
Not All Should Go South: The Pragmatic AI Strategy for Secondary Powers

Kauã Santos
Insper

Bruno Nobre
Independent

Adre Novais
Federal University of São Carlos

Erickson Leon Kovalski
Federal University of Paraná

With
Apart Research

Abstract

National AI strategies across the Global South increasingly consist in the pursuit of "AI sovereignty", especially domestic training of frontier foundation models. We argue this represents a fundamental strategic error that stems from conceptually mixing the three distinct stages of the AI value chain: model development (pretraining), deployment (hosting, post-training, serving), and usage (application building and adoption). Given the "winner-takes-most" dynamics of frontier development, rapid commoditization of foundational models, and the structurally discrepant cost-benefit of open-weight models, competing at the foundational models layer is both unrealistic and economically irrational for "secondary powers" (i.e, nations other than US and China). In contrast, the deployment and usage layers present faster and possibly greater opportunities for profit and social impact through technological leapfrogging. To correct this strategic misallocation, we propose a directional framework, and tie each component to concrete AI-safety rationales. Ultimately, we offer this pragmatic reframing to guide the future analysis and design of National AI Plans, shifting the focus from the illusion of full-stack sovereignty to actionable downstream resilience.

1. Introduction

If you can't define it, you can't program it. – Naval Ravikant

When people invoke an "AI strategy," their vision is frequently blurred by a failure to distinguish the distinct layers of the artificial intelligence ecosystem. This confusion pushes national plans toward objectives that are, for the vast majority of countries, neither achievable nor desirable, so we shall separate them explicitly. **Development** is the creation of a model from scratch: most demanding, pretraining a foundational model, which requires frontier-scale compute, vast data, and concentrated research talent. **Deployment** is making an existing model usable: hosting it for inference, adapting it through post-training or fine-tuning, and operating the surrounding infrastructure. **Usage** (or **adoption**) is applying models to real problems – building products on top of them and integrating them into firms, governments, and daily life.

The threat model here is not a single misaligned system but a strategic **failure mode**. Driven by the buzzword of "**Sovereign AI**", governments (such as [Brazil](#) and [India](#)) are increasingly attempting to build their own **foundation models** from scratch. A Global South government that pours their scarce resources into this development layer chasing sovereignty will almost certainly **fail to reach the frontier, abdicate the faster returns available downstream, and remain dependent on foreign providers anyway**, thus accumulating geopolitical and supply-chain vulnerability without the compensating benefit.

The safety stakes are concrete. A country that adopts AI without a deliberate deployment posture inherits the risks of dependence (gradual disempowerment, an enlarged attack surface, exposure to provider decisions and export controls) while capturing fewer of the benefits and exerting none of the influence over how the underlying models are built.

This is a position-and-framework paper. We present a reframing and a structured plan, grounded in the current economics of the AI stack and meant to directly improve National AI Plans.

Our contributions are:

1. A conceptual disaggregation of "AI strategy" into development, deployment, and usage, with a **market-grounded argument for initially concentrating on the latter two**;
2. A four-part directional **framework**, each component tied to an explicit safety rationale; and
3. A proposed **agenda for testing** the framework against existing national plans, such as tools to quickly evaluate proposed plans.

2. Related Work

"If you want to go fast, go alone. If you want to go far, go together." – African Proverb

Our argument sits at the intersection of three literatures.

AI and the Global South. Prior work documents the structural exclusion of the Global South from frontier AI and argues for proactive engagement with frontier risk ([Chan et al. 2021](#); [Abungu et al. 2024](#)), and analyzes how compute access shapes who can govern AI ([Hawkins et al. 2024](#)). Furthermore, middle powers risk being entirely sidelined as superintelligence concentrates geopolitical and military power within US jurisdictions and datacenters ([Davidson 2024](#)). This vulnerability makes full-stack "AI sovereignty" highly appealing, yet strategic analyses warn that total technology autarky is an expensive illusion, and, instead, nations must pivot toward a framework of downstream "AI resilience" ([Lang et al. 2026](#)).

The economics of the AI stack. Whether foundation models will commoditize, and where value will then accrue, is a live debate among industry leaders. OpenAI's Sam Altman has articulated a "middle layer" thesis – that few will train base models from scratch, and that durable value comes from taking a large existing model and doing the small fraction of training that matters for a vertical ([Hoffman 2022](#)); the same scenario was treated as serious, not fringe, in a [20VC discussion](#) with Altman, OpenAI's Brad Lightcap, and Mistral's Arthur Mensch. This rapid commoditization is empirically supported by inference costs dropping by orders of magnitude annually ([Epoch 2024](#)). Sequoia's ["Services: The New Software"](#) argues that AI lets a vendor sell the *work* rather than the tool, capturing the services budget. We extend these arguments from venture strategy to the situation of resource-constrained states.

The empirical economics of AI productivity. Macro estimates remain modest ([Acemoglu 2025](#)), while micro evidence on software development is strong but heterogeneous ([Cui et al. 2025](#); [METR 2025](#)). We use this literature to discipline our adoption argument.

Where we differ. Most Global South AI Governance literature focuses on frontier governance, multilingual evaluation, or harm taxonomies. We address a fundamental prior question: which layer of the AI stack a secondary power should invest in at all. We argue that this industrial choice is itself the foundational safety decision. By reframing "sovereignty" from the expensive illusion of full-stack development to the pragmatic reality of downstream resilience, we show that the layer a country chooses directly fixes its systemic risk exposure and its leverage over foreign labs.

3. Methods

Plans are nothing; planning is everything. – Dwight D. Eisenhower

Our method is conceptual analysis: we state premises about the AI market, derive strategic implications, and assemble them into a framework, making the premises explicit so a reader who rejects one can see which conclusions fall with it.

3.1 Premise: foundational-model development is a bad investment for the Global South

Three compounding factors support this. First, **no comparative advantage**: limited compute and concentrated know-how mean Global South actors cannot competitively pretrain at the frontier, and export controls further restrict the most capable accelerators ([CSIS 2024](#); [Brookings 2026](#)). Second, **winner-takes-most economics**: frontier development concentrates returns among a few leaders, so the expected return on a development bet is low under any reasonable scenario. Third, **a cheaper substitute exists**: open-weight models have closed most of the capability gap – by early 2026 the gap on knowledge benchmarks such as MMLU had effectively reached zero against the best closed models ([Stanford AI Index 2025, as reported](#)), and the lag to match the best closed model fell from ~27 weeks (H1 2024) to ~13 weeks (H1 2025) ([EE Times 2026](#)) – and they are far cheaper to serve, roughly 87% lower cost per token in one study ([MIT Sloan 2026](#)). So deploying open models, post-training included, is the smarter move.

Despite these structural barriers, the current policy zeitgeist leans heavily toward domestic development. Middle powers are increasingly confusing "sovereignty" with "pretraining," evidenced by national plans that dedicate from hundreds of millions to a few billion dollars of public budget to building native foundational models rather than deploying existing ones.

One facet of US-China tensions illustrates the point: the alleged "AI dumping." Because AI investment has become a dominant driver of recent US growth, a US economy this exposed to foundation models is correspondingly exposed when those models commoditize. A [US–China Economic and Security Review Commission analysis](#) argues Chinese labs deliberately undercut US rivals through permissive licensing and aggressive pricing, [explicitly analogized to steel dumping](#). For a player that cannot leap ahead on raw performance, competing on cost against subsidized, commoditizing open models is fruitless but *consuming* them is not.

3.2 The framework

The opportunity is not what we thought, but instead may be even better: a first-mover window remains open at the deployment and usage layers. This is attractive for profit (faster, often larger returns, since usage $\approx$ leapfrogging) and for impact (creating a greater final effect – amplifying the **bad**, but also the **good**). Four components follow.

(A) Increase adoption to justify deployment and accumulate expertise. The economy-wide productivity case is, for now, a long shot: [Acemoglu \(2025\)](#) estimates a total-factor-productivity gain of no more than ~0.66% over a decade – possibly under 0.53% – once harder-to-learn tasks are weighed. We therefore do *not* premise the framework on an imminent growth miracle. What *is* nearly undisputed is the effect on software development specifically: a [randomized study across 4,867 developers](#) at Microsoft, Accenture, and a Fortune 100 firm found a ~26% increase in completed tasks, with the effect being larger for junior developers. On the other hand, honesty requires counter-evidence: [METR \(2025\)](#) found *experienced* developers on familiar codebases ~19% *slower* with AI. That asymmetry is itself the lever: if governments subsidize model access for developers, they will build capability precisely among the people who will later hold the comparative advantage to develop applications – converting cheap inputs today into domestic capacity tomorrow.

(B) Deploy models to mitigate supply-chain and military risks. Deploying domestic technology is a matter of resilience to supply-chain risk and national security. Accepting this transition may be worth it: as dependence on AI applications deepens – a gradual disempowerment ([Kulveit et al. 2025](#)) – geopolitical vulnerability to foreign deployers and their providers grows. This matters because unlike in most other contexts, in the military domain AI sovereignty transcends nationalist posturing and becomes a rational priority. The response requires not only datacenters but (i) inference compute proportional to projected usage, especially to make local access viable, and (ii) mission-specific hardware for military self-sufficiency (such as GPUs and NPUs). A focused success here – for instance, agentic white-hat security tooling such as CodeWall – could be an internationally relevant AI-safety contribution *without any foundational-model development*.

(C) Develop applications to yield the greatest benefits. This is the most overlooked bet by the community and also the most powerful, on two fronts. First, *profit through a revolution in services* (e.g., Enter): following [Sequoia's thesis](#), service businesses can be rebuilt around AI that sells the outcome rather than the tool, and because the services market dwarfs the software market (about six dollars spent on services per dollar on software), AI-native applications can also be *exported*, generating revenue that justifies further deployment. Second, *impact through leapfrogging* (e.g., Horizon 1000): poorer economies skip intermediate technologies and adopt advanced ones directly – as many did by going straight to mobile phones without building out landlines. Applied to AI, application adoption delivers the most value fastest, since it needs neither frontier compute nor years of R&D, which can be spectacular for the poor, unequal economies that are the rule across the Global South. A parallel front is government efficiency: in the spirit of [Palantir's pitch](#), AI can cut bureaucracy, sharpen decision-making, and lower the cost of a public machine that consumes a large share of national resources. The obvious risk in both cases is the same – gradual disempowerment and an enlarged attack surface – and must be priced in.

(D) Influence frontier safety from outside the pipeline. Even without building frontier models, Global South actors can shape how safely they are built. This influence rests on two distinct levers: accumulating intellectual capital by producing AI-safety research and growing a domestic

community of AIS researchers whose work feeds back into frontier labs; and – a point underexplored in the Global South AI literature – exercising the **bargaining power of the buyer and intermediary**, much like the EU uses its massive market power to export its regulatory standards.

4. Results

To avoid the [streetlighting effect](#), our work took a first-principles analysis, our methodology doesn't involve quantitative data analysis. Still, to optimize reasoning transparency, let's organize arguments structures and sensitivity to evidence, distinguishing observation from interpretation.

Observation 1: the open–closed capability gap is small and shrinking, and open models are substantially cheaper to serve. Interpretation: the marginal value of *domestic development* is low and falling, while the marginal value of *deployment and usage* is preserved or rising. This inference depends not on the exact numbers, but rather on the *direction* of convergence and the *sign* of the cost gap, both corroborated across independent sources ([MIT Sloan 2026](#); [EE Times 2026](#)), not on the exact numbers.

Observation 2: macro gains are modest, but software-development gains are real and concentrated among less-experienced developers. Interpretation: an adoption subsidy targeted at developers is better justified than a broad "AI will grow GDP" claim, which is why the framework rests on the narrower, better-evidenced one. The [METR \(2025\)](#) slowdown is the main threat to this reading; we keep it explicit, noting it concerns *expert* developers on *familiar* code – a different population from the new entrants a subsidy would target.

The central claim is thus robust to conservative readings: it survives even pessimistic productivity estimates, because its load-bearing premise is the *commoditization* of development, not the magnitude of downstream gains.

5. Discussion on Limitations and Dual-Use Considerations

Implications for AI Safety

A replicable, pragmatic playbook for the Global South. The dominant safety conversation focuses on constraining frontier development in the US and China. For most of the world, the more *tractable levers* sit at the deployment and usage layers: managing dependence, building resilience and military sovereignty, and using buyer leverage as well as a domestic research community to influence frontier labs.

Potentially increasing systemic risks. This work's central recommendation is itself dual-use. The same leapfrogging that delivers public-service and economic gains also enlarges each country's attack surface and deepens the risk of gradual disempowerment through depending on systems

built elsewhere. In particular, government-efficiency applications concentrate sensitive data and decision-making in systems whose providers may be subject to foreign export controls or intelligence laws. A naive reading like "adopt aggressively, worry later" would amplify precisely the risks we aim to mitigate. Therefore, we stress that deployment must be paired from the outset with the resilience measures of component (B) and the influence and research activities of component (D), because adoption without sovereignty and safety capacity is a net increase in risk.

Limitations

Lacking country-specific context, by design. Since the framework is directional and omits country-specific calibration, it should not be applied mechanically: copying the sequencing without local analysis of compute, talent, and institutional capacity could misallocate scarce public resources. The framework is a starting orientation for National AI Plans, not a substitute for the contextual analysis each country must perform.

Short-term dependence on human engineering know-how. Our adoption-and-application argument assumes human engineering expertise will remain a critical, scarce input in the near term. Subsidizing developer access builds durable national capability rather than merely renting it. If AI agents substantially automate software engineering sooner than expected, the mechanism weakens, and the argument would shift toward orchestration and domain expertise rather than engineering headcount.

Lacking deeper application to current analysis. We set out to test the framework against existing National AI Plans but did not execute that comparative analysis. For that reason, until we have future work, the confidence in our framework is validated by argument and external evidence, not yet by application to real plans

Future Work

1. The immediate next step is to execute the deferred analysis of current National AI Plans (e.g., those from Brazil and India) against our framework to identify in greater detail their differences, and from that derive concrete suggestions to the plans. This should be then reviewed by more senior governance researchers, and, after polished, delivered to equip analysts and decision-makers designing these strategies.
2. A more ambitious extension is creating a tool for automated evaluation of National AI Plans. One approach would be using an LLM-as-a-judge to score plans against rubrics distilled from existing governance frameworks (like our own, but not only), and measuring the adherence of National Plans to its proposals. This could be an interesting project to compare frameworks quantitatively. A similar, but more speculative step would be using swarm intelligence tools, like [MiroFish](#), to simulate reactions from and interactions between multiple agents, in order to estimate outcomes.

- a. These extensions, however, are mere adjacent aspirations from the team, compared to the clear, critical step of evaluating plans adherence to the framework.

6. Conclusion

For most of the Global South, the question is not whether to engage with AI but where to spend scarce political and fiscal capital. The default answer, **"build sovereign foundation models", is the worst of the available options**. It spends a country's most limited resource on the one layer where it has no comparative advantage, cannot win, and will remain dependent anyway, all while forgoing the faster returns and the influence over safety that are actually within reach.

Our framework turns that critique into a sequence.

A National AI Plan that takes this analysis seriously should, in practice: (A) subsidize developer access to open models before subsidizing any domestic AI training, building human capacity where the productivity evidence is strongest; (B) scale inference compute rather than training compute, in proportion to projected usage, treating sovereignty as a deployment question rather than a development one; (C) channel public resources toward AI-native applications, especially in service delivery and government efficiency, where leapfrogging pays off fastest; and (D) invest, from the outset and in parallel, in a domestic safety-research community while using the bargaining power of the buyer to shape how frontier models are built. Component (D) should not be treated as an afterthought because without it and the resilience of (B), aggressive adoption is a net increase in risk, not a strategy.

Most importantly, this choice of layer is itself the safety decision. The posture a country adopts does not merely determine its economic returns. **It fixes its exposure to gradual disempowerment**, the size of its attack surface, and its leverage over foreign labs.

Treating AI strategy as a purely industrial question, divorced from safety, is the mistake this paper sets out to correct. **The argument is robust, holding even under the most pessimistic productivity estimates**, because its load-bearing premise is the commoditization of development, not the magnitude of downstream gains. What remains is to test it against the plans already being written, and for the policymakers writing them, to start at the right layer.

Code and Data

N/A. As of now, this work is purely qualitative.

References

1. Acemoglu, D. (2025). **The Simple Macroeconomics of AI**. *Economic Policy*, 40(121), 13–58. <https://doi.org/10.1093/epolic/eiae042>
2. Epoch AI (2024). **Trends in Artificial Intelligence**. <https://epoch.ai/trends>
3. Davidson, T. (2024). **Superintelligence, Geopolitics, and Middle Powers**. X (formerly Twitter). <https://x.com/TomDavidsonX/status/2057500418910154831>
4. Abungu, C., et al. (2024). **Why Global South Countries Need to Care About Highly Capable AI**. CIGI. <https://www.cigionline.org/>
5. Brookings (2026). **Ball Game's Over – the US Is Out of the AI Chip Market in China**. <https://www.brookings.edu/articles/ball-games-over-the-us-is-out-of-the-ai-chip-market-in-china/>
6. Chan, A., Okolo, C., Turner, Z., & Wang, A. (2021). **The Limits of Global Inclusion in AI Development**. <https://arxiv.org/abs/2102.01265>
7. CSIS (2024). **Choking Off China's Access to the Future of AI**. <https://www.csis.org/analysis/choking-chinas-access-future-ai>
8. Cui, Z., et al. (2025). **The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers**. *Management Science*. <https://doi.org/10.1287/mnsc.2025.00535>
9. EE Times (2026). **Open-Source AI Gains Ground as Closed Model Costs Rise**. <https://www.eetimes.com/open-source-ai-gains-ground-as-rising-costs-push-shift-to-smaller-models/>
10. Hawkins, W., et al. (2024). **AI Compute Sovereignty: Infrastructure Control Across Territories, Cloud Providers, and Accelerators**. <https://arxiv.org/abs/2402.08431>
11. Hoffman, R. (2022). **Sam Altman: AI for the Next Era (Greylock)**. <https://greylock.com/greymatter/sam-altman-ai-for-the-next-era/>

12. Kulveit, J., et al. (2025). **Gradual Disempowerment: Systemic Existential Risks from Incremental AI Development.**<https://arxiv.org/abs/2501.16946>
13. MIT Sloan (2026). **AI Open Models Have Benefits. So Why Aren't They More Widely Used?**
<https://mitsloan.mit.edu/ideas-made-to-matter/ai-open-models-have-benefits-so-why-arent-they-more-widely-used>
14. METR (2025). **Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity.**
<https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/>
15. Sequoia Capital (Bek, J., 2026). **Services: The New Software.**
<https://sequoiacap.com/article/services-the-new-software/>
16. Stanford HAI (2025). **AI Index Report 2025, open-vs-closed convergence as reported.**
<https://aibuzz.blog/open-source-vs-closed-source-ai-models/>
17. U.S.–China Economic and Security Review Commission (2026). **Two Loops: How China's Open AI Strategy Reinforces Its Industrial Dominance, as reported in Computerworld.**
<https://www.computerworld.com/article/4149313/chinas-use-of-open%E2%80%91source-ai-threatens-the-us-lead-in-ai-development-us-commission-warns.html>
18. 20VC (2024). **Will Foundation Models Be Commoditised? Is the Value in the Infrastructure or Application Layer?** (*Altman, Lightcap, Mensch et al.*).
<https://www.thetwentyminutevc.com/ai-compilation>
19. Palantir Technologies. **About Palantir.** <https://www.palantir.com/about/>
20. Ministry of Science, Technology and Innovation (MCTI), Government of Brazil (2024). **Plano Brasileiro de Inteligência Artificial (PBIA) 2024-2028.**
<https://www.gov.br/mcti/pt-br/acompanhe-o-mcti/noticias/2024/07/plano-brasileiro-de-ia-tera-supercomputador-e-investimento-de-r-23-bilhoes-em-quatro-anos>
21. Ministry of Electronics and Information Technology (MeitY), Government of India (2024). **IndiaAI Mission.**
22. Lang, N., Langione, M., Iyer, S., Das, A., & Zuluaga Martínez, D. (2026). **For Most Countries, AI Sovereignty Is an Illusion. Resilience Is Real.** Boston Consulting Group.
<https://www.bcg.com/publications/2026/ai-sovereignty-is-an-illusion-resilience-is-real>

23. Singh, V. V. (2025). Sovereign AI: Pathways to strategic autonomy. The International Institute for Strategic Studies (IISS).
<https://www.iiss.org/online-analysis/charting-cyberspace/2025/08/sovereign-ai-pathways-to-strategic-autonomy/>

LLM Usage Statement

We used Claude Opus 4.8 for “vibe-editing”: checking language and logical pitfalls, and locating supporting literature we loosely mentioned. The final version of the edition was verified by us, the authors – the argument, framework, and conclusions are our own.