



TAKSHASHILA
INSTITUTION

Should the US Sell Advanced GPUs to China? An Indian Perspective

A Causal Analysis

Pranay Kotasthane

Takshashila Issue Brief 2026-15.
Version 1.0, April 2026.

India benefits from the existence of multiple competitive AI ecosystems, including China's.

Recommended Citation:

Pranay Kotasthane, "Should the US Sell Advanced GPUs to China? An Indian Perspective", Takshashila Issue Brief 2026-15., Version 1.0, April 2026., The Takshashila Institution

©The Takshashila Institution, 2025

Contents

1 Executive Summary	2
2 An Unresolved Debate	2
3 Chip Denial Works Narrowly	4
4 Where India Sits in the Five-Layer Stack	4
5 Considerations for India	6
6 Causal Structure of the Debate	9
7 What India Should Do	11

1 Executive Summary

This document examines the Nvidia-China question from India's vantage point. It makes four arguments:

1. US chip export controls on China are having diminishing returns. They restricted leading-edge chips but failed to restrict AI capability. DeepSeek and China's broader AI ecosystem are the empirical evidence. The causal chain of "deny chips, deny AI" is broken. While restricting top-end lithography machines might still make limited strategic sense, restricting chips doesn't.
2. India is a consumer-integrator economy in AI: strong at the applications layer, embedded in chip design, but dependent on imported compute and models at other layers. This position is structurally different from that of either the US or China, and India's interests do not map onto either side of the American debate.
3. India benefits from the existence of multiple competitive AI ecosystems, including China's. A monopolar AI supply structure makes India more dependent, not less. Open-source Chinese models are usable inputs for India precisely because they are separable from Chinese hardware.
4. India's policy priorities should be to secure reliable access to compute through Pax Silica, fund open-source GPU-agnostic middleware to reduce CUDA lock-in, and maintain model pluralism across American, Chinese open source, and sovereign Indian models.

2 An Unresolved Debate

The question of Nvidia selling its advanced GPUs to China has become geopolitically charged in recent years. In October 2022, the Biden administration imposed export controls that restricted China's access to advanced AI chips, such as Nvidia's A100.¹ In October 2023, these were broadened to cover chips built on the same architectures.² Nvidia engineered the H20, a deliberately downgraded chip, to comply with the restrictions and continue selling to China. In April 2025, the Trump administration then banned even the compliant H20, costing Nvidia a \$5.5 billion writedown on unsold inventory.³ In July 2025, it reversed this move and allowed H20 sales to resume, after lobbying from Jensen Huang and trade negotiations with Beijing.⁴

It was China's turn to put restrictions now. It instructed companies, including ByteDance, Alibaba, and Tencent, to halt H20 purchases, citing security concerns about potential backdoors.⁵ By August 2025, Nvidia had asked its suppliers to suspend H20 production.⁶ In January 2026, Trump announced a new arrangement: Nvidia could sell the more powerful H200 to approved Chinese customers, with a 25 per cent surcharge paid

Lokendra Sharma's editorial comments improved the paper substantially.

AI use disclosure: Claude Opus 4.6 Skills were used to review the paper, brainstorm logical inconsistencies, and for generating causal loops.

to the US government.⁷ As of February 2026, the US government had approved the sale of a small number of H200s to China, but it was unclear whether Chinese regulators would approve the imports.⁸

The underlying question remains unresolved: should the US sell advanced AI chips to China? For most Indian policymakers and technology strategists, this registers as an American debate about American interests. There are stridently opposing views in the US technology industry on this question. Jensen Huang wants to sell; Dario Amodei wants to restrict.⁹

However, the outcome of this debate will shape the structure of the global AI supply chain for the next decade. It will determine whether India has one source of advanced compute or several, whether the open-source model ecosystem continues to receive contributions from Chinese labs, and whether CUDA remains the only viable computing stack for AI workloads. Each of these outcomes has direct consequences for India's AI ambitions.

CUDA is Nvidia's platform for accelerated computing, which enables applications to harness the power of GPUs. See the explainer video [here](#). CUDA is a full *software ecosystem* that makes Nvidia GPUs the easiest, safest bet for anyone doing heavy compute. Due to network effects built over a decade, switching costs are high, and competitors have been unable to break through. Because of CUDA, Nvidia doesn't just sell GPUs, but it sells a full AI platform: hardware, drivers, libraries, SDKs, reference designs, and a massive knowledge base.

In a recent conversation with Dwarkesh Patel, Nvidia CEO Jensen Huang argued that AI is a "five-layer cake" of energy, chips, infrastructure, models, and applications.¹⁰ American policy, in his telling, sacrifices an entire layer (chips) to protect another (models), while China has enough compute, energy, and talent to build competitive AI regardless. On the other side, Anthropic CEO Dario Amodei has compared selling AI chips to "selling nuclear weapons to North Korea and then bragging that the missile casings are made by Boeing."¹¹ Amodei's position underpins much of the current US export control regime, which is: deny compute, deny capability.

Both sides disagree on the *American* strategy. Missing from the conversation is the perspective of countries like India, which are embedded in the semiconductor supply chain and affected by the outcome but have no say in it. India does not want China to dominate all layers of AI. But India also does not benefit from a world where access to advanced chips is contingent on the policy mood of the US administration. India's strategic calculus is shaped by several factors — embeddedness in the design layer of the semiconductor supply chain, dependence on imported compute infrastructure, a large market for AI applications, and recent policy volatility from the US.

The [interview](#) with Jensen Huang on the *Dwarkesh Podcast* was hotly debated.

3 Chip Denial Works Narrowly

The logic behind US export controls on China's semiconductor sector rests on this causal chain: deny advanced chips and chip manufacturing equipment, deny compute, deny AI capability. Each link in this chain needs to hold for the policy to work. The evidence now suggests that the final link is broken.

The controls have succeeded at their immediate objective. China remains unable to manufacture chips at the leading edge. Without access to EUV lithography machines, SMIC is stuck at 7nm, while TSMC and its customers have moved to 3nm and beyond.¹² In this narrow sense, the denial strategy has worked.

But the implied objective of constraining China's AI ecosystem has been achieved only partially. DeepSeek and Kimi's emergence as competitive frontier AI labs is the most visible evidence.¹³ But it is not the only evidence. Huawei reported revenue of CNY 880.9 billion in 2025, its second-highest annual revenue ever, and invested a record CNY 192.3 billion in R&D.¹⁴ Chinese contributions to open-source AI models are among the largest in the world. China's AI researcher base constitutes roughly half the global talent pool.

What went wrong with the causal chain? Huang's five-layer framework offers one explanation. AI capability is not a linear function of chip quality. It is a product of interactions across energy, chips, computing stack, models, and applications. The controls perhaps did buy some time for the US in the semiconductor domain, perhaps two to three years on leading-edge nodes. Thus, it is quite likely that at least some Chinese AI firms are constrained by compute and must innovate in other layers to offset the crunch.¹⁵ Some of these firms have compensated for weaker chips with abundant energy (through brute-force scaling of less efficient hardware), algorithmic innovation at the model layer (mixture-of-experts architectures, attention-mechanism improvements), and a large engineering workforce that optimised within constraints. The controls restricted one layer, and the other layers adapted.

4 Where India Sits in the Five-Layer Stack

If we adopt Huang's five-layer framework, India's position is uneven. It is strong at one layer, present at another, and thin in the other three.

Energy. India is constrained. Unlike China, which Huang describes as having abundant energy and fully powered but empty data centres, India cannot compensate for chip quality with energy quantity. India's per-capita energy consumption remains among the lowest of major economies, and its data centre capacity, while growing, is orders of magnitude smaller than that of the US or China. Despite the promise that renewables

hold for India, the current regulatory architecture for electricity prevents rapid growth and transition. This is a binding constraint on sovereign AI compute ambitions.

Chips. India designs but does not yet manufacture at a commercial scale. Nearly every top-25 fabless semiconductor company has a design centre in India, and the country's engineering workforce is embedded in the global chip design supply chain.¹⁶ Eight OSAT plants and one commercial CMOS fab are under construction. India's accession to the Pax Silica grouping in February 2026 is designed to secure access to chips, chip-making software, and manufacturing equipment.¹⁷ But as of now, India remains primarily a design services provider. The IP, with few exceptions, belongs to firms headquartered elsewhere.¹⁸

India's engineering workforce is also deeply embedded in the CUDA ecosystem, Nvidia's software platform for programming its GPUs. CUDA's dominance means Indian engineers are globally competitive. But it also means India's AI software stack is locked into a single chip vendor's ecosystem. This is a case of the distinction between *supply access* and *capability ownership*.¹⁹ India has extensive supply access to the CUDA ecosystem. It has little ownership over it.

Infrastructure. Huang defines this layer as land, power delivery, cooling, networking, and the systems that orchestrate processors into AI factories.²⁰ India's total data centre capacity is estimated to be 1.5GW, an order of magnitude less than that of China and the US. Moreover, given the higher temperatures in India, it needs to invest in advanced cooling systems, reduce water consumption, and shift new demand to renewables through better pricing structures. The delays related to land and construction are other weaknesses in this layer.

Models. India now has an early sovereign model capability. Sarvam AI's 105B-parameter model, trained on domestic compute infrastructure under the IndiaAI Mission, was released in March 2026 with open-source weights.²¹ It uses a mixture-of-experts architecture and has been trained specifically on the corpus of the ten most-spoken Indian languages. This is a significant milestone because it's the first Indian model trained from scratch at this scale, rather than fine-tuned on a foreign base model. It does not match frontier models in scale, but it demonstrates that sovereign model development is feasible for specific use cases, particularly multilingual applications, government deployments, and defence.

Beyond sovereign development, India benefits from model pluralism: using American proprietary models, Chinese open-source models (DeepSeek, Qwen, Kimi K2), and domestically developed models for different purposes. The model layer is where India has the most optionality, precisely because open-source models are separable from the hardware infrastructure on which they were trained.

Applications. This is India's strongest layer. Indian IT services firms are positioned to become the global delivery vehicle for

Check this [paper](#) by Akshay Jaitly and Ajay Shah documenting the issues of electricity reform in India.

The list of Indian fabless firms that own strategic semiconductor design intellectual property is small. The notification of the Design Linked Incentive (DLI) scheme mentions that "the cumulative annual revenue of domestic semiconductor design companies estimated to be less than ₹150 Crore." Some firms which are developing domestic IP through the DLI scheme are listed [here](#).

enterprise AI deployment. India is one of the world's largest data markets and will be a decisive arena for AI adoption across sectors. As Kai-fu Lee's distinction between "internet AI" and "enterprise AI" suggests, the economic value of AI will increasingly flow through integration, deployment, and customisation.²² These are the capabilities India's services sector has built over three decades.

The picture that emerges from this discussion is that India is a *consumer-integrator* AI economy. India consumes chips and models produced elsewhere, and integrates them into data centres, applications and services. This is not inherently weak; consumer-integrators can capture significant value, as India's IT services industry demonstrates. But it creates a strategic calculus different from either the US (which dominates production at multiple layers) or China (which is building indigenous capability across all layers). India's strategic priority is to ensure reliable access across all layers while gaining leverage at the application layer.

This calculus has a direct implication for the Nvidia-China debate. Since India cannot build all five layers itself in the foreseeable future, India has an interest in the existence of *multiple* competitive AI ecosystems globally. A world in which China develops a capable AI stack — and, more specifically, releases key components of it as open source — is a world in which India has alternatives.

DeepSeek and Qwen are not threats to India's AI ambitions; they are inputs. The more competitive the global model ecosystem, the less leverage any single country holds over India's access. India should therefore be cautious about endorsing the US denial strategy because its success would leave India more dependent on a single supplier. A supplier who knows it has no competition has less incentive to offer reliable access on reasonable terms. From India's perspective, a thriving Chinese open-source AI ecosystem is preferable to a monopolar one, even when accounting for the security concerns that Chinese hardware dependencies legitimately raise.

5 Considerations for India

India's AI hardware strategy must address four specific considerations.

The American stack is optimal, but access is unreliable. India's AI ecosystem runs on American infrastructure: Nvidia GPUs, CUDA libraries, AWS and Azure cloud services, OpenAI and Anthropic APIs. This is the most capable stack available.

The problem is access reliability. The Biden administration's AI Diffusion framework, issued in January 2025, created tiered access categories that would have capped India's chip imports at 50,000 H100-equivalents from 2025 to 2027. The Trump

administration rescinded the rule in May 2025, calling it “overly complex, overly bureaucratic,” and indicated it would replace it with a country-specific approach based on bilateral negotiations. As of this writing, no replacement has been finalised. For a country building long-term AI capability, this policy volatility is itself a strategic risk.

India’s Pax Silica membership is meant to mitigate this risk. But membership in a grouping is not the same as guaranteed access at specific volumes and price points. The terms of India’s compute access, including new cloud compute commitments and model accessibility, should be a central negotiating priority within Pax Silica.

The Chinese hardware stack is not an option, but Chinese models are usable, and their proliferation is in India’s interest.

India will not adopt Huawei Ascend chips for its AI infrastructure. This is a strategic choice driven by security concerns about hardware-level dependencies on a geopolitical rival. The Indian government’s exclusion of Huawei from 5G networks and its broader defensive decoupling from Chinese technology supply chains make this position clear. But hardware and models are separable. Open-source Chinese AI models run on any compliant hardware and are available for fine-tuning. For commercial and enterprise deployment — the bulk of India’s near-term AI activity — fine-tuning on open weights addresses the principal security concerns: behavioural misalignment can be corrected, censorship artefacts can be overridden, and outputs can be calibrated for the Indian context.

The calculus is different for defence, intelligence, and critical infrastructure, where epistemic concerns override utility. Pre-training data composition cannot be audited retrospectively. More critically, research on adversarially trained language models has shown that some kinds of backdoor triggers can survive standard fine-tuning, i.e., a model conditioned to behave differently under specific inputs can retain that behaviour through subsequent alignment procedures. For applications where the adversary may have had prior access to the training process, this residual uncertainty is an instant disqualifier.²³

More broadly, India has an interest in Chinese AI capability being expressed through open-source releases rather than proprietary, hardware-locked systems. Every open-source model released by a Chinese lab is a model India can use without depending on either China’s hardware stack or America’s goodwill. This is how a consumer-integrator benefits from supplier competition.

The policy implication is that India should resist pressure to join any coalition aimed at restricting the global release of open-source AI models. Such restrictions would reduce India’s optionality and hence are not desirable.

CUDA dominance is a bottleneck India should work to loosen.

Indian engineers benefit from CUDA’s dominance today; it makes

their skills globally portable. But India's long-term interest is in a more competitive accelerator ecosystem, because competition creates optionality. If the only viable AI compute stack is CUDA, India's AI capability is hostage to a single company's product roadmap and a single government's export policies.

Indian firms are trying to build hardware accelerators from scratch, and the India AI Mission and the India Semiconductor Mission must shift their focus to supporting such products. But India also has another asymmetric option at hand: fund open-source projects that operate at the layer between hardware and frameworks like PyTorch, creating GPU-agnostic middleware that makes it easier to switch between accelerators. India's AI Mission should make this a funding priority. The cost is modest compared to chip manufacturing subsidies, and the strategic payoff of reduced dependence on any single hardware vendor is significant.

Sovereign capability is necessary for a narrow set of applications. Full-spectrum AI sovereignty is neither achievable nor desirable. But for military AI, critical infrastructure protection, and intelligence analysis, India needs models and inference infrastructure that are not dependent on any foreign provider's continued goodwill. This does not require training frontier models from scratch. It requires sovereign fine-tuning of reliable open-source models, deployed on infrastructure that India controls. Sarvam 105B's development on IndiaAI Mission compute demonstrates that this path is workable. The invest-in-everything approach is wasteful; a targeted approach to sovereignty, focused on the narrow set of applications where dependency becomes a strategic vulnerability, is both realistic and more effective.

Apache's TVM, MLIR (within LLVM), and ONNX Runtime from Microsoft are examples of projects operating at different layers that aim to make the GPU a switchable backend layer.

The recent US debate over chip exports has been shaped by Anthropic's Mythos model, whose cyber-offensive capabilities led Anthropic to delay its public release while American companies patched their software. This framing has three problems. First, it follows the narrative that AI is akin to nuclear weapons, which means that getting there first at all costs and then denying others access to it is the only viable strategy. However, that's an incorrect frame because developments like Mythos don't structurally change determinants of national power as other companies will get there soon—as Qihoo 360's launch illustrates. For countries like India, it's better to think of AI as a vital, general-purpose technology. Second, it assumes that China would use such capabilities offensively while the US would use them defensively. Both countries have offensive cyber capabilities, and both have used them. The threat is not China-specific malice; it is the absence of credible governance mechanisms for AI deployment, a problem neither side has solved. Third, the Mythos episode reveals a vulnerability that exists regardless of who trained the model. If AI models can find thousands of unpatched vulnerabilities in major software systems, then every country running those systems is exposed, including India. The lesson is about defensive readiness, not about chip exports.

6 Causal Structure of the Debate

The arguments in this paper can be represented as a causal loop diagram (Figures 1 and 2).

The system has eight variables. “US chip exports to China” sits at the top as the exogenous variable whose value is set by American policy. When exports increase, China develops more on the US stack. When exports are restricted, China’s development on the US stack falls.

This variable is the branching point from which the system’s two loops originate.

The left branch traces Jensen Huang’s logic. China’s development on the US tech stack feeds US chip industry revenue and R&D capacity (+), which reinforces exports (+), forming R2, the ecosystem flywheel.

The right branch traces the denial logic. China’s development on the US tech stack has an opposite-direction (–) relationship with the intensity of China’s substitution efforts, i.e., when development on the US stack falls, substitution efforts intensify. This feeds China’s AI output on own hardware (+). China’s AI output feeds back into substitution efforts (+), forming R1, the indigenous acceleration loop. Once triggered, R1 is self-sustaining. It is the causal structure behind the argument from *When the Chips Are Down* that keeping an adversary dependent on your products is an underrated strategy: the US abandoned this approach and triggered a self-reinforcing cycle it cannot easily reverse.

A cross-connection links the two loops: China’s AI output on own hardware erodes US chip industry revenue and R&D capacity (–). R1 actively undermines the conditions for R2.

Two channels feed “open-source models India can use.” China’s development on the US tech stack feeds it positively (+)—Chinese developers building on CUDA produce models natively compatible with India’s infrastructure. China’s AI output on own hardware also feeds it positively (+) as open-source model weights are separable from the hardware they were trained on, as DeepSeek and Qwen demonstrate. India gets usable models regardless of which path dominates.

This is the key insight. India’s model availability is not the variable that differentiates the two scenarios. Whether China builds on the US stack or builds on its own hardware, India gets open-source models either way. What differentiates the scenarios is the health of R2, the ecosystem flywheel. When R2 operates, India’s primary compute infrastructure supplier remains competitive, access terms are commercially motivated, and models produced on the US stack are natively compatible with India’s infrastructure. When R1 is triggered and erodes R2, India faces access uncertainty from a weakened supplier with less commercial incentive to maintain reliable terms.

Figure 1: the system as it operates today

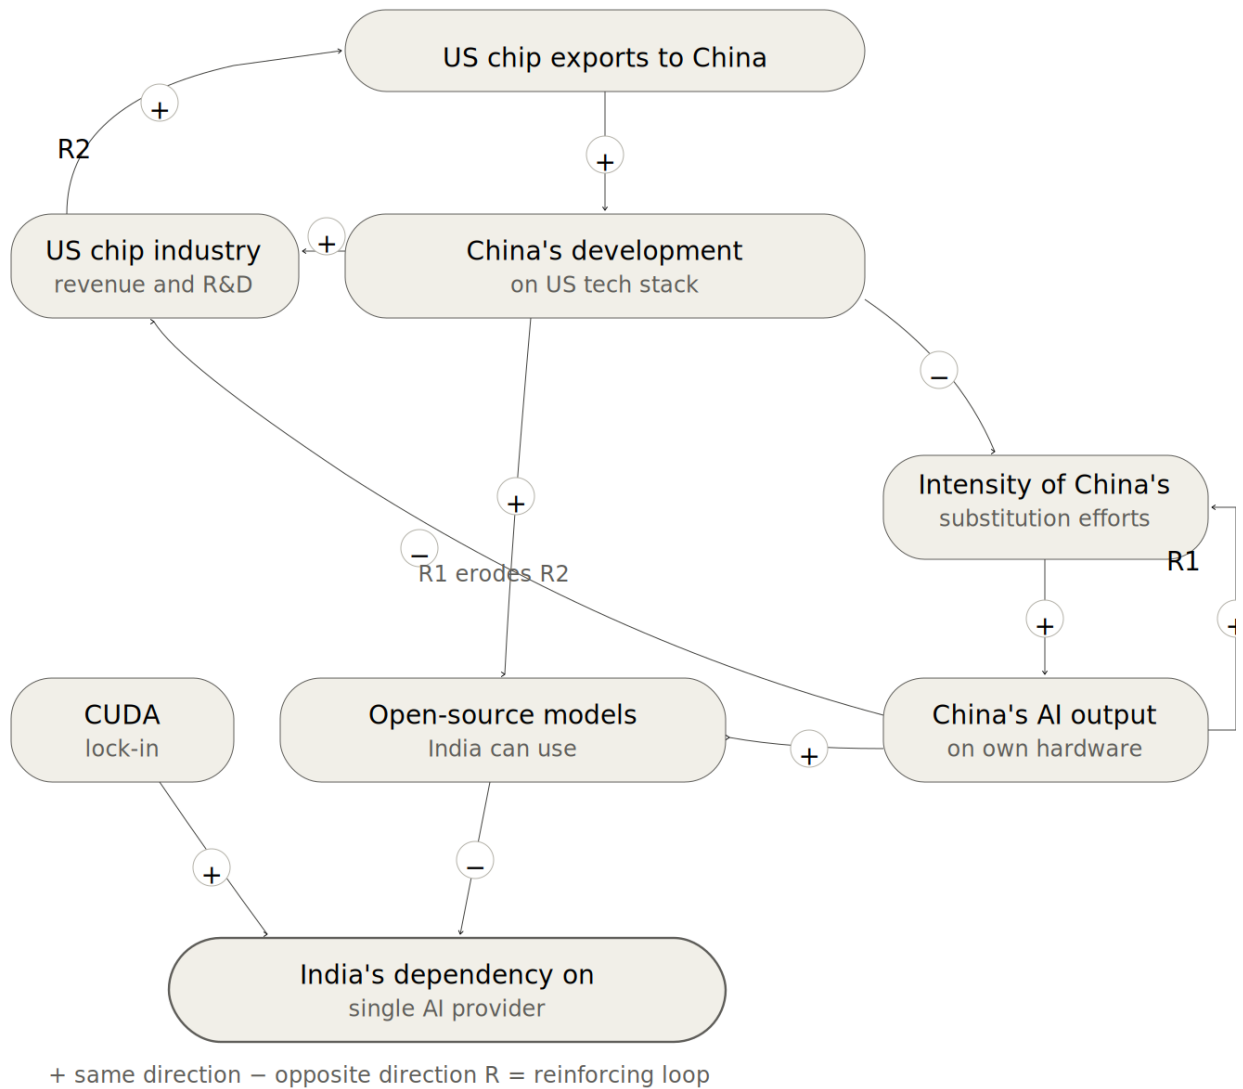


Figure 1: Figure 1: The system as it operates today

India sits at the bottom as a downstream outcome variable. Two forces act on India's dependency on a single AI provider: open-source models India can use pushes dependency down (-), while CUDA lock-in pushes dependency up (+).

Figure 2 shows the same system with three exogenous inputs India can add. GPU-agnostic middleware reduces CUDA lock-in (-). Sovereign model development (Sarvam 105B, fine-tuned open-source models) adds to models India can use (+), independently of what happens in either branch. Pax Silica access guarantees reduce India's dependency directly (-), through institutional commitments rather than market dynamics.

Figure 2: the system with India's interventions

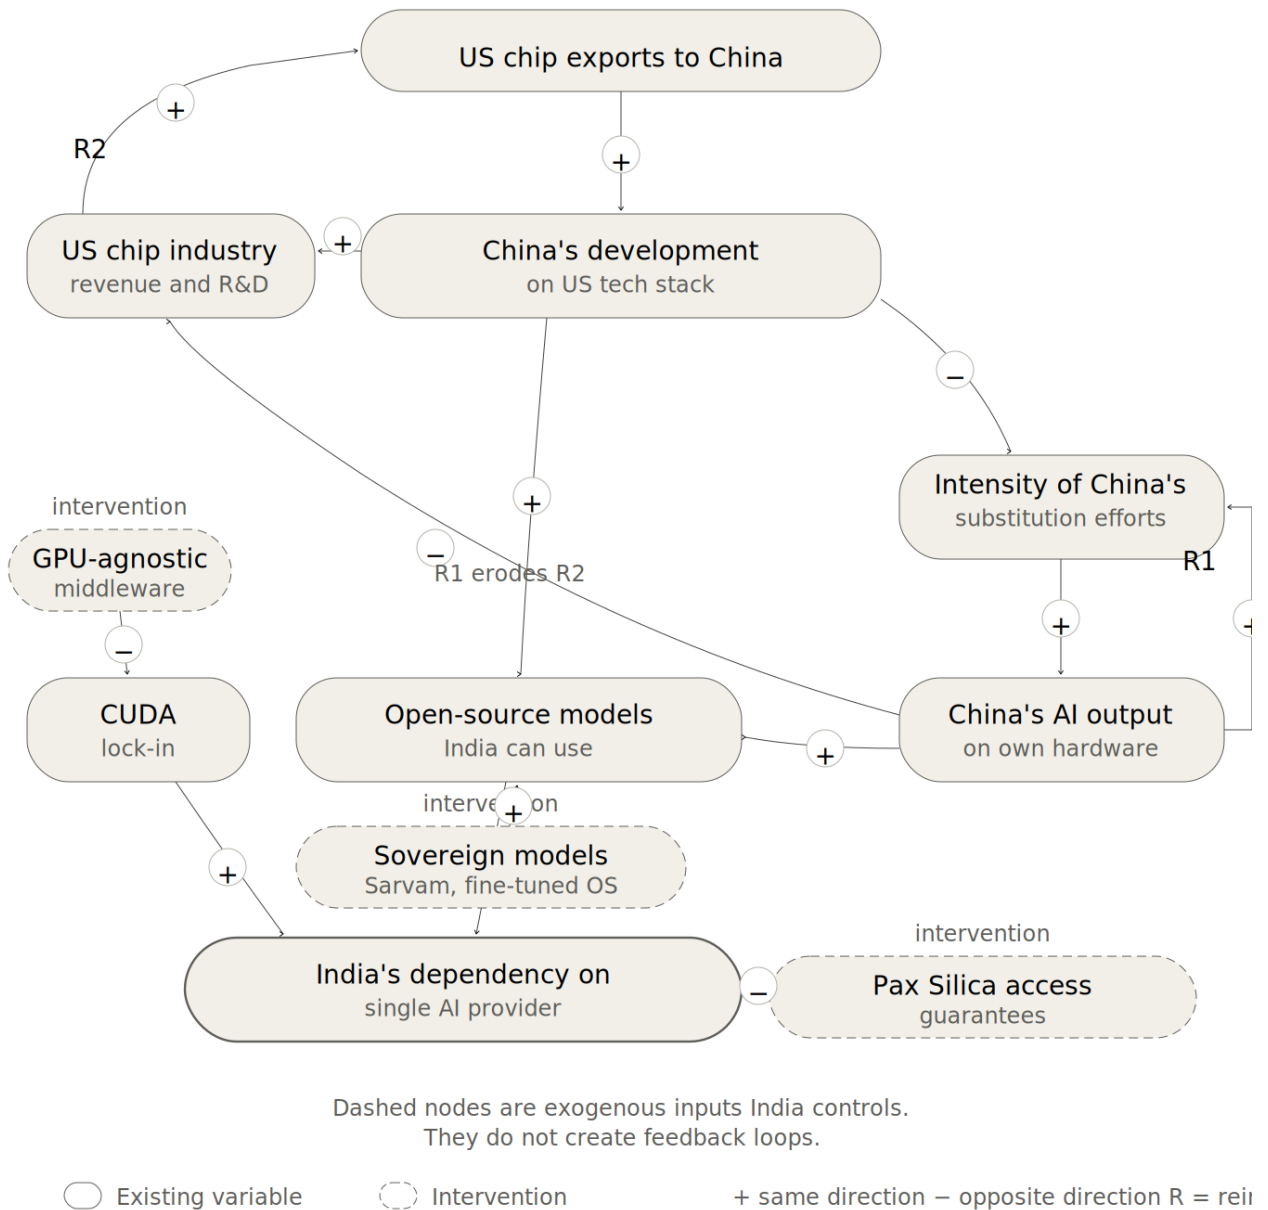


Figure 2: Figure 2: the system with India's interventions

These are leverage points that reduce India's exposure to a system it does not control.

7 What India Should Do

India's strategic posture on the Nvidia-China question should be guided by a principle: maximise access across all five layers while

minimising dependence on any single provider.

First, India's primary economic play in AI is at the applications layer. Policy should be oriented toward accelerating the diffusion of AI in the Indian industry and public services.

Second, India should use its Pax Silica membership to negotiate guaranteed compute access, including commitments to cloud compute capacity and model access on predictable terms.

Third, India should invest in open-source GPU-agnostic middleware projects that reduce CUDA lock-in. This is a high-leverage, low-cost intervention that creates optionality across the compute stack. India's AI Mission should identify and fund specific projects operating at the layer below PyTorch that enable hardware-agnostic AI workloads.

Fourth, India should maintain model pluralism as a deliberate strategy. Use American proprietary models where they are best. Use Chinese open-source models where they are useful. Build sovereign models like Sarvam for applications where dependency on foreign providers is unacceptable. The model layer is where India has the most freedom to operate.

Endnotes

1. Gregory C. Allen, "Choking Off China's Access to the Future of AI," *CSIS*, October 2022, [Link](#); Pranay Kotasthane and Abhiram Manchi, *When the Chips Are Down* (Bloomsbury, 2023), chapter 8.
2. Bureau of Industry and Security, "Export Controls on Semiconductor Manufacturing Items," Interim Final Rule, *Federal Register* 88, no. 73424, October 17, 2023. The revised controls targeted chips engineered to circumvent the 2022 thresholds, including Nvidia's A800 and H800, by imposing new performance density and interconnect bandwidth criteria. [Link](#)
3. "Nvidia Says It Will Record \$5.5 Billion Charge Tied to H20 Processors Exported to China," *CNBC*, 15 April 2025. [Link](#)
4. Dylan Butts, "Nvidia Says It Will Resume H20 AI Chip Sales to China 'soon,' Following U.S. Government Assurances," *CNBC*, 15 July 2025. [Link](#)
5. "China Cautions Tech Firms over Nvidia H20 AI Chip Purchases, Sources Say," *Reuters*, 12 August 2025. [Link](#)
6. "Nvidia Orders Suppliers to Halt Work on China-Focussed H20 AI Chip, The Information Says," *Reuters*, 22 August 2025. [Link](#)
7. Kif Leswing, "Trump Administration Clears Way for Nvidia H200 Chip Sales to China with a 25% Surcharge," *CNBC*, 14 January 2026. [Link](#)
8. "Nvidia Gets US Licence for Small Amount of H200 Exports to China," *Moneycontrol*, 26 February 2026. [Link](#)
9. Dario Amodei, "The Adolescence of Technology," accessed 21 April 2026. [Link](#)
10. Dwarkesh Patel, "Jensen Huang – TPU Competition, Why We Should Sell Chips to China, & Nvidia's Supply Chain Moat," 7 April 2026. [Link](#)
11. Dario Amodei, "The Adolescence of Technology," accessed 21 April 2026. [Link](#)
12. "Watch Huawei Mate 60 Pro Teardown Reveals Kirin 9000s Chip in China Breakthrough," *Bloomberg*. [Link](#)
13. DeepSeek's V3 and R1 technical reports document the hardware constraints and algorithmic workarounds used to train competitive models on H800-equivalent hardware.
14. Huawei, *2025 Annual Report*, 31 March 2026. [Link](#)
15. Pranay Kotasthane, "#339 Divergence and Assimilation," *Anticipating the Unintended*, 30 March 2026. [Link](#)
16. Douglas Fuller, "Chip Design in China and India: Multinationals, Industry Structure and Development Outcomes in the Integrated Circuit Industry," *Technological Forecasting and Social Change* 81 (2012): 1–10. [Link](#)
17. Pranay Kotasthane and Tannmay Kumarr Baid, "Inside the Tent: India and the Limits of Pax Silica," *India's World*, 16 January 2026. [Link](#)
18. Invest India, "Million Chips, Billion Dreams – II (Developing Fabless Ecosystem in India)," 27 July 2021; Kotasthane and Manchi, *When the Chips Are Down*, pp. 136–137.
19. Kotasthane and Manchi, *When the Chips Are Down*, chapter 8.
20. "AI Is a 5-Layer Cake," *NVIDIA Blog*, 10 March 2026. [Link](#)
21. Sarvam AI, "Open-Sourcing Sarvam 30B and 105B," 6 March 2026. [Link](#)

22. Kai-Fu Lee, *AI Superpowers: China, Silicon Valley, and the New World Order* (Houghton Mifflin Harcourt, 2018).
23. "Persistent Backdoor Attacks under Continual Fine-Tuning of LLMs," *arXiv*, 2023. [Link](#)



TAKSHASHILA
INSTITUTION

The Takshashila Institution is an independent centre for research and education in public policy. It is a non-partisan, non-profit organisation that advocates the values of freedom, openness, tolerance, pluralism, and responsible citizenship. It seeks to transform India through better public policies, bridging the governance gap by developing better public servants, civil society leaders, professionals, and informed citizens.

Takshashila creates change by connecting good people, to good ideas and good networks. It produces independent policy research in a number of areas of governance, it grooms civic leaders through its online education programmes and engages in public discourse through its publications and digital media.

©The Takshashila Institution, 2025