Leaky Rents at the Frontier: Political Embeddedness and Frontier Performance in China's Civilian Generative AI Sector

Cameron Keith

Dartmouth College

cameron.s.keith.26@dartmouth.edu

March 12, 2026

This paper examines how China's governance system shapes competition among AI labs. Firms with stronger political ties before breakthrough often receive better access to state-mediated inputs, yet those ties do not consistently translate into stronger efficiency-adjusted frontier performance. Any regulatory edge also proves limited: advantages fade, while outsiders that succeed are often brought into the system afterward. Drawing on a structured comparison of Chinese AI labs by level of pre-breakthrough embeddedness, the paper finds a recurring pattern. Highly embedded firms benefit from access, but less embedded firms have frequently delivered stronger efficiency-adjusted frontier advances. The findings suggest that the system exerts its strongest political influence not by shutting outsiders out in advance, but by absorbing them once they have succeeded.

Word count: 2773

1. Introduction

In late 2024, DeepSeek, the AI subsidiary of the Chinese hedge fund High-Flyer, achieved what China's most politically connected labs had failed to do for years. DeepSeek released V3, an open-source 671-billion-parameter large language model that matched or exceeded frontier Western models on standard benchmarks. According to DeepSeek, the total training time was 2.788 million H800 GPU-hours. However, Meta's Llama 3.1 with 405 billion parameters required 30.84 million H100 GPU training hours (DeepSeek-AI, 2024; Meta AI, 2024). DeepSeek reported the final run cost at $5.6 million, although total costs were almost certainly much higher; SemiAnalysis estimates roughly $1.6 billion in server capital expenditure and approximately $2.5 billion overall (Patel et al., 2025). Despite the differences in training costs, the broader fact remains: a firm with no national-team designation, no state funding, and no venture capital reached state-of-the-art performance ahead of labs that had enjoyed years of preferential state access. DeepSeek then open-sourced the model weights, which helped trigger price cuts across the industry.

This paper asks whether China's AI governance system allocates frontier opportunities efficiently or produces rent-seeking distortions. My hypothesis is that greater political embeddedness is associated with priority access to state-sponsored inputs, but not with stronger efficiency-adjusted model performance in the long run. Short-run regulatory advantages should be observable. However, successful outsiders should become more politically embedded after they achieve breakthroughs.

AI development is capital intensive, and key inputs like compute, regulatory approval, and procurement contracts are politically allocated in China. Those conditions make the sector a useful place to test whether political ties improve resource allocation or simply give some firms an early edge. I call those temporary advantages *leaky rents*: state-conferred benefits that matter initially but erode under competitive pressure (Krueger, 1974).

The analysis focuses on publicly benchmarked civilian generative AI in China from 2023 to 2026 and compares leading labs along a gradient of pre-breakthrough political embeddedness.

Section 2 develops the analytical framework. Section 3 defines the scope, case selection, and measures. Section 4 describes the institutional channels. Section 5 evaluates the evidence. Section 6 addresses counterarguments. Section 7 discusses future research.

2. Analytical Framework

2.1. Industrial Policy Expectation

Industrial policy scholars argue that strategic state coordination can improve outcomes in sectors characterized by large fixed costs, externalities, and coordination failures. Juhasz, Lane, and Rodrik (2024) survey the new industrial policy literature and show that there is a serious economic case for intervention under some conditions, while also emphasizing the importance of governance, measurement, and context. Cherif and Hasanov (2019) distill three core principles of successful industrial policy: targeting sophisticated industries, maintaining export discipline, and enforcing competition with accountability. Mazzucato (2013) similarly argues that the state can play an entrepreneurial role in innovation by bearing risks that private actors are unwilling to take.

These arguments appear directly in debates over Chinese AI today. Zhang (2025) contends that China's AI regulation is not purely restrictive, but can also serve enabling and coordinating functions, including by sending strong pro-growth signals to the industry. Chan et al. (2025) likewise present Chinese AI development as the product of interaction between state support and private-sector innovation, with policy most effective when it addresses real bottlenecks in the ecosystem. Scanlon (2025) argues that DeepSeek is best understood not as an isolated outlier but as a product of China's broader state-backed AI ecosystem. I return to this claim in Section 6. If this industrial-policy view is correct, state-sponsored firms should convert that support into stronger frontier performance.

2.2. Rent-Seeking Expectation

Rent-seeking theories predict weaker performance effects from political privilege. Krueger (1974) argues that state-created scarcity can redirect effort away from productive innovation and toward competition for rents. Baumol (1990) and Murphy, Shleifer, and Vishny (1991) similarly show that institutional environments can channel entrepreneurial talent toward politically rewarded but less productive activity, with negative consequences for aggregate growth. Huang and Zhu (2025) provide China-specific evidence consistent with this concern: the Ten Industry Revitalization Plan reduced innovation efficiency among Chinese listed firms, in part through the misallocation of innovation resources. If this view is correct, privileged access should be observable, but it should not reliably predict stronger efficiency-adjusted frontier performance.

2.3. Analytical Synthesis

These two perspectives are not mutually exclusive. A policy regime can be designed to solve coordination failures and still generate rent-seeking effects when political criteria shape the allocation of scarce inputs. The question, then, is not whether Chinese AI policy is best described as purely developmental or purely distortive. Rather, the empirical issue is whether political embeddedness predicts access more strongly than it predicts performance. That is a claim that I can test even if the underlying intent of policy remains ambiguous.

2.4. Observable Implications

If the argument is correct, four patterns should appear in the evidence. First, more politically embedded labs should receive greater access to state-mediated inputs such as compute, approvals, procurement, and political recognition. Second, those access advantages should not reliably translate into stronger efficiency-adjusted frontier performance. Third, early regulatory or allocation advantages should operate as leaky rents: they may produce real short-run benefits, but those benefits should erode once more technically capable rivals enter. Fourth, firms with initially low embeddedness that later achieve technical breakthroughs should become more politically entangled

after success rather than before it. The evidence section evaluates each implication in turn.

3. Scope, Case Selection, and Measures

The units of analysis are AI labs or model teams, not parent firms. That distinction matters because corporate political connections and lab-level technical autonomy do not always align. Alibaba is an original 2017 national team member, yet the Qwen team operates with meaningful commercial autonomy through Cloud/DAMO. I classify on lab-level operational embeddedness, not parent-company designation alone.

I define *political embeddedness* as the degree to which a lab's pre-breakthrough operations depend on state capital or state-guided finance, national-team designation, government procurement, state-mediated compute access, or direct political steering. I base the classification on pre-breakthrough conditions, not on ties that formed after success. The resulting gradient, summarized in Table 1, has three tiers.[1]

Tier	Labs	Basis
Highly embedded	Baidu, SenseTime, iFlytek	Government contracts, national-team designation, political mandates
Semi-embedded	Alibaba/Qwen, ByteDance/Doubao, MiniMax, Moonshot/Kimi, Z.ai/Zhipu	State-connected capital, but labs retain technical autonomy
Least embedded	DeepSeek	Self-funded, no state capital, no national team lineage

Table 1: Pre-breakthrough political embeddedness gradient.

I measure efficiency-adjusted frontier performance through benchmark competitiveness, speed of frontier catch-up, architectural contribution, and serious developer adoption. These indicators test whether access advantages translate into frontier results. I also track the durability of early regulatory advantage and whether successful outsiders become more politically entangled after breakthrough.

[1] Tencent, also an original 2017 national team member (medical imaging), is excluded because its generative AI models have not been prominent in public frontier benchmarks during the study period.

4. Institutional Background: Three Channels of State Influence

China's AI governance system operates through three overlapping channels. The first is compute mediation. Access to compute, the most strategically important input, is shaped by state-owned infrastructure, government compute vouchers distributed across at least seventeen cities (typically worth $140,000 to $280,000 per company), and state-funded data centers that formally bar foreign AI chips (Chan et al., 2025; James, 2025; Scanlon, 2025).

The second channel is the Cyberspace Administration of China (CAC) filing and registration process, which the Interim Measures for Generative AI, issued on July 10, 2023, and effective August 15, 2023, formalized. Although the process is not formally equivalent to a license, it operates as a gatekeeping mechanism that imposes timing costs, compliance burdens, and uncertainty. Those costs favor firms with stronger regulatory capacity. In practice, regulators defer acceptance through multiple review rounds with "poorly formalized and constantly changing" procedures (Welch, 2024). By March 2025, 346 services had filed; by August 31, 2025, 538 services had completed filing and 263 applications or functions had completed registration (Zhang, 2025; Zhao, 2025; Schaefer, 2025).

The third channel is national-team designation. In November 2017, the Ministry of Science and Technology designated Baidu (autonomous driving), Alibaba (smart city), Tencent (medical imaging), and iFlytek (speech recognition) as National AI Open Innovation Platforms. SenseTime joined in 2018, and the list expanded to roughly fifteen members by late 2019. Benefits include government contracts, pilot access, public-sector data opportunities, and influence over standards and ecosystem development (Larsen, 2019; Ding, 2019; Cave, Ryan, and Xu, 2019). Together, these three channels create a sector in which political access to scarce inputs shapes frontier competition alongside engineering quality and capital.

Figure 1 summarizes the paper's core mechanism, while Figure 2 situates the three channels in temporal sequence. Table 2 compares the institutional logic of each channel.

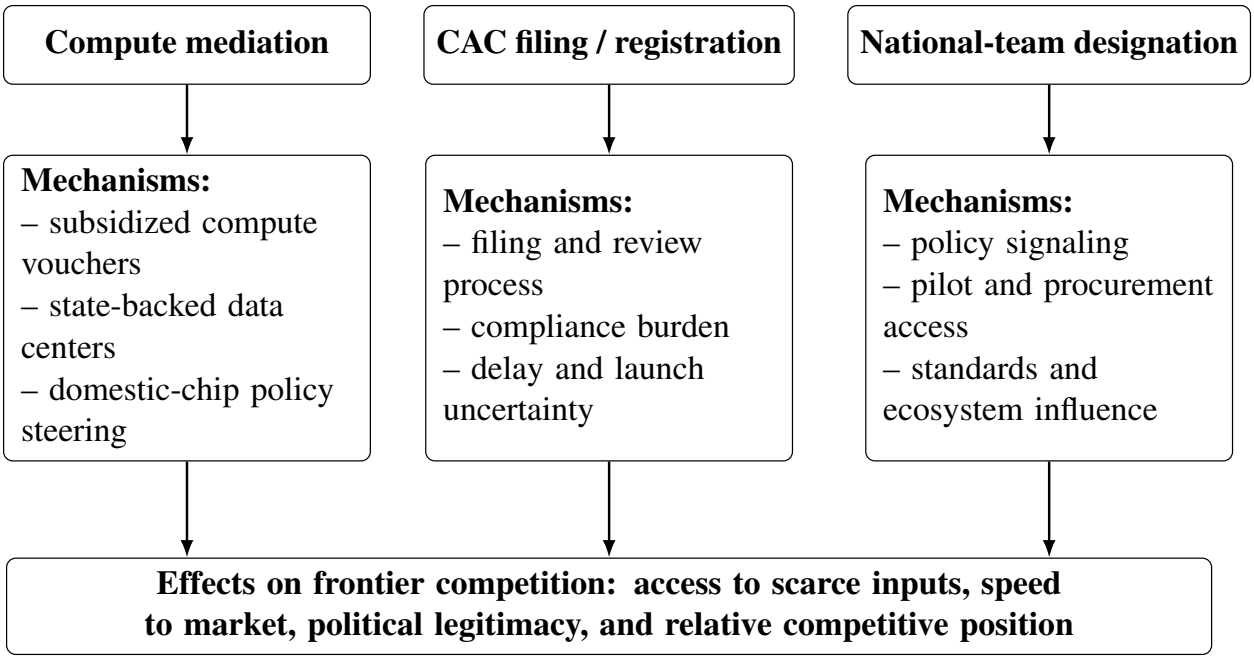

Figure 1: Three channels of state influence in China's AI sector. State influence operates through compute mediation, CAC filing and registration, and national-team designation. These channels shape firms' access to scarce inputs, regulatory timing, and political legitimacy, thereby affecting frontier competition alongside engineering quality and capital.

5. Evidence

The evidence is organized around the four observable implications developed in Section 2.

5.1. Access Tracks Embeddedness

More politically embedded labs enjoy greater access to state-mediated inputs before any assessment of technical breakthrough. The RAND "Full Stack" report warns that if scarce AI chips are not allocated efficiently, resources risk diversion from more productive private users (Chan et al., 2025). Trivium China finds that state-owned enterprises registered approximately 22% of generative AI tools but "are not actually building these projects by themselves," instead partnering with private firms (Schaefer, 2025). For example, PetroChina's Kunlun model required assistance from China Mobile, Huawei, and iFlytek, while State Grid relied on DeepSeek distillation (Schaefer, 2025). Econometric evidence confirms that political connections drive procurement, with rent-

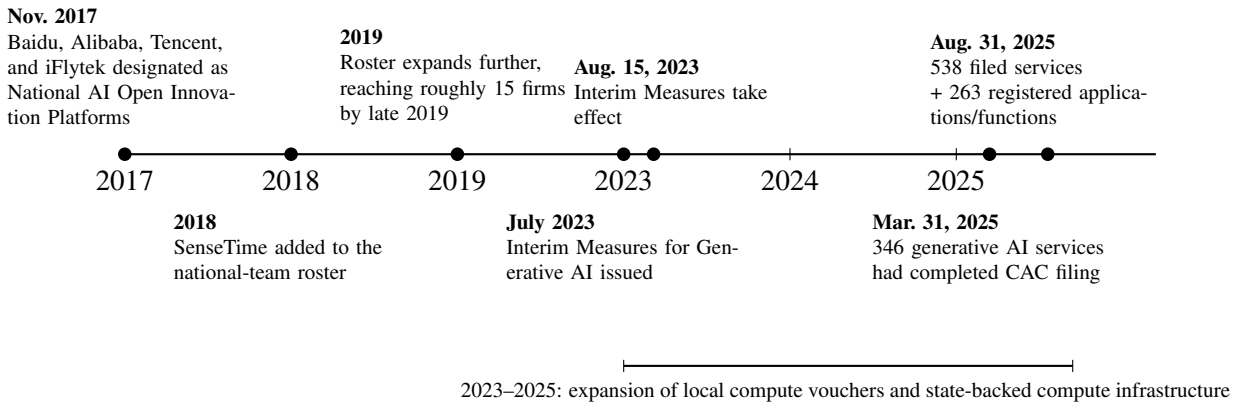

Figure 2: Institutional buildup of state influence in Chinese AI, 2017–2025. National-team designation began in 2017, the CAC generative AI regime was formalized in 2023, and by 2025 the filing and registration system had scaled substantially. Over the same period, state-backed compute support expanded through local subsidies and infrastructure policy.

seeking as a mediating mechanism (Wang et al., 2025). Moreover, following DeepSeek's success, venture capital for AI startups declined approximately 50% as state-affiliated actors assumed a leading role, and provincial authorities pursued AI projects to "demonstrate administrative competence" (Singer and Sheehan, 2025).

Semi-embedded labs occupy an intermediate position. Alibaba/Qwen retains the autonomy noted in Section 3. MiniMax retains operational independence despite state-connected capital. ByteDance/Doubao competes at the frontier on commercial terms despite deep regulatory entanglement. DeepSeek, the least embedded lab in the sample (Table 1), stands at the opposite end.

Those patterns are consistent with broader findings in political economy. Zhong and Zheng (2025) find that Chinese firms that lost political connections subsequently filed more patent applications, which suggests that connections may inhibit innovation rather than enable it. Huang and Zhu (2025) provide corroborating China-specific evidence (discussed in Section 2), and Akcigit, Baslandze, and Lotti (2023) provide a general model of how political connections reshape innovation dynamics. The first observable implication holds: access tracks embeddedness. The next subsection asks whether that access translates into stronger frontier outcomes.

Channel	Lead state actor	Mechanism	Resource affected	Firms most advantaged
Compute mediation	Local governments, state-linked infrastructure actors	Compute vouchers, subsidized data centers, chip-access steering	Compute and infrastructure access	Firms with capital, scaling needs, and political access
CAC filing / registration	Cyberspace Administration of China	Filing, review, iterative compliance, launch delays	Market access and deployment timing	Firms with stronger legal, policy, and compliance capacity
National-team designation	Ministry of Science and Technology and related state bodies	Official platform status, policy signaling, pilot projects, ecosystem coordination	Contracts, public-sector access, legitimacy, standards influence	Politically trusted incumbents and strategic sector leaders

Table 2: Comparison of the three channels of state influence. Each channel operates through a distinct institutional mechanism, but all three shape firms' access to scarce resources or strategic advantages relevant to frontier AI competition.

5.2. Embeddedness Does Not Predict Frontier Performance

Greater pre-breakthrough embeddedness does not reliably produce stronger efficiency-adjusted frontier results. When the first SuperCLUE benchmark ran in May 2023, iFlytek's SparkDesk, from the highly embedded tier, scored highest among Chinese models at 53.58, compared to GPT-4's 76.67. However, the frontier shifted within eighteen months. By December 2024, DeepSeek-V3 (least embedded), Doubao, GLM-4-Plus from Z.ai, and Qwen (all semi-embedded) led the field. The highly embedded original national team members, despite years of preferential access, fell behind labs with greater technical autonomy.

DeepSeek serves as the critical case, though I do not treat it as the sole basis of the argument. V3 consumed 2.788 million H800 GPU-hours versus Llama 3.1's 30.84 million H100 GPU-hours (DeepSeek-AI, 2024; Meta AI, 2024). Its key architectural innovations (Multi-head Latent Attention, FP8 native training, auxiliary-loss-free mixture-of-experts load balancing) originated entirely outside the national team framework. Developer adoption reinforces the picture: Qwen accumu-

lated over 600 million Hugging Face downloads by late 2024, and Chinese open-source models rose from near zero to approximately 30% of OpenRouter usage between late 2024 and mid-2025 (Chen, 2026; Meinhardt et al., 2025). Epoch AI estimates the average China-US frontier gap at seven months, with less embedded labs driving most of the convergence (Emberson, 2026).

I must acknowledge counterevidence. SenseTime topped SuperCLUE in early 2025. iFlytek led the earliest benchmarks. Baidu claimed GPT-4 parity with ERNIE 4.0 in October 2023, though that claim is contested.

5.3. Regulatory Advantages Are Real but Temporary

The CAC filing system generates measurable but temporary advantages that erode under quality competition. The rent was real. Batch 1 approval on August 31, 2023, covered Baidu, SenseTime, ByteDance, Baichuan, and Z.ai (then Zhipu AI) (Zhao, 2025; Schaefer, 2025). Ernie Bot surged to number one on Apple's App Store that same day (Soo, 2023). Baidu later reported that the service had surpassed 200 million users by April 2024 and 300 million by June 2024 (Ye, 2024; Mo and Hall, 2024). Regulatory timing translated directly into market access.

However, the rent leaked. By January 2025, Ernie Bot had about 13 million monthly active users, compared with 78.6 million for ByteDance's Doubao and 33.7 million for DeepSeek, according to data compiled by AI tracker Aicpb.com (Reuters, 2025a). Early regulatory timing created a real first-mover advantage, but it did not produce durable insulation once stronger rivals entered.

The mechanism follows Krueger (1974). Filing involved the opaque, resource-intensive procedures described in Section 4, sufficiently burdensome to favor well-resourced firms. Only approximately 45% of 305 known models had successfully registered by mid-2024. Regulatory requirements consumed real resources and advantaged incumbents, yet later entrants with technically superior products still captured the market. I use commercial adoption metrics here not as measures of frontier quality but as indicators of whether early regulatory advantage translated into durable insulation. It did not.

5.4. Successful Outsiders Are Absorbed after Breakthrough

The most suggestive finding concerns what happens to successful outsiders after their breakthroughs. The temporal sequence matters. Before its breakthrough, DeepSeek exhibited low embeddedness by every measure in the gradient (Table 1). The breakthrough came with V3 in December 2024, followed by the release of the reasoning model R1 in January 2025 (DeepSeek-AI, 2024, 2025b,a). R1 reinforced the firm's global significance: in the ensuing selloff, Nvidia lost roughly $589 billion in market value in a single day as investors reassessed the compute economics of frontier AI (Reuters, 2025b).

After breakthrough, political entanglement escalated along multiple dimensions. DeepSeek had received a national high-tech enterprise (HTNE) tax designation from Zhejiang province in December 2023, but thousands of Chinese technology firms hold HTNE status, which makes it categorically different from national team membership. More significant developments followed the breakthrough: Liang Wenfeng attended Premier Li Qiang's symposium in January 2025, where political recognition clearly followed technical achievement (Singer and Sheehan, 2025). Exiger documented DeepSeek researchers on 396 PLA-funded AI research projects with affiliations to 42 government talent recruitment programs (Exiger, 2025). The Jamestown Foundation identified dozens of PLA procurement documents that called for DeepSeek tools (Cheung and Lau, 2025). Chinese authorities also reportedly pressured DeepSeek to use Huawei Ascend chips for R2 training (Olcott and Wu, 2025). Scanlon (2025) reads the endpoint and projects it backward; I read the sequence forward. Independent innovation came first. State absorption came second.

The evidence is consistent with post-hoc absorption, but I cannot yet establish that this mechanism has altered DeepSeek's innovation incentives or technical trajectory. Whether the independence that enabled the breakthrough will survive growing state entanglement remains the most consequential open question. I flag it as a direction for future research.

6. Counterarguments and Limitations

6.1. The Pro-Growth Interpretation

Zhang (2025) argues that China's AI regulations are "markedly business-friendly" and function as "an enabler by sending a strong pro-growth signal." Li and Ban (2025) argue that government guidance funds represent a "good enough" institutional solution. I do not dispute that pro-growth intent exists. However, pro-growth intent can coexist with rent-seeking distributional outcomes. The allocation patterns I document appear more consistent with privileged access than with efficient selection of the strongest frontier performers. Zhang's own evidence partially supports this reading: she acknowledges that the regulations "offer little protective value to the Chinese public" (Zhang, 2025). If the system is neither protecting consumers nor blocking firms, the question becomes what distributional consequences it produces. The embeddedness gradient provides an answer.

6.2. DeepSeek as a Product of the Chinese Ecosystem

Scanlon (2025) argues that DeepSeek emerged from "the vast political and financial ecosystem China has built to accelerate AI innovation." This is partly true. I do not claim that DeepSeek arose outside China's broader institutional environment. The narrower claim is that DeepSeek's breakthrough did not depend on the strongest forms of pre-breakthrough political embeddedness observed elsewhere, and that the more substantial forms of political incorporation followed success rather than caused it. The HTNE designation is available to thousands of firms. The 2021 crackdown on speculative investment targeted finance, not AI. Regarding distillation, OpenAI and Anthropic have alleged that DeepSeek used distillation from their models (CNBC, 2026). However, the core architectural innovations documented in the V3 technical report are original contributions (DeepSeek-AI, 2024). The efficiency argument in this paper rests on GPU-hours and architectural novelty, not on training data provenance.

6.3. Visibility, Measurement, and Scope Limitations

Three further limitations constrain the analysis. First, if the state's most capable AI work remains classified, public benchmarks understate state-linked capability. I limit my claims to publicly benchmarked civilian generative AI. However, if state-linked labs deploy superior models only internally, that pattern itself suggests that public-market competition is not the system's sole objective. Second, compute figures are imperfect, benchmarks are incomplete, and some evidence comes from media rather than peer-reviewed studies. Those limitations justify bounded claims, not abandonment of the pattern. Third, results from surveillance AI should not transfer to frontier generative AI. Beraja et al. (2023) find that state procurement increased innovation in surveillance AI by approximately 49%, but surveillance contracts provide defined specifications and stable demand. Frontier generative AI requires open-ended exploration that political mandates may constrain.

7. Conclusion

China's AI governance system is best understood as a mixed system in which political embeddedness shapes access to scarce inputs and early regulatory advantages but does not reliably produce stronger efficiency-adjusted frontier performance. The resulting rents are real but leaky: they distort allocation without providing durable protection to favored firms against technically stronger competitors. Less embedded labs have frequently played a disproportionate role in pushing the frontier, while successful outsiders have then faced increasing political absorption after breakthrough.

The broader implication is not that industrial policy is absent or that the Chinese state is irrelevant to AI. It is that political embeddedness appears more effective at distributing opportunities than at selecting the most efficient frontier innovators. One further implication, which I can only suggest rather than prove, is that tighter external compute constraints (such as U.S. export controls) may increase the domestic value of state-mediated infrastructure and therefore strengthen the po-

litical leverage attached to compute allocation. That dynamic would push even independent labs into the political orbit. U.S. export controls, the very policy designed to slow Chinese AI, may thereby strengthen the very feature of the system that this paper identifies as the most durable rent. Future research should test whether post-success absorption changes the innovation incentives of firms such as DeepSeek over time.

References

Akcigit, U., S. Baslandze, and F. Lotti. 2023. "Connecting to Power: Political Connections, Innovation, and Firm Dynamics." *Econometrica* 91 (2):529–564.

Baumol, W. J. 1990. "Entrepreneurship: Productive, Unproductive, and Destructive." *J.P.E.* 98 (5, pt. 1):893–921.

Beraja, M., A. Kao, D. Y. Yang, and N. Yuchtman. 2023. "AI-tocracy." *Q.J.E.* 138 (3):1349–1402.

Cave, D., F. Ryan, and V. X. Xu. 2019. "Mapping More of China's Tech Giants: AI and Surveillance." Tech. rep., Australian Strategic Policy Institute. https://www.aspi.org.au/report/mapping-more-chinas-tech-giants.

Chan, K., G. Smith, J. Goodrich, G. DiPippo, and K. F. Pilz. 2025. "Full Stack: China's Evolving Industrial Policy for AI." RAND Perspectives PE-A4012-1, RAND Corporation. https://www.rand.org/pubs/perspectives/PEA4012-1.html.

Chen, C. 2026. "What's Next for Chinese Open-Source AI." *MIT Technology Review* https://www.technologyreview.com/2026/02/12/1132811/whats-next-for-chinese-open-source-ai.

Cherif, R. and F. Hasanov. 2019. "The Return of the Policy That Shall Not Be Named: Principles of Industrial Policy." IMF Working Paper WP/19/74, International Monetary Fund.

Cheung, S. and K.-s. Lau. 2025. "DeepSeek Use in PRC Military and Public Security Systems." *China Brief* https://jamestown.org/program/deepseek-use-in-prc-military-and-public-security-systems/.

CNBC. 2026. "Anthropic Joins OpenAI in Flagging 'Industrial-Scale' Distillation Campaigns by Chinese AI Firms." *CNBC* https://www.cnbc.com/2026/02/24/anthropic-openai-china-firms-distillation-deepseek.html.

DeepSeek-AI. 2024. "DeepSeek-V3 Technical Report." Tech. rep., DeepSeek-AI. https://arxiv.org/abs/2412.19437.

———. 2025a. "DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning." Tech. rep., DeepSeek-AI. https://arxiv.org/abs/2501.12948.

———. 2025b. "DeepSeek-R1 Release." Tech. rep., DeepSeek-AI. https://api-docs.deepseek.com/news/news250120.

Ding, J. 2019. "ChinAI #51: China's AI 'National Team'." ChinAI Newsletter. https://chinai.substack.com/p/chinai-51-chinas-ai-national-team.

Emberson, L. 2026. "Chinese AI Models Have Lagged the US Frontier by 7 Months on Average since 2023." https://epoch.ai/data-insights/us-vs-china-eci.

Exiger. 2025. "DeepSeek's Deception: How the Chinese Military and Government Funded DeepSeek's AI Research." https://exiger.com/perspectives/deepseek-chinese-government-funded-ai.

Huang, X. and R. Zhu. 2025. "Selective Industrial Policy and Innovation Efficiency in Chinese Listed Companies." *Econ. Transition and Inst. Change* 34 (1):117–136.

James, L. 2025. "China Officials Overseeing Allocation of High-End AI Chips, Prioritizing Homegrown Options." *Tom's Hardware* https://www.tomshardware.com/tech-industry/semiconductors/chinas-ai-chip-pivot-accelerates-as-us-export-restrictions-bite.

Juhasz, R., N. Lane, and D. Rodrik. 2024. "The New Economics of Industrial Policy." *Ann. Rev. Econ.* 16 (1):213–242.

Krueger, A. O. 1974. "The Political Economy of the Rent-Seeking Society." *A.E.R.* 64 (3):291–303.

Larsen, B. 2019. "Drafting China's National AI Team for Governance." https://digichina.stanford.edu/work/drafting-chinas-national-ai-team-for-governance.

Li, X. and C. Ban. 2025. "Financing Technological Innovation in China: Neo-Developmental Financial Statecraft through Government Guidance Funds." *Rev. Int. Polit. Econ.* 33 (1):131–161.

Mazzucato, M. 2013. *The Entrepreneurial State: Debunking Public vs. Private Sector Myths*. London: Anthem Press.

Meinhardt, C., S. Nong, G. Webster, T. Hashimoto, and C. D. Manning. 2025. "Beyond DeepSeek: China's Diverse Open-Weight AI Ecosystem and Its Policy Implications." Tech. rep., Stanford HAI. https://hai.stanford.edu/policy/beyond-deepseek-chinas-diverse-open-weight-ai-ecosystem-and-its-policy-implications.

Meta AI. 2024. "The Llama 3 Herd of Models." Tech. rep., Meta AI. https://arxiv.org/abs/2407.21783.

Mo, L. and C. Hall. 2024. "Baidu Launches Upgraded AI Model, Says Ernie Bot Hits 300 Million Users." *Reuters via Yahoo Tech* https://tech.yahoo.com/ai/articles/baidu-launches-upgraded-ai-model-103349292.html.

Murphy, K. M., A. Shleifer, and R. W. Vishny. 1991. "The Allocation of Talent: Implications for Growth." *Q.J.E.* 106 (2):503–530.

Olcott, E. and Z. Wu. 2025. "DeepSeek's Next AI Model Delayed by Attempt to Use Chinese Chips." *Financial Times* https://www.ft.com/content/eb984646-6320-4bfe-a78d-a1da2274b092.

Patel, D., A. Kourabi, D. O'Laughlin, and R. Knuhtsen. 2025. "DeepSeek Debates: Chinese Leadership on Cost, True Training Cost, Closed Model Margin Impacts." https://newsletter.semianalysis.com/p/deepseek-debates.

Reuters. 2025a. "Baidu to Launch Updated Ernie AI Model by Mid-March." *Reuters via Investing.com* https://www.investing.com/news/stock-market-news/baidu-to-launch-updated-ernie-ai-model-by-midmarch--reuters-3894902.

———. 2025b. "DeepSeek Sparks AI Stock Selloff; Nvidia Posts Record Market-Cap Loss." *Reuters via Yahoo Finance* https://finance.yahoo.com/news/chinas-deepseek-sets-off-ai-071458415.html.

Scanlon, R. 2025. "Beyond DeepSeek: How China's AI Ecosystem Fuels Breakthroughs." Lawfare. https://www.lawfaremedia.org/article/beyond-deepseek--how-china-s-ai-ecosystem-fuels-breakthroughs.

Schaefer, K. 2025. "Seeking the Next DeepSeek: What China's Generative AI Registration Data Can Tell Us about China's AI Competitiveness." https://triviumchina.com/research/seeking-the-next-deepseek-what-chinas-generative-ai-registration-data-can-tell-us-abo

Singer, S. and M. Sheehan. 2025. "China's AI Policy at the Crossroads: Balancing Development and Control in the DeepSeek Era." Tech. rep., Carnegie Endowment for International Peace. https://carnegieendowment.org/research/2025/07/chinas-ai-policy-in-the-deepseek-era.

Soo, Z. 2023. "China's Baidu Makes AI Chatbot Ernie Bot Publicly Available." *Associated Press* https://apnews.com/article/baidu-ai-chatbot-ernie-chatgpt-627bd09608816847907d41f44da235d9.

Wang, Z. et al. 2025. "Corporate Political Connection, Rent Seeking and Government Procurement Order." SSRN, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5210652.

Welch, N. 2024. "SB 1047 with Socialist Characteristics: China's Algorithm Registry in the LLM Era." *ChinaTalk* With reporting by pseudonymous contributor Bit Wise. https://www.chinatalk.media/p/sb-1047-with-socialist-characteristics.

Ye, J. 2024. "Baidu Says AI Chatbot Ernie Has Attracted 200 Million Users." *Reuters via Yahoo Tech* https://tech.yahoo.com/ai/articles/baidu-says-ai-chatbot-ernie-023012481.html.

Zhang, A. H. 2025. "The Promise and Perils of China's Regulation of Artificial Intelligence." *Columbia J. Transnational Law* 63 (1). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4708676.

Zhao, P. 2025. "China's Generative AI: Building Innovation through Regulation." Hello China Tech. https://hellochinatech.com/p/chinas-generative-ai-regulation-2025.

Zhong, W. and Y. Zheng. 2025. "Do Political Connections Stifle Firm Innovation? Natural Experimental Evidence from China's Anti-Corruption Campaigns." *Strat. Mgmt. J.* 46 (5):1183–1216.