinafter Inkling Announcement]; Bellan, *Industry Urges Against Broad Open-Weight Restrictions*, *supra* note 62.
[73] Lambert, *supra* note 58.
[74] Joint Distillation Advisory, *supra* note 14, at 1–2; NSTM-4, *supra* note 38; Anthropic Distillation Report, *supra* note 16.
[75] *See* Chan et al., *supra* note 18; *Tech War: China Creates US$8.2 Billion AI Investment Fund Amid Tightened US Trade Controls*, S. China Morning Post (Jan. 20, 2025); Lo, *supra* note 18; *see also infra* Annex B.

frontier capabilities, while PRC firms can use those capabilities, backed by state support, to bring cheaper open-weight models to market.

There are signs that American open-weight developers, including Meta, NVIDIA, Thinking Machines Lab, Arcee AI, Poolside, and Reflection AI, are starting to catch up.[76] For example, on October 5, 2026, Reflection AI released Beam, which it says matches leading PRC open models on advanced reasoning benchmarks while using substantially less inference compute, although those claims have not yet been independently verified.[77] Allied developers are advancing as well. In October 2026, France-based Mistral AI launched Mistral Large 4, which the company says ranks among the strongest open-weight models available and outperforms PRC models in certain areas, including cyber.[78] On the policy side, the Administration's AI Action Plan calls for a "supportive environment for open models" aimed at making American models global standards.[79]

Finding 3. Much of the PRC's apparent progress rests on capabilities extracted from American models, and American firms are helping supply the data that trains PRC models.

The PRC's gains have occurred despite U.S. controls intended to restrict access to the most advanced semiconductors.[80] Some analysts argue that those constraints pushed PRC developers toward more efficient architectures and open-weight distribution.[81] However, export controls were only part of the story. U.S. government agencies and American frontier developers found that several leading PRC AI labs also accelerated their models by extracting capabilities from American systems, while at least one developer obtained advanced American chips it was not authorized to acquire.[82]

By early 2026, those extraction campaigns were operating at industrial scale. In February 2026, American frontier developers disclosed that DeepSeek, Moonshot AI, and MiniMax had collectively generated more than 16 million exchanges with American AI systems through approximately 24,000 fraudulent accounts.[83] In April 2026, the White House Office of Science and Technology Policy (OSTP) described the broader activity as "deliberate, industrial-scale campaigns," principally originating from the PRC, to distill American frontier AI systems.[84] OSTP also warned that developers could use the resulting models to reproduce

[76] Inkling Announcement, *supra* note 72; Rebecca Bellan, *Reflection Debuts Beam, an Open-Weight AI Model to Rival Chinese Models at Lower Compute Cost*, TechCrunch (Oct. 5, 2026); Artificial Analysis Index, *supra* note 12.
[77] Bellan, *Reflection Debuts Beam*, *supra* note 76.
[78] Elizabeth Howcroft & Leo Marchandon, *France's Mistral Launches AI Model It Says Outperforms Some Chinese Rivals*, Reuters (Oct. 6, 2026).
[79] The White House, *Winning the Race: America's AI Action Plan* 4 (July 23, 2025), https://perma.cc/826C-EEQR [hereinafter AI Action Plan].
[80] *See* Rajic & Williams, *supra* note 58; Caiwei Chen, *China's Open-Source Bet*, MIT Tech. Rev. (Apr. 21, 2026).
[81] Rajic & Williams, *supra* note 58.
[82] Joint Distillation Advisory, *supra* note 14; Michael Kratsios (@mkratsios47), X (July 22, 2026); *see* Kai Nicol-Schwarz, *Moonshot AI Accessed Nvidia's Chips Despite Chinese Export Ban, White House Official Says*, CNBC (July 23, 2026).
[83] Anthropic Distillation Report, *supra* note 16; OpenAI Memorandum, *supra* note 16.
[84] NSTM-4, *supra* note 38.

advanced capabilities while "deliberately strip[ping] security protocols" from them.[85]

In September 2026, a joint advisory from the NSA, CISA, and FBI laid out how PRC developers sustained these operations, including by routing queries through proxy transfer stations, leveraging bulk subscriptions, obfuscating request origins, and deploying automated filters to bypass distillation checks.[86] Public disclosures by American frontier developers confirm these patterns, as summarized in **Table 4**.[87]

Table 4. Illicit Distillation Activity Attributed to PRC Developers

Developer	Date(s) of Reported Activity	Targeted U.S. Models	Reported Activity	Capabilities Targeted
DeepSeek	Late 2024; 14-day campaign in July 2026	Claude, Gemini, GPT, Grok	12 million+ exchanges	Advanced Reasoning; Agentic Functions
Moonshot	Mid-2025; May–July 2026	Claude Fable 5	23 million+ exchanges	Advanced Reasoning
Alibaba	Late 2025; May–July 2026	Claude, GPT-5	151 million+ exchanges; peak of nearly 3 million per day	Software Engineering; Agentic Functions
Zhipu	By mid-2026; 17-day campaign in June–July 2026	GPT-5.5, Claude Opus 4.8	3.4 million+ exchanges; Billions of tokens distilled	Advanced Reasoning; Cyber
MiniMax	Late 2025	American frontier and coding models	Large-scale distillation and use of proxy infrastructure	Advanced Reasoning; Software Engineering
StepFun	Late 2025–early 2026	Claude and GPT-family models	Distillation of Claude and GPT-family models	Coding; Agentic Functions
Xiaomi	March–April 2026	American frontier models	400,000+ requests across 1,500+ accounts	Not publicly specified

Sources: Joint Distillation Advisory, *supra* note 14; Anthropic September 2026 Report, *supra* note 87.

The joint advisory identified DeepSeek, Moonshot AI, Alibaba, MiniMax, StepFun, and Z.ai, the international brand of Zhipu, as participants in large-scale extraction campaigns involving billions of tokens from American frontier systems, likely with PRC government awareness, and described distillation as "the core" of the development strategy of several leading PRC AI labs.[88] The advisory also

[85] *Id.*
[86] Joint Distillation Advisory, *supra* note 14.
[87] *See* Anthropic Distillation Report, *supra* note 16; OpenAI Memorandum, *supra* note 16; Anthropic, *Detecting and Countering Misuse of AI: September 2026* (Sept. 10, 2026) [hereinafter Anthropic September 2026 Report].
[88] Joint Distillation Advisory, *supra* note 14, at 1–2.

concluded that DeepSeek's widely cited $5.6 million training-cost figure was "misleading" because it excluded the value of data acquired through distillation.[89]

An American frontier developer found that a distilled model can acquire dangerous capabilities, including in biological and cyber domains, even when the harvested exchanges contain little information about those subjects, while safeguards limiting misuse of the original model "do not transfer" to the distilled model.[90] The same developer reported that Zhipu employees stopped trying to distill the cyber capabilities of one American model after its safeguards interfered with the requested attacks, then shifted to other American models they believed had weaker protections.[91] Those extracted capabilities are not confined to commercial use, as discussed further in Finding 4.

PRC developers are also obtaining American training data through ordinary commercial relationships.[92] Reporting in August 2026 found that leading PRC AI laboratories collectively pay U.S. data vendors, including Surge AI, Mercor, AfterQuery, and Turing, roughly $500 million a year for expert reasoning, coding, evaluation, and other training data.[93] Some of those vendors also serve American frontier laboratories and the U.S. Army and Air Force, while reported PRC customers include Tencent, ByteDance, Alibaba, Ant Group, and Moonshot.[94] At least one U.S. data company reportedly declined a ByteDance contract over national security concerns, while another does not sell outside the United States.[95]

The Committees do not suggest that these sales are currently unlawful. They do, however, raise a separate policy question of whether American expert training data, including data from firms serving U.S. defense customers, should remain available without restriction to PRC laboratories designated by the U.S. government as military companies or placed on the Entity List.[96]

PRC developers are drawing on American resources through two separate channels. Illicit distillation transfers capabilities directly from American frontier models, while commercial purchases provide the high-quality data used to train and refine PRC systems. Both shorten the path to models that can compete with American systems.

Finding 4. Open-weight distribution is a deliberate PRC strategy to expand global dependence, backed by state policy and capital, and the monetization phase has begun.

[89] *Id.* at 1–2.
[90] Anthropic September 2026 Report, *supra* note 87.
[91] *Id.*
[92] Tong, *supra* note 15.
[93] *Id.*
[94] *Id.*
[95] *Id.*
[96] *See* U.S. Dep't of War, *Entities Identified as Chinese Military Companies Operating in the United States in Accordance with Section 1260H of the William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021* (June 8, 2026) [hereinafter Section 1260H List]; Addition of Entities to and Revision of Entry on the Entity List, 90 Fed. Reg. 4617, 4618–20 (Jan. 16, 2025).

The PRC has made open-weight distribution central to its effort to expand the global reach and influence of its AI ecosystem. The rapid international adoption described above suggests that the strategy is gaining traction outside the PRC.[97] PRC policy documents call for an open-source AI ecosystem with "international influence," while researchers at a state academy have described the approach as pairing "low-cost technology exports" with "high-stickiness ecosystem binding."[98] A 2025 article published by the Cyberspace Administration of China urged PRC developers to "attract global developers to participate in technological iteration" and increase Beijing's "international discourse power over technical rules."[99] The National Office for Philosophy and Social Sciences went further, arguing that open-source community rules could shape the "underlying structure" of global AI governance and urging the PRC to "substantively intervene in and shape" those rules.[100] In July 2026, Xi Jinping also endorsed open-source AI, calling for "open-source, openness, collaboration and sharing" in AI.[101]

That direction from Beijing is already reflected in the strategies and public commitments of leading PRC AI developers. Moonshot AI describes open-weight models as the foundation of a "broad and open ecosystem,"[102] while DeepSeek founder Liang Wenfeng has said that the company's "strongest model will probably also be open sourced."[103] Zhipu founder Tang Jie released GLM-5.2 under the permissive MIT license, arguing that "true safety is not built on technological closure and barriers."[104] Alibaba's chief executive likewise pledged to keep Qwen open-source and to turn it into the "operating system of the AI era."[105]

The PRC has committed substantial state financing to support this strategy. In January 2025, the PRC established an $8.2 billion[106] National AI Industry Investment Fund backed by state investors,[107] while the state-owned Bank of China announced plans to provide approximately $138 billion in AI-related financing over five years.[108] Leading model developers have also benefited from direct subsidies, state investment, and government ownership:

---

[97] *See* Microsoft AI Econ. Inst., *Global AI Adoption in 2025: A Widening Digital Divide* (Jan. 2026), https://perma.cc/NRD5-G852 [hereinafter Microsoft AI Diffusion Report]; Caiwei Chen, *supra* note 80.
[98] State Council AI+ Opinions, *supra* note 17, § III(11); Chen & Zhu, *supra* note 17; Xi Jinping, *Joining Hands to Build a Just and Equitable System for Global AI Governance* [携手构建公正合理的全球人工智能治理体系], Ministry of Foreign Affs. of China (July 17, 2026), https://perma.cc/WS7L-AX8B.
[99] Cyberspace Admin. of China, *Vigorously Promote the Development of China's Artificial Intelligence Large Models* [大力推动我国人工智能大模型发展] (May 8, 2025), https://perma.cc/V5PE-F2PW.
[100] Nat'l Off. for Phil. & Soc. Scis., *Constructing an Internally and Externally Integrated Artificial Intelligence Rule-of-Law Paradigm* [构建内外贯通的人工智能法治范式] (July 4, 2025), https://perma.cc/47ZP-TFHQ.
[101] Xi, *supra* note 98.
[102] Moonshot AI, *Kimi K3 Open Day: Model Weights, Technical Report, and Key Infrastructure Technologies Released Simultaneously* [Kimi K3 开放日：模型权重、技术报告和关键 Infra 技术同步开放] (July 27, 2026), https://perma.cc/FH3X-PRVC.
[103] Liang Wenfeng, Founder, DeepSeek, *Investor Meeting Remarks* [梁文锋投资人会议文字稿修正稿], Sumsec (May 20, 2026), https://perma.cc/4D6G-AUFF.
[104] *Zhipu's Tang Jie Says in Internal Letter That the Shift from AGI to ASI Is Irreversible, Launches the "Reach Higher" Initiative* [智谱唐杰内部信称从AGI到ASI不可逆转，开启"摸高计划"], Sina Finance (July 12, 2026), https://perma.cc/78C2-XUK5.
[105] Alibaba Cloud, *Alibaba Cloud's Apsara Conference 2025: Full Stack AI + Cloud Leads the Way to the Future of AI* (Sept. 29, 2025), https://perma.cc/8MFJ-Y55T.
[106] For this report, U.S. dollar equivalents are calculated using an exchange rate of approximately RMB 7.25 per U.S. dollar.
[107] *Tech War*, *supra* note 75.
[108] Chan et al., *supra* note 18.

- **Alibaba** received approximately $1.05 billion in company-wide government grants in fiscal year 2024,[109] while Hangzhou separately awarded Alibaba DAMO Academy about $4.7 million specifically for Qwen.[110]
- **Zhipu** disclosed at least $13.8 million in AI-related R&D subsidies, received a $29.8 million state investment, and is 4.9 percent owned by government bodies, with another 3.5 percent held by Tsinghua University.[111]
- **DeepSeek** received roughly $146.2 million from China's National AI Industry Investment Fund in 2026.[112]
- **MiniMax** reported about $2.3 million in deferred government grants at the end of 2025,[113] with an additional, undisclosed amount of state capital provided through multiple financing rounds led by Shanghai state-capital firms.[114]
- **Moonshot** is 18.1 percent indirectly owned by eighteen Chinese government bodies through investment funds and partnerships.[115]
- **ByteDance** has a 1 percent "golden share" in its key China operating entity held by Beijing through WangTou Zhongwen (Beijing) Technology Co., Ltd.[116]

Annex B documents additional provincial, municipal, and district subsidies for compute, deployment, taxes, talent, and customer adoption. Together, these programs reduce the costs of developing and commercializing PRC AI models. The scale of financial losses reported by leading developers illustrates the scale of investment required to compete. When Zhipu and MiniMax listed in Hong Kong in January 2026, Zhipu disclosed approximately $330 million in losses on $27 million in revenue during the first half of 2025, while MiniMax reported $512 million in losses on $53 million in revenue during the first nine months.[117] Independent analysts nevertheless identify private technology firms as the primary drivers of PRC AI progress, with state support most effective when directed toward those firms.[118] Government financing and subsidies give

[109] Alibaba Grp. Holding Ltd., *Fiscal Year 2024 Annual Report* (May 23, 2024), https://perma.cc/KP96-NXDM [hereinafter Alibaba FY2024 Annual Report]. Of note, the filing does not attribute those grants specifically to Qwen, compute, or other AI-related activities. By 2026, Alibaba no longer disclosed a standalone grant total, instead folding government support into "other income" and expense offsets while describing it broadly as support for "technology development" and "investments in local business districts." *See* Alibaba Grp. Holding Ltd., *Fiscal Year 2026 Annual Report* (May 20, 2026), https://perma.cc/59FN-PKL9.

[110] *2024 Hangzhou Artificial Intelligence Industry Funding Subsidy Project List Announced* [2024年杭州市人工智能产业资金补助项目名单公布], Hangzhou SME Serv. Network (Dec. 16, 2024), https://perma.cc/TH2T-YAYC; Alibaba FY2024 Annual Report, *supra* note 109.

[111] Knowledge Atlas Tech. JSC Ltd., *Prospectus* [招股書] (Dec. 30, 2025), https://perma.cc/HLY8-RQBM [hereinafter Zhipu Prospectus]; *Announcement on the Public Selection of a Management Institution for the Beijing Artificial Intelligence Industry Investment Fund* [北京市科学技术委员会、中关村科技园区管理委员会、北京市政府投资引导基金关于公开遴选北京市人工智能产业投资基金管理机构的公告], Nat'l Sci. & Tech. Innovation Ctr. (Sept. 18, 2023), https://perma.cc/H2NQ-2GDK; *Beijing Artificial Intelligence Industry Investment Fund Adds RMB 200 Million Investment in Zhipu (Z.ai)* [北京市人工智能产业投资基金追加投资智谱（Z.ai）2亿元], Beijing Mun. Sci. & Tech. Comm'n & Zhongguancun Sci. Park Admin. Comm. (Apr. 18, 2025), https://perma.cc/4EBJ-MRK8.

[112] *DeepSeek Completes First Financing Round*, *supra* note 18; *"It's Not That You Want to Invest and Liang Wenfeng Will Take Your Money": DeepSeek Raises RMB 51 Billion* ["不是你想投，梁文锋就会要你的钱"，DeepSeek融资510亿元，他为何选了腾讯、宁德时代、网易、京东？], Nat'l Bus. Daily (June 18, 2026), https://perma.cc/R4L4-FVB2.

[113] MiniMax Grp. Inc. [稀宇科技], *Annual Report 2025* (2026), https://perma.cc/72RT-WG54 [hereinafter MiniMax Annual Report].

[114] *MiniMax Lists on the Hong Kong Stock Exchange*, *supra* note 18.

[115] *China & AI*, WireScreen Briefings (Aug. 6, 2026), https://wirescreen.ai/briefings/china-ai.

[116] Huatai United Securities Co., Ltd., *Verification Opinion* [核查意见] (Jan. 17, 2023), https://perma.cc/6J5Q-TWD8.

[117] Lo, *supra* note 18.

[118] Chan et al., *supra* note 18.

developers greater latitude to absorb losses, lower prices, and prioritize market share over near-term profitability.

The relationship between leading developers and the PRC state also extends beyond financing. Several have roots in research networks connected to state security and intelligence institutions, with some ties continuing through cybersecurity partnerships. Zhipu, for example, traces part of its research origins to a Tsinghua University project with ties to the PRC intelligence apparatus.[119] In 2017, Zhipu co-founder Li Juanzi led a National Natural Science Foundation of China project on person profiling and data analysis under a joint fund with the China General Technology Research Institute, an entity that supports the Ministry of State Security's (MSS) technical espionage capabilities.[120] Tsinghua identified XLORE 2.0 and NewsMiner as outputs of the project, designed to analyze "China-related key persons and events."[121] The project drew on Tsinghua's Knowledge Engineering Group (KEG) and involved researchers who later joined Zhipu,[122] including co-founder Tang Jie,[123] CEO Zhang Peng,[124] and researchers Xin Lv and Zijun Yao.[125]

Moonshot AI founder Yang Zhilin also emerged from Tsinghua's KEG, where he conducted research under Tang Jie with support from government projects focused on "public security and social management."[126] He later contributed to Zhipu's GLM model lineage through the Beijing Academy of Artificial Intelligence and, through his company Recurrent AI, collaborated with Huawei Cloud to develop the PanGu NLP model using computing resources from Peng Cheng

[119] *See* Wu Guozheng et al., *Review of the 2017 Funding Evaluation Work of the Department of Information Sciences* [2017年度信息科学部基金评审工作综述], Chinese Sci. Found. (Dec. 14, 2017), https://perma.cc/4CAW-Y5E7; MI5, *MI5 Issues Espionage Alert – CGTRI 中国通用技术研究院* (Sept. 30, 2026), https://perma.cc/55FG-3BVT [hereinafter MI5 Espionage Alert].

[120] Qinghua Wen et al., *A Multi-Strategy Approach to Chinese Open Relation Extraction* [多策略中文开放关系抽取方法], Journal of Chinese Information Processing (Jan. 2023), https://perma.cc/K4TF-2VZR; Li Juanzi, *Faculty Profile* [教师简介], Tsinghua Univ. Dep't of Comput. Sci. & Tech., https://perma.cc/F9S6-ZT7A [hereinafter Li Juanzi Faculty Profile]; Wu et al., *supra* note 119; Cui Lin et al., *Review of the 2017 Funding Evaluation Work of the Department of Chemical Sciences* [2017年度化学科学部基金项目评审工作综述], Chinese Sci. Found. (Jan. 11, 2018), https://perma.cc/5HEH-Q578; MI5 Espionage Alert, *supra* note 119.

[121] Li Juanzi Faculty Profile, *supra* note 120.

[122] *Knowledge Engineering Laboratory of the Department of Computer Science Jointly Releases New Academic Search Product with Sogou and Zhipu AI* [计算机系知识工程实验室联合搜狗与智谱AI发布全新学术搜索产品], Tsinghua Univ. Dep't of Comput. Sci. & Tech. (June 3, 2021), https://perma.cc/J56T-ZQ7K.

[123] Jifan Yu et al., *Course Concept Expansion in MOOCs with External Knowledge and Interactive Game*, 57th Ann. Meeting Ass'n for Computational Linguistics (Aug. 2019), https://perma.cc/Z8LE-KLQX; Tony Peng, *Zhipu AI Chief Scientist Tang Jie: Making Machines Think Like Humans*, Recode China (June 30, 2026), https://perma.cc/K4XS-V4UW.

[124] Hailong Jin et al., *XLORE2: Large-Scale Cross-Lingual Knowledge Graph Construction and Application*, Data Intelligence (2019), https://perma.cc/C5QK-R3TT; *First Zhipu AI Technology Open Day Successfully Held* [首届智谱 AI 技术开放日成功举办], Zhipu AI (Jan. 16, 2024), https://perma.cc/39WL-Q222.

[125] Xin Lv et al., *Is Multi-Hop Reasoning Really Explainable? Towards Benchmarking Reasoning Interpretability*, in Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (Nov. 2021), https://perma.cc/W3BV-HKCX; Zijun Yao et al., *Interpretable and Low-Resource Entity Matching via Decoupling Feature Learning from Decision Making*, in Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics (Aug. 2021), https://perma.cc/BYV2-BFQD; Xin Lv, *About Me*, https://perma.cc/3JUU-JFJG; Zhipu & Tsinghua Univ., *GLM-5: From Vibe Coding to Agentic Engineering*, arXiv (Feb. 24, 2026), https://perma.cc/E3LF-SSD7.

[126] *Yang Zhilin Interview: "I Have Not Done Enough to Deserve This Much Attention"* [杨植麟接受采访："我还不足以得到这么多关注"], Shantou Special Econ. Zone Evening News [汕头特区晚报] (Nov. 20, 2014), https://perma.cc/AFC2-YETC; Liu Yukun, *Eleven-Year-Old Tsinghua Defense Goes Viral: Kimi's Yang Zhilin Received Perfect Scores, and His Recommender Later Founded Zhipu* [11年前清华一场答辩刷屏：Kimi杨植麟全满分，推荐人后来创办智谱], Phoenix Tech. (July 25, 2026), https://perma.cc/KE6C-DM9T. *See, e.g.*, Zhilin Yang, Jie Tang & Yutao Zhang, *Active Learning for Streaming Networked Data*, in Proceedings of the 23d ACM International Conference on Information and Knowledge Management (Nov. 2014), https://perma.cc/SX6G-AP2P; *973 Project Approved by the Ministry of Science and Technology* [973项目获科技部批准], Tsinghua Univ. Dep't of Comput. Sci. & Tech. (Jan. 13, 2014), https://perma.cc/2RLC-UJ8P. In China, "social management" is often a euphemism for surveillance and political control. *See, e.g.*, Peter Mattis & Samantha Hoffman, *Plenum Document Highlights Broad Role for Social Management*, China Brief (Oct. 28, 2011); Samantha Hoffman, *Managing the State: Social Credit, Surveillance and the CCP's Plan for China*, China Brief (Aug. 17, 2017).

Laboratory.[127] Additional information points to contacts between Zhipu and elements of the PRC's military cyber apparatus, as detailed further in **Annex A**.[128]

Zhipu and Alibaba Cloud also participate in the PRC's government cybersecurity and vulnerability-reporting systems. In April 2026, a Zhipu executive participated in the launch of the National Artificial Intelligence Security Vulnerability Database, established under the China National Vulnerability Database of Information Security, which is maintained by the MSS's 13th Bureau.[129] Zhipu subsequently reported vulnerabilities identified through its GLM models to the database.[130] Alibaba Cloud has also maintained longstanding ties to the PRC's national cybersecurity apparatus. In 2019, the National Computer Network Emergency Response Technical Team/Coordination Center of China designated Alibaba Cloud a national-level cybersecurity emergency service support unit alongside firms including Qihoo 360 and NSFOCUS.[131] In 2021, PRC regulators suspended Alibaba Cloud from a government threat-sharing platform for disclosing the Log4j vulnerability to the Apache Software Foundation before notifying PRC authorities.[132]

U.S. government designations and intelligence assessments have also raised concerns about the military and intelligence ties of leading PRC AI developers. The Department of War designated Alibaba a Chinese military company in June 2026, while Zhipu has been on the Commerce Department's Entity List since January 2025.[133] A senior U.S. official stated that DeepSeek "has willingly provided and will likely continue to provide support to China's military and intelligence operations," citing more than 150 procurement records involving the PLA and entities affiliated with the PRC defense industrial base, along with technology services provided to PLA research institutions.[134] The Committees' review of PLA procurement records likewise identified military applications involving several PRC AI model families, including:

---

[127] Chang Zheng, *New Breakthrough in Chinese AI Research: BAAI Releases Wu Dao 1.0* [中国AI研究新突破：智源"悟道1.0"发布], Beijing Academy of Artificial Intelligence [北京智源人工智能研究院] (Mar. 20, 2021), https://perma.cc/7JMH-MUTU; Zhang Qian, *Detailed Explanation of Wu Dao-Wenhui* [悟道·文汇详解], Beijing Acad. of AI (Mar. 24, 2021), https://perma.cc/253L-BTYX; Zhengxiao Du et al., *GLM: General Language Model Pretraining with Autoregressive Blank Infilling*, arXiv (Mar. 18, 2021), https://perma.cc/6MQQ-KH22; Zhang Qian, *Targeting GPT-3's Commercialization Problem, the First 100-Billion-Parameter Chinese Model "PanGu" Debuts for Enterprise Applications* [瞄准GPT-3落地难题，首个千亿中文大模型"盘古"问世，专攻企业级应用], Machine Intelligence (Apr. 25, 2021), https://perma.cc/X587-UWTN; *Recurrent AI and Huawei Cloud Jointly Launch the World's Largest Chinese-Language Model* [循环智能与华为云联合推出全球最大中文语言模型], Consumer Daily (Apr. 26, 2021), https://perma.cc/53LQ-FUQZ; *Why Has Beijing Become the Source of AI Innovation?* [北京缘何成为AI创新之源？], Beijing Sci. & Tech. Innovation Ctr. Off. (July 20, 2026), https://perma.cc/CR2A-MCBP.

[128] *See infra* Annex A; Committee Records, *supra* note 37.

[129] *National Artificial Intelligence Security Vulnerability Database Launched* [国家人工智能安全漏洞库启动运行], China Nat'l Vulnerability Database (Apr. 23, 2026), https://perma.cc/54JX-JUHB; *China Information Technology Security Evaluation Center*, China Info. Tech. Sec. Evaluation Ctr. (Aug. 2, 2017), https://perma.cc/NRS9-6LAD; U.S.-China Econ. & Sec. Rev. Comm'n, *China's Cyber Capabilities: Warfare, Espionage, and Implications for the United States* (2022).

[130] Z.AI, *Disclosure Ledger* [披露账本], https://perma.cc/TE2W-AR9L.

[131] Nat'l Computer Network Emergency Response Tech. Team/Coordination Ctr. of China, *List of the Eighth Cohort of CNCERT Cybersecurity Emergency Service Support Units* [第八届CNCERT网络安全应急服务支撑单位名单] (July 17, 2019), https://perma.cc/QJ98-ENEX.

[132] Liu Shidong, *Person Close to MIIT Confirms: Alibaba Cloud Suspended as Partner Unit of Cybersecurity Threat Information Sharing Platform* [接近工信部人士确认：阿里云被暂停网络安全威胁信息共享平台合作单位], China Daily (Dec. 22, 2021), https://perma.cc/59CY-TAW5; Ministry of Indus. & Info. Tech. Cybersecurity Admin. Bureau, *Cybersecurity Risk Alert on the Major Security Vulnerability in the Apache Log4j2 Component* [关于阿帕奇Log4j2组件重大安全漏洞的网络安全风险提示] (Dec. 17, 2021), https://perma.cc/5EUG-5WCX.

[133] Section 1260H List, *supra* note 96; Addition of Entities to and Revision of Entry on the Entity List, 90 Fed. Reg. 4617, 4618–20 (Jan. 16, 2025).

[134] Michael Martina & Stephen Nellis, *Exclusive: DeepSeek Aids China's Military and Evaded Export Controls, US Official Says*, Reuters (June 23, 2025).

- A 2026 procurement sought servers for an "intelligent command information system" and required support for the full development workflow of Alibaba's Qwen and DeepSeek models.[135]
- A 2026 cyber range procurement required integrating DeepSeek to support "intelligent attack" and "intelligent penetration."[136]
- A 2026 model-upgrade project required DeepSeek, Qwen, Moonshot's Kimi, and MiniMax models to operate entirely on an internal network without internet access for military AI agent and software development applications.[137]
- A 2026 military medical procurement called for locally deploying DeepSeek or Qwen to support medical record generation, quality control, and clinical decision support.[138]
- A 2024 procurement called for deploying and fine-tuning Zhipu's GLM models on a military office network.[139]

These records show PRC-developed models being incorporated into PLA computing environments and operational workflows, including through purchases of servers, software platforms, and local deployment infrastructure configured around specific model families. Additional procurements for integrated foundation model platforms and dedicated local deployment infrastructure show that this use extends beyond the examples above.[140]

The same state-backed ecosystem that supports these models at home is also helping them gain adoption abroad, creating a separate risk that low-cost access today becomes dependence tomorrow.[141] There are already signs that PRC developers are seeking to convert adoption into commercial leverage. Moonshot AI's Kimi K3 license requires model-as-a-service providers with more than $20 million in annual revenue to negotiate a separate commercial agreement with the company.[142] Alibaba reportedly plans to charge large commercial users of its next Qwen release, while Moonshot AI was reportedly seeking up to 30 percent of revenue generated from Kimi K3 services offered through major American cloud providers.[143]

Dependence on a PRC model family could also leave users exposed to decisions made in Beijing. In July 2026, Reuters reported that PRC authorities, in

[135] *Procurement Intent Announcement for Computing Servers* [算力服务器采购采购意向公告], Mil. Procurement Network (Jan. 19, 2026), https://archive.org/details/20260119_202610.
[136] *Second Procurement Intent Announcement for a Certain Cybersecurity Professional Capability Generation Support System* [某网络安全专业能力生成支持系统采购意向公告（二次）], Mil. Procurement Network (July 10, 2026), https://archive.org/details/20260710_202610.
[137] *Request for Comments on Upgrade of a Certain Large Model (First Round)* [某大模型升级征求意见公告（第一次）], Mil. Procurement Network (May 7, 2026), https://archive.org/details/20260507_202610.
[138] *Procurement Intent Announcement for Supporting Software Procurement Project* [支撑软件采购项目采购意向公告], Mil. Procurement Network (July 2, 2026), https://archive.org/details/20260702_202610.
[139] *Procurement Intent Announcement for AI Large-Model Inference Operating Environment* ["AI大模型推理运行环境"采购需求 意向公开采购意向公告], Mil. Procurement Network (Oct. 14, 2024), https://archive.org/details/20241014_202610.
[140] *See, e.g.*, *Tender Announcement for AI Large-Model Foundation Platform Information-System Integration (Goods) Procurement Project* [AI大模型基础平台信息系统集成（物资类）采购项目招标公告], Mil. Procurement Network (June 18, 2025), https://archive.org/details/20250618_202610; *Procurement Intent Announcement for Hardware Equipment Procurement* [硬件设备采购采购意向公告], Mil. Procurement Network (Apr. 29, 2025), https://archive.org/details/20250429_202610; *Service Requirements Public Notice for a Certain Unit's Local Deployment of the DeepSeek Model (2025-JHBNJB-W4001) (Package 1)* [某单位deepseek模型本地化部署项目服务需求公示(2025-JHBNJB-W4001)(第1包)], Mil. Procurement Network (Apr. 25, 2025), https://archive.org/details/20250425_202610.
[141] *See* Microsoft AI Diffusion Report, *supra* note 97; Caiwei Chen, *supra* note 80.
[142] Kimi K3 License, *supra* note 19.
[143] Nellis & Baptista, *supra* note 19; Baptista et al., *supra* note 19.

meetings led by the Ministry of Commerce, had discussed restricting overseas access to the country's most advanced AI models, including Alibaba's Qwen and Zhipu's GLM-5.2.[144] For American companies that build products around a PRC model family, this creates a continuity risk controlled not only by the developer but also by the PRC government. Even if previously downloaded weights remain available, companies could lose access to future releases, security updates, or improvements. DoorDash's estimate that replacing a single PRC-developed model embedded in a larger system could take up to six months illustrates how quickly technical adoption can become operational dependence.[145]

That dependence would not be limited to the United States. A January 2026 Microsoft analysis found that DeepSeek had gained meaningful market share across a number of developing countries, including shares of roughly 11 to 14 percent in several African markets and substantially higher shares in countries including Belarus, Cuba, and Russia.[146] Wider adoption of PRC-developed models abroad would extend the same dependency questions to foreign governments, companies, and digital ecosystems with which American firms interact.

These dynamics complicate the U.S. policy response. As **Findings 2 and 3** show, the United States does not yet have a comparably competitive open-weight ecosystem, while PRC developers have also benefited from capabilities extracted from American frontier models.[147] Policymakers therefore face the challenge of reducing dependence on PRC-developed models while managing the cost of replacing tools on which American companies have already come to rely. In July 2026, after the release of Moonshot AI's Kimi K3, administration officials publicly raised concerns about illicit distillation and unauthorized access to advanced American chips, while press reports indicated that restrictions on PRC-developed models were under consideration.[148] The Secretary of the Treasury warned that "open source is not open season on American IP" and indicated that sanctions remained possible.[149] At the same time, administration officials continued to support a competitive ecosystem that includes open-weight models.[150] In October 2026, DHS made public an order, issued in July under the Federal Acquisition Supply Chain Security Act of 2018, that excludes DeepSeek products and services from federal civilian agency procurement and requires agencies to remove them from their information systems, including third-party products and services built using DeepSeek.[151] The order does not apply to private use or to Department of War systems.[152]

---

[144] Potkin, *supra* note 19.
[145] DoorDash Briefing, *supra* note 6; DoorDash Nov. 14 Email, *supra* note 6.
[146] Microsoft AI Diffusion Report, *supra* note 97; Chan Ho-Him & Matt O'Brien, *DeepSeek's AI Gains Traction in Developing Nations, Microsoft Report Says*, Associated Press (Jan. 8, 2026).
[147] Joint Distillation Advisory, *supra* note 14; Artificial Analysis Index, *supra* note 12; Lambert, *supra* note 58.
[148] Maria Curi, *The Secret Trump Administration Battle to Fight Chinese AI*, Axios (July 20, 2026); Kratsios, *supra* note 82; *see* Nicol-Schwarz, *supra* note 82.
[149] Scott Bessent (@SecScottBessent), X (July 22, 2026), https://perma.cc/7MPR-XYLS.
[150] AI Action Plan, *supra* note 79, at 4; Sophia Cai, *Startup Founders Urge Trump Not to Shut Off Chinese Open Weight AI*, Politico (July 22, 2026).
[151] 91 Fed. Reg. at 66,346; 41 U.S.C. §§ 1321–1328; Lee & Nicol-Schwarz, *supra* note 23.
[152] *See* Lee & Nicol-Schwarz, *supra* note 23; Bill Toulas, *DHS Bans DeepSeek Products Across US Federal Civilian Agencies*, BleepingComputer (Oct. 20, 2026).

Industry warned against broad restrictions. Nearly 200 startups and technology companies urged the administration not to cut off access to PRC open-weight models, arguing that many young companies had already become dependent on them.[153] That figure likely understates the scale of U.S. industry exposure. As **Finding 1** explains, open-weight models can be downloaded, modified, and embedded without an ongoing commercial relationship or other conventional record of use, making much of that dependence difficult to detect. The number of American companies relying directly or indirectly on PRC-developed models may therefore be substantially larger than public disclosures suggest, and a disruption in access to future releases, updates, or related services could impose costs across a much broader segment of the American economy. Major American technology companies separately cautioned against "premature restrictions," emphasizing the value of capable open models for innovation and cybersecurity.[154]

The resulting policy challenge is how to reduce that dependence quickly while ensuring that trusted American and allied open-weight alternatives are available to replace what is removed.

Finding 5. The risks are increasing as PRC models become more capable, while their behavior and security properties can differ in ways ordinary evaluations may not detect.

Independent testing indicates that constraints introduced during development can persist in ways that conventional performance benchmarks may not reveal. CrowdStrike found that DeepSeek's flagship reasoning model was up to 50 percent more likely to produce code with severe security flaws when a prompt referenced topics the CCP considers politically sensitive.[155] Booz Allen Hamilton found that, in more than 2,800 trials, three of four PRC models it tested produced significantly more vulnerable code when the user identified as a member of the U.S. government, and that the vulnerabilities were highly obfuscated.[156] NIST found that DeepSeek models echoed CCP narratives four times as often as American reference models and were far more susceptible to hijacking in agent workflows.[157] CrowdStrike concluded that the behavior it observed was more likely an unintended consequence of training the model to adhere to CCP political requirements, like "Core Socialist Values."[158] The security consequence is nevertheless the same. A model can perform normally on standard evaluations yet degrade under conditions that a user may not know to test.[159]

---

[153] Cai, *supra* note 150.
[154] Bellan, *Industry Urges Against Broad Open-Weight Restrictions*, *supra* note 62.
[155] Stein, *supra* note 20.
[156] Booz Allen Press Release, *supra* note 20.
[157] Ctr. for AI Standards & Innovation, Nat'l Inst. of Standards & Tech., *Evaluation of DeepSeek AI Models* (Sept. 30, 2025), https://perma.cc/CY4X-YMWH [hereinafter NIST DeepSeek Evaluation].
[158] Stein, *supra* note 20; *see also* Cyberspace Admin. of China et al., Interim Measures for the Management of Generative Artificial Intelligence Services [生成式人工智能服务管理暂行办法] arts. 4, 14, 17 (July 13, 2023), https://perma.cc/7LHH-HR4A [hereinafter Interim Measures].
[159] Stein, *supra* note 20; Booz Allen Press Release, *supra* note 20; *see also* Evan Hubinger et al., *Sleeper Agents: Training Deceptive LLMs That Persist Through Safety Training*, arXiv:2401.05566 (Jan. 10, 2024), https://arxiv.org/abs/2401.05566.

The stakes have also increased as PRC open-weight models have become more capable. In November 2025, an American frontier developer disclosed that a group it assessed with high confidence to be a PRC state-sponsored actor had used its AI system in a largely autonomous cyber espionage campaign against approximately 30 targets, with AI agents performing most of the intrusion activity.[160] The Committee on Homeland Security responded with oversight letters and a December 2025 hearing.[161] Committee Members and staff later received classified briefings on advanced cyber-capable models and the risks posed by jailbroken systems.[162] Then, in April 2026, the same developer announced Mythos Preview, a model capable of autonomously discovering and exploiting serious vulnerabilities in critical software. It concluded that the capability was too dangerous for general release and limited access to selected cyber defenders.[163] Five months later, NIST and the developer separately concluded that Zhipu's openly available GLM-5.3 had reached comparable offensive cyber capability in important respects.[164]

NIST described GLM-5.3 as capable of autonomously building end-to-end cyber exploits and found that it trailed the American frontier by only about four months on an aggregate of government cyber benchmarks.[165] Because the model's weights are public, researchers were able to reduce its refusal rate on harmful requests from approximately 95 percent to between 3 and 6 percent using roughly 2,200 hours of graphics processor time at an estimated cost of about $4,400.[166] The same American developer warned that the advanced cyber capabilities it had restricted in April "have now arrived" in open weights.[167] The progression from AI-assisted intrusion activity to downloadable models capable of autonomous exploit development sharply reduces the time between the emergence of a frontier capability and its broad availability.

Deployment architecture changes some of these risks, but not all of them. The companies that the Committees examined largely avoided sending data to PRC-operated services by running models on their own infrastructure or through U.S.-based cloud and inference providers.[168] Those deployment choices substantially reduce the most direct form of data exposure. They do not, however, resolve the broader provenance and dependency risks described in Findings 1 and 4, or change how the underlying model was trained and aligned.

---

[160] Anthropic, *Disrupting the First Reported AI-Orchestrated Cyber Espionage Campaign* (Nov. 13, 2025), https://perma.cc/J2RH-QG8U [hereinafter Anthropic Espionage Disclosure].

[161] Press Release, H. Comm. on Homeland Sec., *Homeland Republicans Request Anthropic, Google, Quantum Xchange Testimony Following Report of AI-Assisted, Partially Autonomous PRC Cyber Operation* (Nov. 26, 2025), https://homeland.house.gov/2025/11/26/homeland-republicans-request-anthropic-google-quantum-xchange-testimony-following-report-of-ai-assisted-partially-autonomous-prc-cyber-operation/ [hereinafter Testimony Request Press Release]; December 2025 Hearing, *supra* note 36.

[162] *See* Sabin & Gold, *supra* note 37; DiMolfetta, *House Homeland Panel*, *supra* note 37.

[163] Anthropic, *Project Glasswing: Securing Critical Software for the AI Era* (Apr. 7, 2026), https://perma.cc/BT4U-GCEK [hereinafter Project Glasswing]; Nicholas Carlini et al., *Assessing Claude Mythos Preview's Cybersecurity Capabilities*, Anthropic (Apr. 7, 2026), https://perma.cc/2VJ3-Z43J.

[164] NIST GLM-5.3 Assessment, *supra* note 21; Anthropic GLM-5.3 Analysis, *supra* note 21.

[165] NIST GLM-5.3 Assessment, *supra* note 21; Anthropic GLM-5.3 Analysis, *supra* note 21.

[166] Anthropic GLM-5.3 Analysis, *supra* note 21.

[167] John Sakellariadis, *Anthropic Says New Chinese Model Can Hack at Mythos Levels*, Politico (Sept. 30, 2026); Anthropic GLM-5.3 Analysis, *supra* note 21.

[168] DoorDash Response, *supra* note 6; DoorDash Production, *supra* note 6, attach. B; Airbnb Response, *supra* note 31, at 3–4; Anysphere, Inc., *Security*, Cursor (Aug. 25, 2026), https://perma.cc/8YPP-SQ8U [hereinafter Cursor Security Page]; Anysphere Briefing, *supra* note 33.

Users who access PRC models through PRC-operated services face additional risks. Requests sent to those services may place prompts, proprietary documents, and personal data on infrastructure controlled by PRC companies subject to PRC intelligence and data security laws.[169] At scale, those interactions can create an intelligence collection opportunity across American industry, research, and government.[170] Recent evidence illustrates how those risks can materialize. In September 2026, an American frontier developer reported that PRC laboratories had silently relayed customer requests to third parties, including sessions originating from routing services used in the United States and Europe that contained names, email addresses, company information, and other sensitive data.[171] The Select Committee separately found that DeepSeek transmitted American user data to the PRC through infrastructure connected to state-owned China Mobile.[172]

Exposure can also persist through the personnel, affiliates, and services surrounding a deployment. Airbnb retained a PRC office focused on product development after winding down its domestic lodging business there in 2022, and current job listings indicate that its PRC-based team develops AI for customer service, including fine-tuning and evaluating large language models.[173] The company disclosed in 2021 that PRC government agencies may require or request that it share information about users who use or offer services in the PRC through its platform.[174] According to Airbnb, its PRC-based employees have access only to a limited subset of Airbnb-approved models, and its AI models do not send user data or related information to any PRC-based entity for processing.[175] The Committees have not established whether PRC-based personnel evaluate, fine-tune, deploy, or maintain the PRC-developed models identified in Airbnb's production, or whether they can access data about users outside the PRC processed by those systems.[176]

Large technology companies have greater capacity to mitigate some of these risks through extensive testing, fine-tuning, and layered deployment controls. DoorDash described treating model inputs, outputs, and proposed actions as untrusted, validating model files with cryptographic hashes, permitting only secure file formats, and relying on overlapping guardrails.[177] Airbnb described a cross-functional review process and additional safety checks for customer-facing

[169] National Intelligence Law of the People's Republic of China [中华人民共和国国家情报法] art. 7 (adopted June 27, 2017, amended Apr. 27, 2018), https://perma.cc/H76S-NJTU; Data Security Law of the People's Republic of China [中华人民共和国数据安全法] art. 35 (adopted June 10, 2021, effective Sept. 1, 2021), https://perma.cc/83KP-5W5P.
[170] *See* Nat'l Counterintelligence & Sec. Ctr., *National Counterintelligence Strategy of the United States 2024–2027* (2024), https://www.odni.gov/files/NCSC/documents/features/NCSC_CI_Strategy-pages-20240730.pdf [hereinafter National Counterintelligence Strategy]; H. Select Comm. on the Strategic Competition Between the U.S. & the Chinese Communist Party, *DeepSeek Unmasked: Exposing the CCP's Latest Tool for Spying, Stealing, and Subverting U.S. Export Control Restrictions* (Apr. 16, 2025) [hereinafter DeepSeek Unmasked].
[171] Anthropic September 2026 Report, *supra* note 87.
[172] DeepSeek Unmasked, *supra* note 170; Martina & Nellis, *supra* note 134.
[173] *Airbnb to Shut Domestic Business in China from July 30*, Reuters (May 24, 2022); Airbnb, Inc., *Senior/Staff Machine Learning Engineer, Community Support Engineering*, Airbnb Careers, https://perma.cc/W8QZ-P3QV; Airbnb, Inc., *Machine Learning Engineer, Assistance Engineering*, Airbnb Careers, https://perma.cc/RQM7-PWZL; Airbnb, Inc., *Senior Machine Learning Engineer, Assistance Engineering*, Airbnb Careers, https://perma.cc/3JH8-H7H4; Airbnb Briefing, *supra* note 31; Airbnb Nov. 13 Email, *supra* note 4.
[174] Airbnb, Inc., Annual Report (Form 10-K) (Feb. 25, 2021), https://perma.cc/6UHD-QVNA.
[175] Airbnb Nov. 13 Email, *supra* note 4; *see also* Airbnb Briefing, *supra* note 31; Airbnb Response, *supra* note 31, at 2–4.
[176] Airbnb Briefing, *supra* note 31.
[177] DoorDash Response, *supra* note 6; DoorDash Production, *supra* note 6, attach. B.

deployments.[178] Smaller downstream users may lack dedicated AI security teams, the resources to conduct extensive adversarial testing, or the expertise to identify context-triggered behavior.[179] CrowdStrike cautioned that generic open-source benchmarks are insufficient and recommended testing AI-coding assistants within the environments in which they will actually operate.[180]

Those controls, however, do not substitute for evaluating the developer itself. The companies' submissions reflected substantial attention to model performance, deployment architecture, and security controls.[181] Airbnb described reviews covering security, privacy, legal, and open-source licensing.[182] DoorDash stated that it "never assumes a model is 'safe' or 'secure'" and described controls for validating model provenance and integrity.[183] Anysphere described a vendor risk management program and technical controls governing infrastructure and model access.[184] Their submissions provided considerably less detail about how they evaluated the PRC developers themselves.[185]

When the White House publicly accused Moonshot AI of illicitly distilling an American model and acquiring restricted chips, DoorDash told the Committees that it had no independent knowledge of the allegations, that the public disclosures did not provide enough information to identify specific security risks from its use of the model, and that it had not changed the deployment, permissions, or security monitoring of the Moonshot model it uses.[186] Airbnb also continued to list Alibaba models as approved for production use after the Department of War identified Alibaba as a Chinese military company, though the designation itself did not prohibit U.S. companies from using Alibaba's models.[187] By contrast, Anysphere informed the Committees that it plans to explore phasing out PRC-developed open-weight models because of geopolitical risk.[188]

Those differences reflect a broader information problem. Model-level testing can identify how a system behaves under the conditions tested. Developer-level diligence addresses different questions, including how the model was developed, what legal obligations shaped that process, what institutions supported it, and what risks may have been introduced upstream. The federal government currently lacks a dedicated channel for delivering developer-specific AI threat information to industry, while the principal federal cybersecurity information-sharing statute

[178] Airbnb Response, *supra* note 31, at 3–4.
[179] Meyer & Buhler, *supra* note 7; Cai, *supra* note 150.
[180] Stein, *supra* note 20.
[181] Airbnb June 30 Email, *supra* note 31; DoorDash Response, *supra* note 6; DoorDash Production, *supra* note 6, attachs. B, D.
[182] Airbnb Response, *supra* note 31, at 3–4; Airbnb, Inc., *Responsible AI Use Policy*, ABNB_HSCCCP_2026_0000113–0000117 (produced to the Committees).
[183] DoorDash Response, *supra* note 6.
[184] Cursor Security Page, *supra* note 168.
[185] *See generally* Airbnb Response, *supra* note 31; DoorDash Response, *supra* note 6; DoorDash Production, *supra* note 6, attach. D; Cursor Security Page, *supra* note 168.
[186] DoorDash Response, *supra* note 6; DoorDash Nov. 14 Email, *supra* note 6; Kratsios, *supra* note 82; *see also* Nicol-Schwarz, *supra* note 82.
[187] Airbnb Approved Models List, *supra* note 31; Section 1260H List, *supra* note 96; Notice of Availability of Designation of Chinese Military Companies, 91 Fed. Reg. 35,189 (June 10, 2026); Airbnb Nov. 13 Email, *supra* note 4.
[188] Anysphere Briefing, *supra* note 33.

has operated under short-term extensions since its original authorization expired in September 2025.[189]

Together, the evidence shows that secure deployment is necessary but not sufficient. Self-hosting can reduce direct data exposure, and layered controls can reduce operational risk, but neither resolves risks arising from model behavior, rapidly proliferating capabilities, developer provenance, or the legal and institutional environment in which the model was produced. Those risks become more consequential as PRC-developed open-weight models approach the capabilities of the most advanced American systems.

[189] Cybersecurity & Infrastructure Sec. Agency, *JCDC AI Cybersecurity Collaboration Playbook* (Jan. 14, 2025), https://perma.cc/5E5G-2SHV [hereinafter JCDC AI Playbook]; AI Action Plan, *supra* note 79; Cybersecurity Information Sharing Act of 2015, Pub. L. No. 114-113, div. N, tit. I, 129 Stat. 2242 (codified as amended at 6 U.S.C. §§ 1501–1510); David DiMolfetta, *Stopgap Funding Bill Temporarily Extends Key Cyber Info-Sharing Law*, Nextgov/FCW (Sept. 1, 2026); Justin Doubleday, *CR Extends Cyber Info Sharing Law Through December*, Fed. News Network (Sept. 2026).

RECOMMENDATIONS

The findings above present a familiar policy problem. Capable, low-cost PRC technology can become deeply embedded before the United States has competitive alternatives or a practical plan to unwind that dependence. The experiences with Huawei, ZTE, and DJI show the costs of waiting until dependence is already entrenched.[190] Open-weight AI presents many of the same challenges, but with an important difference. Model weights can be copied, modified, and embedded in downstream products in ways that are far harder to inventory or remove than physical equipment.

Table 5. Comparing PRC Technology Dependencies

	Huawei and ZTE	DJI	PRC Open-Weight AI Models
Why users adopted it	Low-cost, capable telecom equipment	Low-cost, capable drones	Low-cost, capable, customizable models
When concern emerged	**2012:** House Intelligence Committee warning	**2017–2025:** Growing federal security restrictions	**2026:** White House and joint federal warnings
What happened next	**2020:** Secure and Trusted Communications Networks Act and removal program	**2025:** Addition to FCC Covered List (December 2025)	Under consideration
How entrenched dependence became	Equipment remained embedded in rural networks	DJI remained widely used by public-safety and commercial users	Full scale is unknown
Cost of unwinding	Est. $4.98 billion	Est. 6× higher average replacement cost	Unknown and growing

Sources: Committees' analysis.[191]

Dependence becomes harder and more expensive to unwind when there are no viable alternatives available. In 2012, the House Permanent Select Committee on Intelligence recommended that American companies avoid Huawei and ZTE equipment because of national security concerns.[192] Many small and rural providers nevertheless continued to rely on that equipment because it was affordable, reliable, and difficult to replace at comparable cost.[193] Congress later enacted the Secure and Trusted Communications Networks Act of 2019 to fund

[190] *See infra* notes 191–199 and accompanying text.
[191] H. Permanent Select Comm. on Intel., 112th Cong., *Investigative Report on the U.S. National Security Issues Posed by Chinese Telecommunications Companies Huawei and ZTE* (Oct. 8, 2012) [hereinafter HPSCI Huawei-ZTE Report]; Secure and Trusted Communications Networks Act of 2019, Pub. L. No. 116-124, 134 Stat. 158 (2020); Letter from Jessica Rosenworcel, Chairwoman, Fed. Commc'ns Comm'n, to Frank Pallone, Jr., Chairman, H. Comm. on Energy & Com., et al. (July 15, 2022) [hereinafter Rosenworcel Letter] (estimating approximately $4.98 billion in reasonable and supported removal-and-replacement costs and a $3.08 billion funding shortfall); U.S. Gov't Accountability Off., GAO-24-106924, *Federal Lands: Effects of Interior's Policies on Foreign-Made Drones* 6 (Sept. 25, 2024) [hereinafter GAO Drone Report]; Public Safety & Homeland Sec. Bureau, Fed. Commc'ns Comm'n, *Public Safety and Homeland Security Bureau Announces Addition of Uncrewed Aircraft Systems (UAS) and UAS Critical Components Produced Abroad, and Equipment and Services Listed in Section 1709 of the FY2025 NDAA, to FCC Covered List*, Public Notice, DA 25-1086 (Dec. 22, 2025) [hereinafter FCC Covered List Notice]; NSTM-4, *supra* note 38; Joint Distillation Advisory, *supra* note 14.
[192] HPSCI Huawei-ZTE Report, *supra* note 191, at iv–vi.
[193] *See* Mike Dano, *FCC Warns of 'Significant Consequences' Without More Rip-and-Replace Funds*, Light Reading (May 8, 2024), https://www.lightreading.com/wireless/fcc-warns-of-significant-consequences-without-more-rip-and-replace-funds; Rosenworcel Letter, *supra* note 191.

removal and replacement.[194] By then, however, the estimated cost had reached approximately $4.98 billion, far exceeding the original $1.9 billion appropriation and leaving a $3.08 billion shortfall.[195] Congress authorized additional funding in December 2024, more than a decade after the original warning, and removal remained incomplete.[196]

The experience with DJI illustrates a related problem. Dependence is hardest to unwind when users have few viable substitutes. After the Department of the Interior restricted its own use of foreign-made drones, the average cost of compliant replacements rose from about $2,600 to more than $15,000 per drone, nearly six times the prior average.[197] Public safety agencies nationwide, including state and local law enforcement, and small businesses likewise resisted broader restrictions on DJI because lower-cost domestic alternatives were not yet available at comparable scale or capability.[198] Recent reporting describes similar challenges as American companies seek to replace PRC components in robotics and other industries.[199]

Dependence on AI systems can be far harder to unwind. Unlike telecommunications equipment or drones, models can be copied, modified, embedded several layers deep in other products, and run anywhere, making even their use difficult to inventory. Hardware restrictions can be enforced at the border or through equipment authorization.[200] No comparable chokepoint exists for open weights. The lesson from Huawei, ZTE, and DJI is therefore to build competitive alternatives, establish visibility, and plan for substitution before dependence becomes too costly to unwind.

A. Build the American Alternative

Recommendation 1. Launch a national American open-weight AI initiative. The most effective response to inexpensive PRC models is a competitive American alternative. That alternative is also what will make broader restrictions on PRC-developed models workable. Startups, technology companies, and cybersecurity firms have warned that they lack American or allied open-weight models with comparable price and performance, and that removing PRC models without those alternatives in place would raise costs and concentrate the market in a small number of closed, proprietary providers.[201] Recent releases from Reflection AI and Mistral AI show that the gap is beginning to close.[202] Congress should authorize, and the executive branch should establish, a public-private initiative to accelerate

[194] Secure and Trusted Communications Networks Act of 2019, Pub. L. No. 116-124, 134 Stat. 158 (2020).
[195] Rosenworcel Letter, *supra* note 191; Fed. Commc'ns Comm'n, *Fiscal Year 2025 Congressional Budget Justification* (2024) [hereinafter FCC FY2025 Budget Justification]; *see also* Consolidated Appropriations Act, 2021, Pub. L. No. 116-260, div. N, § 906, 134 Stat. 1182 (2020).
[196] Servicemember Quality of Life Improvement and National Defense Authorization Act for Fiscal Year 2025, Pub. L. No. 118-159, 138 Stat. 1773 (2024); Wireline Competition Bureau & Office of Managing Director Announce Additional Congressional Funding for the Secure and Trusted Communications Networks Reimbursement Program, Public Notice, DA 24-1279 (Fed. Commc'ns Comm'n Dec. 26, 2024).
[197] GAO Drone Report, *supra* note 191, at 6.
[198] Fed. Commc'ns Comm'n, *Fact Sheet: FCC Updates Covered List to Include Foreign UAS and UAS Critical Components on Going Forward Basis* (Dec. 22, 2025) [hereinafter FCC Drone Fact Sheet]; Didi Tang, *Many Americans Have Come to Rely on Chinese-Made Drones. Now Lawmakers Want to Ban Them*, Associated Press (June 24, 2024).
[199] Stephanie Yang, *America Wants to Wean Itself Off Chinese Technology. Will the Pain Pay Off?*, CNN Business (Oct. 3, 2026).
[200] *See* FCC Covered List Notice, *supra* note 191; FCC Drone Fact Sheet, *supra* note 198.
[201] *See* Cai, *supra* note 150; Bellan, *Industry Urges Against Broad Open-Weight Restrictions*, *supra* note 62; Meyer & Buhler, *supra* note 7.
[202] Bellan, *Reflection Debuts Beam*, *supra* note 76; Howcroft & Marchandon, *supra* note 78.

the development of frontier competitive American open-weight models. The initiative should include an interagency coordination body, access to public computing resources for qualified American developers and researchers, a shared testbed and evaluation infrastructure, and an open-weight AI challenge to accelerate talent and research. DHS, through CISA, should participate in the shared testbed so that the security of American open-weight models, including their resistance to misuse against critical infrastructure, is evaluated alongside their performance. The initiative should also explore a structured framework under which American frontier developers may voluntarily license model outputs to qualified American open-weight developers for training, giving American firms a lawful path to the kind of capability transfer that PRC developers have obtained illicitly.[203] Federal agencies should serve as anchor customers for secure American open-weight models where appropriate. These efforts would advance the AI Action Plan's call for a supportive environment for open models.[204]

Recommendation 2. Make trusted American AI the global default. The executive branch should use export promotion, development finance, and diplomatic engagement to help partner nations adopt trusted American AI, including open-weight models suitable for sovereign deployment. This effort should prioritize countries where PRC models are gaining market share and should coordinate with allied governments that are developing their own AI capacity.[205]

B. Know What Is Running

Recommendation 3. Establish a model provenance standard and an AI bill of materials. Congress should direct NIST, in coordination with DHS through CISA, to develop a standard for documenting the lineage of AI models, including base models, significant fine-tuning, and distillation sources, and a machine-readable AI bill of materials that identifies the models embedded in a product or service. The standard should be required for federal contractors and for AI used in high-risk functions of critical infrastructure, and voluntary elsewhere. NIST should also support research into technical methods for verifying model lineage.[206] Requirements should be scaled to minimize burdens on startups, researchers, and open-source developers.

Recommendation 4. Extend federal restrictions beyond DeepSeek, keep PRC-developed AI models out of federal contract work and high-risk functions of critical infrastructure, and give agencies the visibility to enforce these restrictions. The executive branch has already taken an important step. DHS has ordered federal civilian agencies to exclude DeepSeek products and services from procurement and to remove them from agency information systems, including third-party products built using DeepSeek.[207] Congress has separately prohibited

[203] Meyer & Buhler, *supra* note 7; Lambert, *supra* note 58.
[204] AI Action Plan, *supra* note 79, at 4.
[205] Microsoft AI Diffusion Report, *supra* note 97; Caiwei Chen, *supra* note 80.
[206] Nikolic et al., *supra* note 10; Chang, *supra* note 10.
[207] 91 Fed. Reg. at 66,346; Lee & Nicol-Schwarz, *supra* note 23.

DeepSeek on intelligence community systems.[208] These actions reach one developer. Other leading PRC developers are subject to the same PRC laws, were identified in the same distillation campaigns, and have released openly available models with comparable or greater offensive cyber capability.[209] DHS and the Federal Acquisition Security Council should evaluate those developers for comparable orders.

Congress should prohibit federal agencies from using AI models developed by entities subject to the jurisdiction or control of the PRC, including derivatives of those models, in federal information systems. The No Adversarial AI Act, which was introduced in June 2025, would move in this direction by requiring the Federal Acquisition Security Council to maintain a public list of AI developed by foreign adversaries and directing agencies to consider that AI for exclusion and removal.[210] The prohibition should include a narrow and transparent waiver process for research, evaluation, testing, and defensive cybersecurity work. A prohibition is only as effective as an agency's ability to see what is running on its networks. Congress should therefore codify and modernize CISA's CDM program so that it provides agencies with continuous discovery and inventory of AI systems, models, and agents, identification of unauthorized or unapproved AI deployments, adversarial testing of AI systems before and after deployment, and supply chain risk management for high-risk vendors, models, and software components. These capabilities would allow CISA and agencies to identify PRC-developed models and their derivatives wherever they appear in federal civilian systems, including models embedded in commercial products several layers below what an agency can readily see.

Federal contractors should be held to the same standard. Congress should prohibit contractors from using PRC-developed models in the performance of federal contracts or in information systems that process, store, or transmit federal information. Congress has already taken this approach for DeepSeek, barring Department of War contractors from using it in the performance of their contracts, and federal acquisition rules apply the same scope to TikTok.[211] This requirement should take effect first for the largest contractors, those whose annual federal contract obligations exceed a threshold set in statute, and extend to all contractors after a defined transition period, with additional time and technical assistance for small businesses. The same narrow waiver process should apply, consistent with the waivers Congress authorized for DeepSeek.[212] During the transition, contractors that use PRC-developed models in covered systems should disclose that use to the contracting agency and maintain a plan to replace those models within a defined period. Those plans should account for a PRC government

[208] National Defense Authorization Act for Fiscal Year 2026, Pub. L. No. 119-60, div. F, § 6604 (2025).
[209] *See* National Intelligence Law of the People's Republic of China [中华人民共和国国家情报法] art. 7 (adopted June 27, 2017, amended Apr. 27, 2018), https://perma.cc/H76S-NJTU; Data Security Law of the People's Republic of China [中华人民共和国数据安全法] art. 35 (adopted June 10, 2021, effective Sept. 1, 2021), https://perma.cc/83KP-5W5P; Joint Distillation Advisory, *supra* note 14, at 1–2; NIST GLM-5.3 Assessment, *supra* note 21; Anthropic GLM-5.3 Analysis, *supra* note 21.
[210] No Adversarial AI Act, H.R. 4142, 119th Cong. §§ 2–3 (2025); No Adversarial AI Act, S. 2177, 119th Cong. (2025).
[211] National Defense Authorization Act for Fiscal Year 2026, Pub. L. No. 119-60, § 1532(a)(3) (2025); 48 C.F.R. § 52.204-27(b) (2025) (implementing Consolidated Appropriations Act, 2023, Pub. L. No. 117-328, div. R, § 102, 136 Stat. 4459, 5258 (2022)).
[212] National Defense Authorization Act for Fiscal Year 2026, Pub. L. No. 119-60, § 1532(b) (2025).

decision to restrict access, a material change in license terms, discovery of a serious vulnerability or hidden behavior, and new U.S. legal restrictions.[213]

These restrictions depend on a clear definition of a PRC-developed model. The definition should cover models whose original weights were trained by an entity organized, headquartered, or principally operating in the PRC, or otherwise subject to the jurisdiction or control of the PRC, and models built from those weights through fine-tuning, continued pretraining, merging, or compression. It should not reach an American or allied model solely because some of its training data was generated by a PRC-developed model, although that use should be documented under the provenance standard in Recommendation 3. Developers should also have a way to show that a model is not covered.

The same principle should extend to the most sensitive functions of critical infrastructure. Congress should prohibit the use of PRC-developed AI models, including their derivatives, in high-risk functions of critical infrastructure. High-risk functions should be defined narrowly, in coordination with sector risk management agencies, and should focus on uses where a compromised or manipulated model could cause physical harm or disrupt essential services. These include AI agents with access to operational technology, industrial control and safety systems, and tools that write or review code deployed in those environments. The prohibition should apply immediately to new deployments in those functions. Owners and operators already using PRC-developed models in those functions should identify them to their sector risk management agency and to CISA and remove them on a defined schedule. That information should receive the protections Congress provided for reports under CIRCIA, described in Recommendation 8. Where no trusted alternative yet exists, DHS, acting through CISA, should be able to grant an exception based on a CAISSI evaluation of the specific model and on compensating controls, such as running the model only on infrastructure the operator controls, isolating it from systems it does not need to reach, and requiring human approval before it takes consequential actions. Each evaluation should be conducted once and shared with all owners and operators, so that a utility, hospital, or water system does not have to assess the same model on its own. Exceptions should expire on a fixed date and should be reported to Congress.

The requirement should be built to work in practice. Removal schedules should account for the availability of trusted alternatives and the time needed to test and certify new systems in safety-critical environments, so that owners and operators are not forced to remove systems faster than they can safely replace them. It should be implemented through sector risk management agencies and existing sector regulators in close coordination with CISA. Congress should provide CAISSI and CISA the resources needed to evaluate models at the pace they are released. A single federal process should be established rather than a

[213] Potkin, *supra* note 19; Kimi K3 License, *supra* note 19; DoorDash Briefing, *supra* note 6.

patchwork of state and sector rules that would impose different requirements on the same companies.

C. Share Threat Information and Strengthen Defenses

Recommendation 5. Secure long-term information-sharing authorities and build a developer-level AI threat channel. Any effective public-private response to the risks described in this report depends on the voluntary information-sharing framework established by CISA 2015. Its liability protections, antitrust exemption, and privacy safeguards give companies the legal certainty they need to share cyber threat indicators and defensive measures with DHS and with one another.[214] The joint advisory on PRC distillation campaigns specifically called for cross-organization sharing of infrastructure and behavioral indicators among model developers, cloud providers, and API aggregators, and that sharing rests on the protections CISA 2015 provides.[215] Since the law's original authorization expired on September 30, 2025, it has lapsed once and has otherwise been sustained only through a series of short-term extensions, most recently through December 11, 2026.[216] Short-term extensions create uncertainty for the companies whose participation the framework depends on, at the very moment the AI threat environment demands more sharing, not less.

The Committees' first priority relating to AI cyber threats is therefore a long-term reauthorization of CISA 2015 through the WIMWIG Act. The Committee on Homeland Security approved the WIMWIG Act by a vote of 25 to 0 on September 3, 2025,[217] and the House passed it on July 22, 2026, as part of the National Defense Authorization Act for Fiscal Year 2027.[218] The WIMWIG Act would extend these authorities for ten years and make targeted updates to reflect changes in technology and cybersecurity practice since 2015.[219] Congress should enact it without delay. A durable reauthorization should be prioritized, so that industry and government partners have the stability they need. Once long-term authorities are in place, Congress should examine whether the law's existing definitions reach the threats to AI systems that matter most to national security. Many of those threats, including the theft of model weights, the compromise of AI development infrastructure, malicious activity by AI agents, and the use of fraudulent accounts and proxy networks to extract model capabilities, involve unauthorized activity on information systems and may already fall within the law's definitions of cyber threat indicators and defensive measures.[220] Where they do not, any update should be narrow and targeted. It should be limited to sharing information about malicious actors, their methods, and the measures used to stop them, and it should preserve the law's existing privacy safeguards and its requirement that sharing serve a cybersecurity purpose.

---

[214] *See* 6 U.S.C. §§ 1503(d)–(e), 1504, 1505.
[215] Joint Distillation Advisory, *supra* note 14.
[216] Matishak, *supra* note 25; *Cyber Info-Sharing Legislation Clears House as Part of Defense Authorization Bill*, VitalLaw (July 23, 2026); DiMolfetta, *Stopgap Funding Bill*, *supra* note 189; Doubleday, *supra* note 189.
[217] H.R. 5079; Press Release, H. Comm. on Homeland Sec., *ICYMI: Committee Advances Legislation to Bolster Cyber Resilience, Public Safety, and DHS's Counterterrorism Mission* (Sept. 4, 2025), https://perma.cc/A69E-98CE [hereinafter Markup Press Release]; *House Panel Approves Cyber Information Sharing, Grant Legislation as Expiration Deadlines Loom*, CyberScoop (Sept. 3, 2025).
[218] H.R. 8800; Matishak, *supra* note 25; *Cyber Info-Sharing Legislation Clears House*, *supra* note 216.
[219] H.R. 5079; Markup Press Release, *supra* note 217; Matishak, *supra* note 25.
[220] *See* 6 U.S.C. § 1501(4)–(7).

Congress should also complete a long-term reauthorization of the SLCGP, which CISA administers with the Federal Emergency Management Agency.[221] The House passed the PILLAR Act on November 17, 2025, which would reauthorize SLCGP and expand its scope to address AI systems and operational technology.[222] State, local, tribal, and territorial governments face the same cost pressures that draw private companies to inexpensive PRC models, often with far fewer resources to evaluate them. Congress should prohibit the use of SLCGP funds to acquire, deploy, or maintain PRC-developed AI models or AI services operated by PRC developers, as federal grant rules already do for covered telecommunications equipment and as Congress has done for drones made by covered foreign entities.[223] As a condition of funding, recipients should identify PRC-developed models in grant-supported systems through the cybersecurity plans the program already requires and remove any used in high-risk functions on the schedule described in Recommendation 4.[224] Grant funds should be available to replace PRC-developed models with trusted alternatives, and CISA should make clear in program guidance that replacement is an eligible activity.

Building on these authorities, Congress should direct DHS, acting through CISA, in coordination with the Intelligence Community, the Department of Commerce, and NIST, to establish a mechanism for sharing timely, actionable information with industry about risks associated with specific AI developers and model families. That information should include indicators of illicit distillation, restricted computing acquisitions, and ties to foreign military or intelligence services, delivered at the lowest classification level possible. This mechanism should build on existing partnerships, including the Joint Cyber Defense Collaborative, and advance the AI Action Plan's call for DHS to lead an AI Information Sharing and Analysis Center.[225]

Recommendation 6. Evaluate PRC models systematically, translate the results into guidance, and give the Information Technology Sector a single accountable federal partner. Congress should direct CAISSI to evaluate leading PRC-developed models on a recurring basis, including for safeguard robustness, offensive cyber capability, and context-triggered behaviors such as degraded code security when a user identifies with the U.S. government. CISA should translate those evaluations into practical guidance for federal agencies and critical infrastructure owners and operators, developed with sector risk management agencies.[226] That guidance should advise owners and operators against deploying PRC-developed models in sensitive functions, set clear expectations for moving off models already in use as trusted alternatives become available, and specify the minimum controls expected for any use that continues in the meantime. These evaluations would also inform the exceptions described in Recommendation 4.

[221] *See* 6 U.S.C. § 665g; PILLAR Press Release, *supra* note 26.
[222] H.R. 5078; PILLAR Press Release, *supra* note 26; Colin Wood, *House Passes Bill to Reauthorize State and Local Cyber Grant Program*, StateScoop (Nov. 18, 2025), https://perma.cc/L7YW-CRFA.
[223] *See* 2 C.F.R. § 200.216 (2025); National Defense Authorization Act for Fiscal Year 2024, Pub. L. No. 118-31, § 1825, 137 Stat. 136 (2023).
[224] *See* 6 U.S.C. § 665g(e)(2)(B)(i).
[225] JCDC AI Playbook, *supra* note 189; AI Action Plan, *supra* note 79.
[226] NIST GLM-5.3 Assessment, *supra* note 21; NIST DeepSeek Evaluation, *supra* note 157; Booz Allen Press Release, *supra* note 20; Stein, *supra* note 20.

Guidance should address both closed and open models and should advise owners and operators to confirm that the AI services and routing services they use do not send their data to infrastructure operated by PRC developers.[227] Congress should also codify DHS, acting through CISA, as the Sector Risk Management Agency for the Information Technology Sector, with responsibilities that expressly include AI models and systems, AI compute infrastructure and data centers, cloud services, and software distribution, including open-source repositories. A codified role would give AI developers, cloud providers, and model repositories a single accountable federal partner for threat information, voluntary assessments, and incident coordination.

Recommendation 7. Prepare DHS and CISA to defend at machine speed and secure AI agents on federal networks. The cyber threat environment is entering a new phase. Open-weight models capable of autonomously building end-to-end cyber exploits are now available to anyone who downloads them.[228] Other nations are already confronting the results. In October 2026, South Korean authorities opened an investigation into cyberattacks on several of the country's largest banks, one of which exposed the personal credit information of about 25,000 customers, after the country's president said there were signs that AI models had been used.[229] The United States should expect a near future in which cyberattacks are planned, executed, and adapted by AI agents in minutes rather than months, and in which the volume and speed of attacks exceed what human defenders alone can manage.[230]

Defending federal networks and critical infrastructure in that environment will require defenders who can operate at the same speed, using AI to hunt for threats, detect intrusions, and respond to incidents as they unfold. The expertise needed to build and deploy those capabilities resides largely in the private sector, where frontier AI developers and cybersecurity companies are often months or years ahead of the government in understanding how these systems are built, how they behave, and how they are attacked.[231] The federal government cannot close that gap through traditional hiring alone.

Congress should establish a time-limited pilot program within DHS that uses direct-hire authority and temporary private sector detailees to bring experienced AI and emerging technology professionals into the Department, with a dedicated allocation for CISA, strong ethics and conflict-of-interest safeguards, and regular reporting to Congress. These professionals would help DHS and CISA evaluate models, detect illicit distillation, inform acquisition decisions, and apply AI to threat hunting, detection, and incident response across federal civilian networks and in support of critical infrastructure owners and operators. They would also transfer their knowledge to the career workforce so that it remains in the Department after their service ends.

---

[227] *See supra* notes 171–172 and accompanying text.
[228] NIST GLM-5.3 Assessment, *supra* note 21; Anthropic GLM-5.3 Analysis, *supra* note 21.
[229] Laura Chung & Jin Yu Young, *South Korea Investigates Possible Use of A.I. in Hackings on Its Banks*, N.Y. Times (Oct. 6, 2026).
[230] *See* Anthropic Espionage Disclosure, *supra* note 160; Anthropic GLM-5.3 Analysis, *supra* note 21.
[231] *See, e.g.*, Project Glasswing, *supra* note 163; Sabin & Gold, *supra* note 37.

As agencies deploy AI agents to defend their networks and deliver services, those agents will themselves become targets. Congress should direct CISA to serve as the operational lead for securing AI agents on federal civilian information systems. Agencies should be able to discover every AI agent operating on their networks, assign each one a durable identity and an accountable owner, limit its access to what its task requires, log its actions, and halt it immediately if it deviates from its assigned scope. Agencies should also verify the integrity of the models, tools, and software components that an agent loads or invokes, which would help detect PRC-developed models introduced through third-party agent platforms. These requirements should specify security outcomes rather than particular technologies or vendors.

Recommendation 8. Make CISA the single point of submission for AI-related cyber incident reports and establish an independent board to investigate significant incidents.

The federal government too often learns of significant AI-related cyber incidents from company blog posts, press reports, or voluntary disclosures, sometimes weeks or months after they occur.[232] Frontier AI developers and cybersecurity companies are frequently the first to detect these events, including the theft or attempted theft of model weights, the compromise of AI development and deployment infrastructure, AI agents that access systems without authorization, AI-enabled intrusions, and coordinated foreign efforts to extract model capabilities.[233] That information is most valuable in the hands of the agency that can act on it quickly. CISA can use it to warn other potential victims, share indicators across federal networks and critical infrastructure, and coordinate a defensive response.[234]

PRC-developed models make timely reporting both more important and harder to achieve. The PRC developers whose models are examined in this report are located outside the United States and, once their open weights are downloaded, have no ongoing relationship with the American companies that use them.[235] They are beyond the reach of any U.S. reporting requirement and have no incentive to report problems with their models to the U.S. government. The entities positioned to detect incidents involving PRC-developed models are American ones. When a company's systems are compromised through a PRC-developed model it has deployed, whether through a hijacked AI agent, insecure code the model generated, or behavior the company did not anticipate, the company itself is best placed to identify the model involved. Cloud and inference providers that host these models, and cybersecurity companies that respond to intrusions, see these incidents as well. An effective reporting framework must draw on those entities.

---

[232] *See, e.g.*, OpenAI Incident Disclosure, *supra* note 61; Anthropic Espionage Disclosure, *supra* note 160; Anthropic September 2026 Report, *supra* note 87.
[233] *See, e.g.*, Anthropic Distillation Report, *supra* note 16; OpenAI Memorandum, *supra* note 16; Anthropic Espionage Disclosure, *supra* note 160; OpenAI Incident Disclosure, *supra* note 61; Joint Distillation Advisory, *supra* note 14.
[234] *See* 6 U.S.C. § 681a(a).
[235] *See supra* notes 48–49 and accompanying text.

Congress has already established that framework. CIRCIA places reporting obligations on the entities that experience significant cyber incidents.[236] It directs CISA to require covered entities to report covered cyber incidents within 72 hours and ransom payments within 24 hours, permits supplemental reports as an investigation develops, and allows third parties, such as incident response firms, to submit reports on an entity's behalf.[237] It also provides that reports may not be used to regulate the reporting entity or as evidence against it, exempts them from public disclosure, and preserves legal privileges.[238] As Congress considers how the federal government should learn about AI-related incidents, including those involving PRC-developed models, it should build on CIRCIA rather than create new reporting requirements that would ask the same companies to evaluate the same events against different standards, report them to different agencies, and meet different deadlines. AI-related cyber incidents should be reported once, to CISA, under CIRCIA.

Congress should ensure that CIRCIA captures significant AI-related cyber incidents without stretching it beyond its purpose. As CISA finalizes its rule, it should confirm that substantial cyber incidents involving AI systems are covered.[239] These include the theft or exfiltration of model weights, the compromise of AI training or inference infrastructure, unauthorized access to systems by an AI agent, and intrusions carried out with AI. CISA should also confirm that frontier AI developers and large AI infrastructure providers are covered entities within the Information Technology Sector. Reports should indicate whether an AI system operated by the reporting entity was involved in the incident and, if so, the name of the model's developer, whether its weights were publicly available, and whether the developer was an entity subject to the control of a foreign adversary. Where the reporting entity has reason to believe an attacker used a particular AI system, the report should include that information as well. Over time, these reports would give the government a picture it does not have today of the role PRC-developed models play in incidents affecting the United States. If the final rule leaves any of these questions unresolved, Congress should address them through targeted amendments to CIRCIA.

Reporting tells the government that an incident occurred. Understanding why it occurred requires investigation. The Hugging Face incident, the PRC state-sponsored campaign that used an American AI system to conduct a largely autonomous espionage operation, and the industrial-scale distillation campaigns described in this report each raised questions about root causes that no single company or agency was positioned to answer comprehensively.[240] Congress should establish within CISA an independent, nonregulatory board with the authority to review significant cyber incidents and cyber incidents involving AI,

[236] Cyber Incident Reporting for Critical Infrastructure Act of 2022, Pub. L. No. 117-103, div. Y, 136 Stat. 1038 (codified at 6 U.S.C. §§ 681–681g); *see* 6 U.S.C. § 681b(a).
[237] 6 U.S.C. § 681b(a)(1)–(3), (d).
[238] *Id*. § 681e(a)(5), (b)–(c).
[239] *See* Dawn Halpin, *The CIRCIA Final Rule Is at the White House, with Deadlines Set and Scope Still Open*, Paubox (Oct. 8, 2026), https://www.paubox.com/blog/circia-final-rule-white-house-deadlines-set-scope-open; Cyber Incident Reporting for Critical Infrastructure Act (CIRCIA) Reporting Requirements, 89 Fed. Reg. 23,644 (proposed Apr. 4, 2024) (to be codified at 6 C.F.R. pt. 226).
[240] OpenAI Incident Disclosure, *supra* note 61; Anthropic Espionage Disclosure, *supra* note 160; Joint Distillation Advisory, *supra* note 14.

determine their root causes, and disseminate lessons learned to critical infrastructure owners and operators, state and local governments, industry, and federal agencies. Reviews of AI-related incidents should examine whether the AI system involved had publicly available weights, whether it was developed by an entity subject to the control of a foreign adversary, whether its safeguards had been removed or altered, and how it was distributed and made available in the United States. The board's scope should include coordinated efforts by foreign adversaries to extract the capabilities of American AI systems. The board should have no authority to assign liability, and its membership should be balanced and free of foreign conflicts of interest.

D. Impose Costs on Illicit Extraction

Recommendation 9. Hold illicit distillers accountable and close the gray market. Illicit distillation allows PRC developers to take capabilities that American companies spent billions of dollars to build and offer them back to the market at a fraction of the price. The executive branch should impose sanctions on PRC AI developers that the government determines have engaged in industrial-scale illicit distillation or other malicious cyber-enabled activity against American AI companies.[241] The Department of the Treasury already has authority to sanction foreign persons responsible for cyber-enabled misappropriation of trade secrets for commercial or competitive advantage.[242] The Protecting American Intellectual Property Act separately authorizes sanctions for significant theft of trade secrets belonging to United States persons.[243] Those authorities should be applied wherever the conduct meets their standards. Designations should also reach the transfer stations and fraudulent account brokers that the joint advisory identified as sustaining these campaigns.[244] The executive branch should use the Entity List and information and communications technology supply chain authorities to restrict designated firms' access to American chips and computing resources. Each designation should identify the specific conduct at issue, so that it can withstand legal challenge and make clear to industry what conduct triggers it. The executive branch should also account for American companies that depend on affected model families and provide notice and reasonable transition periods where appropriate, consistent with Recommendation 4.

Congress should examine whether existing trade secret and computer fraud statutes adequately address unauthorized extraction of model capabilities. That review should include whether model outputs and reasoning obtained in violation of a provider's terms of service fall within those statutes. The Department of Commerce, working with DHS and industry, should develop identity verification standards for providers of frontier AI services and cloud computing to help close the gray market of proxy "transfer stations" that enable distillation.[245] DHS,

[241] Bessent, *supra* note 149; Joint Distillation Advisory, *supra* note 14.
[242] *See* Exec. Order No. 13,694, § 1(a)(ii)(D), 80 Fed. Reg. 18,077 (Apr. 2, 2015), *amended by* Exec. Order No. 13,757, 82 Fed. Reg. 1 (Jan. 3, 2017).
[243] Protecting American Intellectual Property Act of 2022, Pub. L. No. 117-336, 136 Stat. 6158 (2023) (codified at 50 U.S.C. § 1709).
[244] Joint Distillation Advisory, *supra* note 14.
[245] *Id.*; Anthropic September 2026 Report, *supra* note 87.

through CISA, should support cross-industry sharing of distillation indicators, as the joint advisory recommends.[246]

Recommendation 10. Examine sales of American training data to PRC AI laboratories. Congress should direct the Department of Commerce, in consultation with the Department of War, DHS, and the Intelligence Community, to assess the scale and national security implications of sales by American data vendors to PRC AI laboratories, particularly those on the Entity List or designated under Section 1260H, and to consider the imposition of disclosure requirements, licensing, or export controls on specific categories of expert training data.[247]

E. Set Durable Rules and Work with Allies

Recommendation 11. Govern both closed and open models and avoid overregulation. Congress should ensure that federal policies, guidance, and programs aimed at securing AI consider the different ways closed and open models are developed and deployed. Congress should also establish a federal baseline for AI transparency and security that provides consistent national rules and does not overburden startups, researchers, or nonprofits.

Recommendation 12. Work with allies and deepen understanding of PRC strategy. The executive branch should expand joint model evaluations with allied AI security institutes, building on the joint United States and United Kingdom assessment of Kimi K3, and should work with allies to shape international standards for model provenance and security.[248] Congress should direct the Intelligence Community, including the DHS Office of Intelligence and Analysis, to provide the relevant committees with an assessment of PRC government involvement in the open-weight strategies of PRC AI developers, including state financing, direction, and access to model weights and user data.

[246] Joint Distillation Advisory, *supra* note 14.
[247] Tong, *supra* note 15; Section 1260H List, *supra* note 96.
[248] Nat'l Inst. of Standards & Tech. & U.K. AI Sec. Inst., *UK AISI / CAISI Preliminary Assessment of Kimi K3's Cyber Capabilities* (July 23, 2026), https://www.nist.gov/news-events/news/2026/07/uk-aisi-caisi-preliminary-assessment-kimi-k3s-cyber-capabilities [hereinafter Kimi K3 Joint Assessment].

CONCLUSION

The question before Congress and the entire federal government is not whether Americans should have access to powerful, affordable AI. They should. The question is whether the AI that powers American businesses, public services, and critical infrastructure will be developed by companies accountable to American law and values, or by companies that must answer to the CCP.

The record assembled by the Committees shows that PRC-developed models have become embedded in American products faster than companies, customers, or the government have been able to understand the implications. That adoption was driven by real advantages in price and flexibility. Those advantages were built in part on capabilities taken from American developers and sustained by state support and a willingness to lose money for market share. The PRC government has made clear that it views open-weight AI as a source of international influence, and it has already considered restricting access to the models on which others now depend.[249]

The United States still has time to change this trajectory. American developers retain the lead at the frontier, and a new generation of American open-weight models is emerging.[250] The Committees' recommendations are designed to accelerate that progress, give companies and the government the information they need, keep PRC models out of the systems where the stakes are highest, impose costs on theft, and prepare for disruptions before they occur. The nation should not wait for another rip-and-replace program to learn the lesson of Huawei, ZTE, and DJI a third time.

The Committees will continue their oversight of these issues.

[249] *See* State Council AI+ Opinions, *supra* note 17, § III(11); Potkin, *supra* note 19.
[250] Emberson, *supra* note 66; Lang et al., *supra* note 66; Stanford AI Index, *supra* note 66; Bellan, *Reflection Debuts Beam*, *supra* note 76.

ANNEX

Annex A. Zhipu and the PLA Cyberspace Force

Information made available to the Committees describes a relationship between elements of the PRC's military intelligence apparatus and Zhipu.[251] According to that information, individuals associated with the 1st Bureau of the PLA Cyberspace Force have been present on multiple occasions at Zhipu's headquarters in Beijing, most recently in August 2026.[252] The 1st Bureau, formerly part of the PLA Strategic Support Force's Network Systems Department, is understood to conduct signals intelligence, cyber espionage and attack, electronic warfare, and psychological operations, and is believed to be co-located with the PLA Cyberspace Force headquarters.[253] The Committees assess with moderate confidence that this pattern indicates a channel between the PLA's cyber intelligence apparatus and a leading commercial PRC AI developer.[254]

Annex B. State Support for Chinese AI Developers

In addition to the direct grants, state investment, and government ownership described above, China's AI labs benefit from a broader set of provincial, municipal, and district policies that reduce the cost of compute, deployment, taxes, and talent while also subsidizing customer adoption.

Compute subsidies target one of the largest recurring costs of model development. At the provincial level, Zhejiang covers 10 to 30 percent of compute costs for Alibaba and a Zhipu affiliate,[255] while Guizhou reimbursed $160,000 for Moonshot's Kimi model.[256] At the municipal level, Beijing discounts compute for Zhipu,[257] Moonshot,[258] and MiniMax,[259] while Hangzhou allocates $37.3 million annually in compute vouchers.[260] At the district level, Shanghai's Xuhui District,

---

[251] Committee Records, *supra* note 37.
[252] *Id.*
[253] *Id.*
[254] *Id.*
[255] Zhejiang Provincial Development & Reform Commission, *Notice on Organizing Applications for Redemption of 2026 Artificial Intelligence Vouchers (Computing Vouchers)* [关于组织开展2026年人工智能券（算力券）兑付申报的通知] (May 1, 2026), https://perma.cc/Y7H3-9UTQ; Zhejiang Provincial Development & Reform Commission, *Notice on Publicizing Projects Proposed for the 2026 Artificial Intelligence Voucher (Computing Voucher) Program* [关于公示拟申请2026年人工智能券（算力券）项目情况的通知] (May 11, 2026), https://perma.cc/JF5Q-B9AT.
[256] *Promoting Computing-Power Transactions to Empower Industrial Development* [推动算力交易，赋能产业发展], People's Daily Overseas Edition [人民日报海外版] (Sept. 27, 2024), https://perma.cc/C83P-EYPM.
[257] *Beijing 2025 National Artificial Intelligence Voucher (Compute Voucher) Project Eligibility Review Results* [北京市2025年国家人工智能券（算力券）项目资格审核结果], Beijing Mun. Bureau of Econ. & Info. Tech. [北京市经济和信息化局] (May 24, 2025), https://perma.cc/5AQU-DMUT; *Notice on Organizing Applications for 2024 Beijing High-Tech Industry Development Funds (Second Batch)* [北京市经济和信息化局 北京市财政局关于组织开展2024年北京市高精尖产业发展资金申请工作（第二批）的通知], Beijing Mun. Bureau of Econ. & Info. Tech. & Beijing Mun. Fin. Bureau [北京市经济和信息化局、北京市财政局] (Apr. 18, 2024), https://perma.cc/XR8E-RRH6.
[258] *Yin Yong Conducts Research on Medical Technology Commercialization and AI Industry Development* [殷勇调研医疗科技成果转化和人工智能产业发展时强调 大力培育发展新质生产力 打造世界级科技创新策源地], Beijing Mun. People's Gov't [北京市人民政府] (Jan. 9, 2026), https://perma.cc/VFS6-G5EM; *Beijing Releases New Policies for AI Agent Development: Haidian's AI Origin Community Explores the Token Economy* [北京发布智能体发展新政：科研智能体跑通全流程 海淀原点社区探路Token经济], Beijing Mun. Dev. & Reform Comm'n [北京市发展和改革委员会] (July 24, 2026), https://perma.cc/B6SB-FFQ2.
[259] *Beijing 2025 Compute Voucher Review Results*, *supra* note 257; *2024 Beijing High-Tech Industry Development Funds Notice*, *supra* note 257.
[260] *Notice of the General Office of the Hangzhou Municipal People's Government on Issuing Several Measures to Support High-Quality Development of the Entire Artificial Intelligence Industry Chain* [杭州市人民政府办公厅关于印发支持人工智能全产业链高质量发展若干措施的通知], Hangzhou Municipal People's Government [杭州市人民政府] (July 18, 2024), https://perma.cc/K54W-HBEN; *Notice on Issuing the 2025 Implementation Plan for Accelerating the Construction of an Artificial Intelligence Innovation Highland* [杭州市人民政府关于印发杭州市加快建设人工智能创新高地实施方案（2025年版）的通知], Hangzhou Mun. People's Gov't [杭州市人民政府] (June 13, 2025), https://perma.cc/5H64-FRPX [hereinafter Hangzhou 2025 AI Innovation Highland Plan].

where MiniMax is based, offers compute subsidies of up to $3 million annually,[261] Beijing's Haidian District targets Zhipu[262] and Moonshot[263] with vouchers, and Hangzhou's Gongshu District provides up to $746,000 for DeepSeek.[264]

To stimulate adoption, governments also subsidize customers and direct demand toward domestic providers. At the municipal level, Beijing channels subsidized purchases to Alibaba Cloud[265] and ByteDance's Volcano Engine,[266] including about $121,000 and $854,000, respectively, in reimbursements for large model applications.[267] At the district level, Beijing's Haidian District also provided Zhipu with a $7.5 million customer voucher pool.[268] Elsewhere, Hangzhou funded $750,000 in "AI+" projects,[269] while Shanghai state platforms connected MiniMax with industry partners for testing and deployment.[270]

China's AI labs also benefit from preferential tax treatment that lowers scaling costs. Qualifying subsidiaries of Moonshot,[271] ByteDance,[272] MiniMax,[273] Zhipu,[274] and Alibaba[275] receive a 15 percent high tech enterprise tax rate instead of the standard 25 percent.[276] Alibaba receives additional multi-year exemptions and, for

[261] MiniMax Releases Successive New Large-Model Products; MiniMax-M1 Series Ranks Second Globally in Long-Text Capability [稀宇科技连续发布大模型新产品 MiniMax-M1系列长文本能力跻身全球第二], Shanghai Mun. People's Gov't [上海市人民政府] (June 22, 2025), https://www.shanghai.gov.cn/nw4411/20250622/75b969a99884465ea0ac64759d587969.html.

[262] *Beijing News Series Report: Large-Model Innovation Leads the Nation as Beijing's Artificial Intelligence Industry Builds Momentum* [《北京新闻》系列报道之《大模型创新领跑全国 北京人工智能产业蓄力腾飞》], Beijing Municipal Science & Technology Commission & Zhongguancun Science Park Administrative Committee [北京市科学技术委员会、中关村科技园区管理委员会] (July 28, 2026), https://perma.cc/RX9N-TS4N; *Haidian District Issues 2025 Artificial Intelligence Enterprise Computing Subsidy Application Guidelines, Offering Subsidies of up to RMB 20 Million* [最高补贴2000万 海淀区发布AI算力补贴申报指南], Open Beijing [开放北京] (Sept. 29, 2025), https://perma.cc/7VMQ-NUB5.

[263] *Kimi Set to Complete $2 Billion Financing!* [Kimi将完成20亿美元融资！], National Science and Technology Innovation Center [国家科技创新中心] (May 8, 2026), https://perma.cc/58Q6-JLNH; *China AI Start-Up Moonshot Snags Funds at US$18 Billion Valuation*, South China Morning Post (Mar. 15, 2026).

[264] *Response to Proposal No. 220 of the Fourth Session of the Twelfth Hangzhou Municipal Committee of the Chinese People's Political Consultative Conference* [对杭州市十二届政协四次会议220号提案的答复], Gongshu District People's Government [杭州市拱墅区人民政府] (Oct. 24, 2025), https://perma.cc/5ZVS-KRNX.

[265] *Beijing 2025 Compute Voucher Review Results*, *supra* note 257; *2024 Beijing High-Tech Industry Development Funds Notice*, *supra* note 257.

[266] *Beijing 2025 Compute Voucher Review Results*, *supra* note 257; *2024 Beijing High-Tech Industry Development Funds Notice*, *supra* note 257.

[267] *Beijing Municipal Bureau of Economy and Information Technology Notice Publicizing Projects Proposed for Support Under the Fourth Batch of 2025 Small and Medium-Sized Enterprise Development Funds* [北京市经济和信息化局关于对2025年第四批中小企业发展资金拟支持项目进行公示的通知], Beijing Municipal Bureau of Economy and Information Technology [北京市经济和信息化局] (Dec. 10, 2025), https://perma.cc/M9AP-PTQT; *Small and Medium-Sized Enterprise Service Voucher Subsidy Projects Proposed for Payment* [中小企业服务券补贴拟兑付项目], Beijing Municipal Bureau of Economy and Information Technology [北京市经济和信息化局] (Dec. 10, 2025), https://perma.cc/B53V-Z6KF; *Notice on Issuing the 2025 Beijing High-Tech Industry Development Project Funds and Small and Medium-Sized Enterprise Development Funds Implementation Guidelines (Second Batch)* [北京市经济和信息化局 北京市财政局关于印发2025年北京市高精尖产业发展项目资金和支持中小企业发展资金实施指南（第二批）的通知], Beijing Municipal Bureau of Economy and Information Technology & Beijing Municipal Finance Bureau [北京市经济和信息化局、北京市财政局] (May 7, 2025), https://perma.cc/TP7T-T38C.

[268] *Zhongguancun Science City Distributes RMB 100 Million in Large-Model Vouchers to Support Enterprise Transformation* [中关村科学城发放亿元大模型券助企转型], Open Beijing [开放北京] (July 13, 2026), https://perma.cc/HNQ4-3CJH; *Yin Yong Conducts Research*, *supra* note 258; *Beijing Releases New Policies for AI Agent Development*, *supra* note 258.

[269] *Hangzhou and Alibaba Group Hold Strategic Cooperation Advancement Meeting* [杭州市与阿里巴巴集团战略合作推进会举行], The Paper [澎湃新闻] (Apr. 2, 2025), https://perma.cc/Q9DS-L9W8; *Hangzhou 2025 AI Innovation Highland Plan*, *supra* note 260; *Launching "DeepSeek+": Gongshu to Build an Open-Source Ecosystem Pilot Zone* [启航"DeepSeek+" 拱墅要建开源生态先导区], Hangzhou Publicity Network [杭州宣传网] (May 22, 2026), https://www.hzxcw.gov.cn/content_47863.html.

[270] *MiniMax Lists on the Hong Kong Stock Exchange*, *supra* note 18.

[271] *Announcement on the Filing of the First Batch of High and New Technology Enterprises Recognized by Beijing's Certification Authorities in 2024* [对北京市认定机构2024年认定报备的第一批高新技术企业进行备案的公告], High and New Technology Enterprise Certification Administration [高新技术企业认定管理工作网] (Oct. 29, 2024), https://perma.cc/XME8-3FPF.

[272] *List of High and New Technology Enterprises in Beijing* [北京市高新技术企业名单], State Administration of Foreign Exchange Beijing Branch [国家外汇管理局北京外汇管理部] (July 19, 2022), https://perma.cc/7XTF-VXQC.

[273] MiniMax Annual Report, *supra* note 113, at 123–24.

[274] Zhipu Prospectus, *supra* note 111.

[275] Alibaba FY2024 Annual Report, *supra* note 109.

[276] *Enterprise Income Tax Law of the People's Republic of China* [中华人民共和国企业所得税法] arts. 4, 28 (as amended Dec. 29, 2018), https://perma.cc/H623-DBE9.

qualifying entities, a 10 percent Key Software Enterprise rate,[277] while Zhipu receives a 100 percent R&D super-deduction.[278]

Additionally, local governments directly subsidize recruitment, housing, and retention to lower talent costs for key firms. In Beijing, Haidian District provides AI workers up to $225 monthly in rent subsidies alongside targeted hiring programs for Moonshot and ByteDance.[279] In Hangzhou, the city offers up to $750,000 for top innovation teams and collaborates with Alibaba on recruitment,[280] while Gongshu District deployed a dedicated task force for DeepSeek.[281] In Shanghai, Xuhui District enrolled MiniMax in a district training alliance connecting the firm directly to university pipelines.[282]

CLASSIFIED ANNEX

A classified annex to this report has been prepared by the Committees and is available to Members of the Committees and to Committee staff who hold the appropriate security clearances. Requests for access from outside the Committees must be submitted in writing to both Chairmen and will be considered on a need-to-know basis.

[277] Alibaba FY2024 Annual Report, *supra* note 109.
[278] Zhipu Prospectus, *supra* note 111.
[279] *Haidian District Strengthens Fiscal Support for AI Industry Development* [海淀区强化财政支持人工智能产业发展], Haidian District People's Government [北京市海淀区人民政府] (Nov. 6, 2025), https://zyk.bjhd.gov.cn/jbdt/auto4495_51790/auto4495_52622/auto4495/auto4495/202511/t20251106_4792267.shtml; *Notice on Applications for 2026 Haidian District Youth Talent Housing Subsidies* [北京市海淀区人才工作局关于开展2026年海淀区青年人才安居补贴申请的通知], Haidian District Talent Work Bureau [北京市海淀区人才工作局] (Feb. 28, 2026), https://zyk.bjhd.gov.cn/zwdt/zcwj/202602/t20260228_4806584.shtml; *Haidian District Human Resources and Social Security Bureau's "Zhongguancun Hundred Universities Alliance" Spring Recruitment Provides 32,000 Quality Jobs for College Graduates* [海淀区人力社保局"中关村百校联盟"春季招聘为高校毕业生提供3.2万个优质岗位], Haidian District Human Resources and Social Security Bureau [海淀区人力资源和社会保障局] (May 8, 2026), https://web.archive.org/web/20261005125515/https://zyk.bjhd.gov.cn/jbdt/auto4496_51791/auto4496_52699/auto4496/auto4496/202605/t20260506_4814045_hd.shtml; *Haidian District Human Resources and Social Security Bureau Holds "Global AI Elite Returnee Talent Matchmaking Conference"* [海淀区人力社保局举办"全球AI菁英归国人才对接会"], Haidian District Human Resources and Social Security Bureau [海淀区人力资源和社会保障局] (July 24, 2026), https://zyk.bjhd.gov.cn/jbdt/auto4496_51791/auto4496_52699/auto4496/auto4496/202607/t20260722_4822819_hd.shtml.
[280] *Hangzhou and Alibaba Group Hold Strategic Cooperation Advancement Meeting*, *supra* note 269; Hangzhou Municipal People's Government, *Notice on Issuing Several Policies for Building the Hangzhou National New Generation Artificial Intelligence Innovation and Development Pilot Zone* [关于印发杭州市建设国家新一代人工智能创新发展试验区若干政策的通知] (Dec. 12, 2019), https://perma.cc/58T9-LBUU.
[281] *Hangzhou's "Six Little Dragons": How Were They Forged?* [杭州："六小龙"如何炼成], China Economic Weekly [中国经济周刊] (Feb. 28, 2025), https://perma.cc/BW5M-Z45K.
[282] *Xuhui District Establishes AI + Skills Training Alliance* [徐汇区成立"AI+技能培训联盟"], Shanghai Municipal Human Resources and Social Security Bureau [上海市人力资源和社会保障局] (Aug. 6, 2026), https://perma.cc/EW3C-Z6LH.