Available on: https://jefrr.ep-journals.org/ E-: 3051-3650, P-ISSN: 3051-3642

Journal of Economic Finance Research and Review

Volume: 02 Issue: 06 June 2026

Toward Layered AI Autonomy: A Strategic Framework for Managing Dependency Risks in Transition Economies A Conceptual Framework with Illustrative Evidence from Vietnam

Pham Duy Minh

Faculty of Development Economics, VNU University of Economics and Business, Vietnam National University, Hanoi, Vietnam

Published Date: June 29, 2026 **Article DOI: 10.65150/EP-jefrr/V2E6/2026-07**

ABSTRACT: The rapid proliferation of artificial intelligence (AI) platforms from a small number of advanced-economy technology firms has created a structural dependency problem for developing nations. Existing concepts of AI sovereignty and digital sovereignty, primarily developed in the context of advanced economies, provide insufficient guidance for resource-constrained transition economies that must simultaneously navigate the imperative of AI adoption and the risks of technological lock-in. This paper proposes the Layered AI Autonomy for Transition Economies (LAAT) framework, a three-layer strategic model that enables transition economies to manage AI dependency risks across four dimensions: financial, technical, legal-contractual, and geopolitical. Drawing on switching costs theory (Shapiro & Varian, 1999), market thickness theory (Roth, 2012), platform economics (Parker et al., 2016), and techno-economic paradigm cycles (Perez, 2002), the framework reconceptualizes AI autonomy not as a binary state but as a multidimensional spectrum achievable through strategic layering. We apply the LAAT framework to Vietnam, a rapidly digitalizing lower-middle-income economy with an ambitious national AI strategy, demonstrating how the framework can guide evidence-based AI policy. The analysis reveals that Vietnam exhibits high dependency vulnerability across all four dimensions, with concentrated risk in the legal-contractual and geopolitical dimensions. The June 12, 2026 export control action against Anthropic's Fable 5 and Mythos 5 models serves as a critical empirical case for the framework's applied significance. We conclude with a three-phase policy roadmap translatable to other transition economies at comparable stages of AI adoption.

KEYWORDS: AI sovereignty, vendor lock-in, transition economies, digital policy, Vietnam, AI governance, platform dependency

1. INTRODUCTION

On June 12, 2026, the United States Bureau of Industry and Security (BIS) issued a directive requiring Anthropic - a leading AI safety company headquartered in San Francisco - to immediately restrict global access to its Fable 5 and Mythos 5 large language models on national security grounds (Koren et al., 2026). The practical consequence was stark: all of Anthropic's international customers lost access to these systems within hours, regardless of their location, industry, or operational criticality. For enterprises that had integrated Anthropic's models into core operations - from customer service automation to financial credit scoring - the disruption was immediate and involuntary. Affected organizations outside the United States had limited immediate operational recourse, even where subsequent legal challenges were possible; no contractual mechanism could override the sovereign export control directive, and no legally mandated transition period was provided. The incident instantiated, with unprecedented clarity, what scholars of platform economics have termed the 'Hotel California' problem (Parker et al., 2016): organizations that integrate deeply into a dominant AI platform cannot exit at will.

For most of the world's nations, this incident highlighted a structural vulnerability that extends well beyond any single AI provider. Governments and firms in developing and transition economies face an acute strategic dilemma. AI adoption is economically imperative: McKinsey Global Institute estimates that AI could add USD 13 trillion to global economic output by 2030, with late adopters facing compounding productivity gaps (Bughin et al., 2018). Yet each integration of foreign AI systems deepens dependencies that may prove difficult, costly, or politically untenable to reverse. Unlike enterprise software or mobile platforms, frontier AI systems create deep organizational interdependencies - they embed into decision-making processes, generate proprietary training data loops that do not transfer to competing platforms, and operate through service architectures that vest unprecedented operational control in the provider.

Despite the urgency of this challenge for the majority of the world's nations, the academic literature offers limited guidance specifically calibrated for transition economies. The concept of AI sovereignty - broadly defined as a state's capacity to independently develop, deploy, and govern frontier AI systems - has gained substantial traction in policy discourse (Bradford, 2023; Csernatoni, 2024; Floridi, 2020). However, full AI sovereignty is effectively unattainable for most transition economies within any strategically relevant timeframe. Training a frontier large language model currently requires computational resources valued at hundreds of millions of dollars, datasets of extraordinary scale and diversity, and specialized engineering talent concentrated in a handful of global research centers (Sevilla et al., 2022). National AI strategies from Singapore

License: This is an open access article under the CC BY 4.0 license: https://creativecommons.org/licenses/by/4.0/

(NAIS 2.0, 2023) and India (IndiaAI Mission, 2024) offer instructive regional precedents, but neither is directly applicable to lower-resource institutional contexts.

This paper addresses this gap through three contributions. First, we develop the Layered AI Autonomy for Transition Economies (LAAT) framework, which reconceptualizes AI autonomy as a spectrum rather than a binary state and enables transition economies to identify strategically appropriate levels of dependency management at each layer of the AI value chain. Second, we propose a four-dimension AI dependency risk model that enables systematic assessment of organizational and national vulnerability across financial, technical, legal-contractual, and geopolitical dimensions. Third, we apply the LAAT framework to Vietnam - a lower-middle-income economy exhibiting some of the fastest AI adoption rates in Southeast Asia, significant data center investment momentum, and an incipient but strategically ambitious national AI ecosystem - to demonstrate the framework's analytical utility and generate concrete policy recommendations translatable to comparable contexts.

The paper proceeds as follows. Section 2 reviews the literature on AI sovereignty, vendor lock-in, platform economics, and technology dependency in developing countries, identifying the gap that motivates the LAAT framework. Section 3 develops the framework and its theoretical foundations. Section 4 presents Vietnam's AI ecosystem as an empirical context. Section 5 applies the framework to assess Vietnam's dependency risk profile. Section 6 derives policy implications structured as a three-phase implementation roadmap. Section 7 concludes with limitations and directions for future research.

This paper employs a theory-building approach to conceptual research (Jaakkola, 2020). The LAAT framework is constructed through theoretically-driven synthesis of four established bodies of literature - switching costs theory, market thickness theory, platform economics, and technology dependency studies - rather than through primary empirical data collection. Following Gregor's (2006) typology of theory in information systems research, the framework belongs to the category of "theory for explaining and predicting": it offers both descriptive propositions about the structure of AI dependency risk and prescriptive propositions about strategic responses available to transition economies. The Vietnam application in Sections 4 and 5 constitutes illustrative case analysis, not empirical case study research: its purpose is to demonstrate the framework's analytical utility and to generate hypotheses for future empirical validation, rather than to make generalizable empirical claims from primary fieldwork. The three organizational cases - MoMo, GreenMind, and Masan Group - were selected through theoretical sampling to illuminate distinct positions within the LAAT framework, not to constitute a representative sample of Vietnamese AI adopters. All data sources are secondary: publicly available company communications, government policy documents, institutional reports, and industry analyses. The limitations of this approach are discussed further in Section 7.

2. LITERATURE REVIEW

2.1 AI Sovereignty and Digital Sovereignty: Conceptual Foundations

The concept of digital sovereignty - broadly defined as a state's capacity to govern its digital infrastructure, data flows, and technology systems - has become a central organizing principle in technology policy discourse (Pohle & Thiel, 2020; Couture & Toupin, 2019). In the AI domain, AI sovereignty has been used to describe the aspiration for independent national or regional capabilities in AI research, development, and deployment (Floridi, 2020; Csernatoni, 2024). Bradford (2023) offers the most comprehensive comparative treatment, framing global digital governance as a competition among three regulatory philosophies: American techno-liberalism (market-driven, permissive), European techno-regulation (rights-based, precautionary), and Chinese techno-nationalism (state-centric, developmentalist). Notably, the framework Bradford develops positions developing nations as largely absent from the conceptual architecture - subject to the regulatory externalities of the three major powers rather than as independent strategic actors.

The empirical record suggests that full AI sovereignty - defined as the capacity to independently develop, train, and operate frontier AI systems - is effectively unattainable for the vast majority of nations within any strategically relevant timeframe. The economics are structurally prohibitive: training frontier-class large language models currently requires computational investments of hundreds of millions of dollars per run (Sevilla et al., 2022), and the ongoing inference costs of operating such models at commercial scale are equally daunting. Industry estimates indicate that OpenAI's inference expenditures exceeded USD 8.4 billion in 2025, with gross margins constrained below 40% despite substantial revenue growth (FutureSearch, 2026). These economics suggest that even wealthy states face structural barriers to meaningful AI sovereignty at the frontier model level, making the aspirational frameworks of digital sovereignty insufficient as operational policy guides for resource-constrained nations.

2.2 Vendor Lock-In and Switching Costs in AI Markets

Shapiro and Varian (1999) identify switching costs - the economic, psychological, and operational costs of transitioning between incompatible technology systems - as a fundamental mechanism of competitive advantage in information markets. They distinguish among contractual lock-in, learning costs, data compatibility costs, and search costs as distinct switching cost categories that can simultaneously accumulate to create strong barriers to platform exit. In AI markets, these mechanisms are amplified by additional dynamics specific to the technology's architecture.

Contemporary analysis of AI-specific vendor lock-in identifies at least four distinct mechanisms beyond those observed in conventional software markets: API dependency (operational workflows built around proprietary function signatures that do not map to alternative platforms), agent framework capture (multi-step AI workflows embedded in provider-specific orchestration tools), data gravity (proprietary training and fine-tuning datasets that accumulate in formats and repositories tied to specific providers), and ecosystem entanglement (complementary tooling, developer communities, and support infrastructure concentrated within a single provider's platform) (CloudZero, 2026). A 2026 survey of 542 U.S. enterprise executives found that 81% expressed concern about their organization's dependency on specific AI vendors - with 29% reporting they are very concerned - while simultaneously deepening their integration with primary AI providers (Flanigan, 2026), a dynamic that Shapiro and Varian (1999) would recognize as rational given near-term productivity gains relative to uncertain future switching costs, but that creates compounding strategic exposure over time.

License: This is an open access article under the CC BY 4.0 license: https://creativecommons.org/licenses/by/4.0/

2.3 Platform Economics and Market Thickness

Roth (2012), drawing on market design theory, identifies market thickness - the participation of sufficient numbers of buyers and sellers on both sides of a platform - as a necessary precondition for platform viability and a source of sustainable competitive advantage. AI model markets exhibit an extreme version of this dynamic: the capital requirements for training frontier AI systems create natural monopoly tendencies (high fixed costs, low marginal costs) that generate winner-take-most outcomes at the foundation model layer (Rochet & Tirole, 2003). Parker, Van Alstyne, and Choudary (2016) extend this analysis, demonstrating that platform businesses systematically extract value through network effects that become self-reinforcing once critical mass is achieved, creating what they term 'pipeline-to-platform' transitions that fundamentally alter competitive dynamics in affected industries.

For developing economies seeking to build competitive AI capabilities, the implications of these dynamics are particularly constraining. The theoretical capacity to build competing AI systems does not translate into a viable market strategy without access to the datasets, compute infrastructure, and distribution networks that establish market thickness. This observation motivates a key analytical move of the LAAT framework: strategic AI autonomy for transition economies cannot be achieved through frontier model competition, but must be pursued at strategically differentiated points in the AI value chain where the resource constraints of transition economies are less binding and where their unique assets - particularly local language data and domain-specific institutional knowledge - confer genuine comparative advantages.

2.4 Technology and Developing Countries: Dependency Dynamics

A substantial interdisciplinary literature documents the ways in which technology transfer and adoption can reinforce rather than reduce international development asymmetries (Avgerou, 2002; Heeks, 2010; Kleine & Unwin, 2009). Dependency theory perspectives applied to digital contexts argue that developing nations' integration into global technology supply chains tends to consolidate core-periphery relationships in digital domains, as technology architectures embed the interests and institutional frameworks of their originating societies (Avgerou, 2002; Couldry & Mejias, 2019).

More recently, scholars have examined how the architecture of AI systems embeds geopolitical power asymmetries. Roberts et al. (2021) document how AI governance frameworks proposed by major powers reflect their domestic regulatory traditions and strategic interests, with developing nations disproportionately subject to externally determined technical norms and ethical frameworks. Couldry and Mejias (2019) introduce the concept of 'data colonialism' to describe the systematic extraction of value from developing nations through digital data flows - a dynamic that AI systems both reflect and amplify, as foundation models trained primarily on English-language and Western-institutional data systematically encode cultural and epistemic biases that disadvantage non-Western adopters.

2.5 National AI Strategies: Comparative Precedents for Middle-Income Economies

Three strategic precedents are particularly relevant for transition economies navigating AI dependency.

Singapore's National AI Strategy 2.0 (NAIS 2.0), released in December 2023, commits SGD 1 billion (approximately USD 743 million) to AI development across nine priority sectors (Smart Nation and Digital Government Office, 2023). Singapore's approach notably avoids the language of sovereignty, framing its objective instead as achieving 'strategic advantage over the AI value chain' - an implicit acknowledgment that a city-state of approximately 6 million cannot realistically pursue full AI self-sufficiency, but can position itself as a node in global AI infrastructure. This 'value chain positioning' strategy is instructive for similarly situated economies.

India's IndiaAI Mission (2024), approved with a budget of Rs 10,371 crore (approximately USD 1.24 billion), articulates a seven-pillar approach to building domestic AI capability while maintaining strategic openness to international partnerships (Ministry of Electronics and Information Technology, 2024). India's National AI Stack - a seven-layer architecture from compute infrastructure to application interfaces - represents the most architecturally detailed national framework for managing AI dependency currently in operation. Crucially, India frames its strategic position as mediating between 'American techno-liberalism and Chinese techno-authoritarianism,' an explicit acknowledgment of the geopolitical dimensions of AI dependency.

The European Union's approach, formalized through the AI Act (2024) and the Cloud and AI Development Act (CADA) included in the June 2026 Tech Sovereignty Package, establishes four levels of sovereignty assurance for critical public sector AI workloads (Inside Global Tech, 2026). While the EU's regulatory capacity and institutional framework are not directly replicable in transition economy contexts, its analytical distinction between different levels of acceptable dependency - from full sovereignty assurance to open-market dependency - provides conceptual scaffolding that the LAAT framework adapts and extends.

2.6 The Gap: A Framework for Transition Economies

The reviewed literature reveals a significant conceptual gap. AI sovereignty frameworks are well-developed for advanced economies with robust public R&D capacity and regulatory institutional infrastructure. Technology dependency dynamics are well-documented for developing nations across multiple technology waves. What is absent is a framework specifically calibrated for the intermediate case of transition economies: countries that possess some but not all of the prerequisites for strategic AI autonomy, that face acute dependency risks given their rapidly deepening AI adoption, and that require pragmatic, resource-constrained policy guidance rather than aspirational sovereignty frameworks designed for different institutional contexts. The LAAT framework is designed to address this gap.

3. THEORETICAL FRAMEWORK: LAYERED AI AUTONOMY FOR TRANSITION ECONOMIES (LAAT)

3.1 Conceptual Logic and Foundations

The LAAT framework rests on a fundamental reframing of the AI autonomy question. Rather than asking whether a country can achieve AI sovereignty - a question that elicits a practically negative answer for most transition economies - the framework asks: at which points in the AI value chain does dependency create unacceptable strategic risk, and at which points does dependency remain strategically acceptable given resource constraints and opportunity costs? This reframing enables policy analysis that is simultaneously principled (grounded in risk assessment) and pragmatic (calibrated to actual resource availabilities).

License: This is an open access article under the CC BY 4.0 license: https://creativecommons.org/licenses/by/4.0/

The framework draws on three theoretical foundations. First, portfolio theory (Markowitz, 1952) suggests that risk management through strategic diversification is more practically attainable than risk elimination - an insight that motivates the framework's emphasis on managed dependency rather than the unattainable goal of dependency elimination. Second, the modularity literature in technology strategy (Baldwin & Clark, 2000) demonstrates that complex technology systems can be decomposed into layers with different competitive dynamics, allowing actors to achieve dominant positions in selected layers rather than across the entire system - motivating the LAAT framework's layered architecture. Third, switching costs theory (Shapiro & Varian, 1999) provides the analytical vocabulary for understanding how dependency accumulates and how it can be strategically limited through architectural choices.

3.2 The Three Layers of AI Autonomy

The LAAT framework identifies three analytically distinct layers of AI activity, each with different dependency dynamics, resource requirements, and strategic implications for transition economies.

Layer 1: Application and Integration (Use Layer). This layer encompasses the deployment of AI systems for specific organizational or sectoral purposes: embedding AI into products, services, workflows, and decision-making processes through API calls, pre-built model interfaces, and off-the-shelf AI tools. At this layer, virtually all organizations in all countries accept some level of dependency on external AI capabilities. The strategic question is not whether to accept dependency but how to structure it to preserve switching optionality - the capacity to transition to alternative providers if necessary. Key instruments include multi-vendor architectures, middleware abstraction layers, and contractual provisions ensuring data portability.

Layer 2: Adaptation and Customization (Adapt Layer). This layer encompasses fine-tuning, retrieval-augmented generation, prompt engineering, and other techniques that customize pre-existing foundation models for local contexts, languages, or domain requirements. For many transition economies, this layer represents the most strategically significant near-term investment opportunity. Building local adaptation capabilities creates portable assets - training datasets, fine-tuning pipelines, domain-specific evaluation frameworks - that retain strategic value independent of the underlying foundation model. Critically, Layer 2 assets are substantially more transferable across foundation model providers than Layer 1 integrations, creating meaningful switching optionality. This optionality must be qualified, however: Layer 2 assets retain a structural dependency on their upstream base model, and the progressive release cycles of foundation model architectures impose recurring compatibility costs as organizations must re-validate and, in many cases, re-fine-tune their models against successive upstream versions. The LAAT framework designates this as Layer 2 maintenance dependency - a form of technical lock-in that is less acute than Layer 1 contractual dependency but non-trivial in aggregate, particularly for organizations operating on constrained engineering budgets.

Layer 3: Foundation Model Development (Build Layer). This layer encompasses the development, training, and deployment of original foundation models. For virtually all transition economies, full Layer 3 capability at the frontier level is not achievable within a 5-10 year strategic horizon given current resource constraints. However, this does not preclude strategic participation: transition economies can develop specialized foundation models for domains where their unique data assets confer comparative advantages (local-language models, sector-specific models for healthcare or agriculture), contribute to open-source foundation model ecosystems, and develop the talent pipelines and compute infrastructure that enable future Layer 3 expansion as the cost frontier of model training declines.

3.3 The Four-Dimension Dependency Risk Model

Across all three layers, dependency creates risk that must be actively managed. The LAAT framework identifies four analytically distinct dimensions of AI dependency risk, each requiring different diagnostic indicators and different management strategies.

Financial Dependency Risk arises from the cost structure of AI adoption. The economics of frontier AI provision - dominated by inference costs that scale continuously with usage - create a dependency dynamic in which deep AI integration generates growing financial exposure to provider pricing decisions. Organizations and governments that integrate AI deeply into operational processes face increasing marginal costs of switching, even where technically feasible alternatives exist, because switching costs compound with the volume and depth of AI usage. This is analogous to but more severe than software licensing dependency, because AI inference costs scale continuously rather than being capped at subscription thresholds. For developing economies where AI API costs represent a significant proportion of technology budgets, this risk dimension can create acute vulnerability to provider pricing changes.

Technical Dependency Risk encompasses the engineering lock-in created by deep integration with specific AI platforms, including API specificity, proprietary function calling conventions, model-specific fine-tuning artifacts, and agent framework dependencies. Unlike financial dependency, technical dependency can theoretically be managed through deliberate architectural choices - abstraction layers, standardized interfaces, modular workflow design - but requires upfront engineering investment that many organizations in resource-constrained contexts have not made. The cost of architectural robustness is typically invisible until a disruption forces it into visibility, at which point the switching costs have already accumulated substantially.

Legal-Contractual Dependency Risk refers to the ways in which provider terms of service, export control regulations, and geopolitical events beyond the dependent organization's control can instantaneously terminate service access. The June 12, 2026 Anthropic incident is the most consequential empirical instantiation of this risk dimension to date: the BIS directive, grounded in the Export Control Reform Act (2018) and Export Administration Regulations, required Anthropic to implement a global access cutoff with no operational recourse available to dependent customers (Koren et al., 2026). This risk dimension is qualitatively different from conventional commercial contract risk because it operates through sovereign legal authority rather than contractual mechanisms, and because it can be exercised without notice, compensation, or transition periods.

Geopolitical Dependency Risk encompasses the broader strategic implications of AI dependency for national security, technological autonomy, and positioning between competing geopolitical blocs. As AI systems become increasingly integrated into critical infrastructure, governmental decision-making, and national security functions, the dependency of a nation's AI ecosystem on foreign providers creates potential vectors for foreign influence, data exfiltration, and service denial that transcend commercial considerations. For countries navigating complex

License: This is an open access article under the CC BY 4.0 license: https://creativecommons.org/licenses/by/4.0/

relationships with multiple major powers - as is the case for many Southeast and South Asian nations - the geopolitical dimension of AI dependency presents particularly complex management challenges.

3.4 Strategic Propositions

From the LAAT framework, three propositions inform the policy analysis and recommendations that follow.

Proposition 1 (Layered Adequacy): Transition economies achieve optimal AI strategic positioning by targeting investment in Layers 1-2 autonomy while selectively building Layer 3 capacity in domains of genuine comparative advantage (particularly local-language and domain-specific models), rather than attempting uniform capability building across all layers. This proposition implies that resource allocation decisions should be differentiated by layer, with the highest investment priority assigned to the layer where the cost-effectiveness of autonomy is greatest given the country's specific asset base.

Proposition 2 (Dimensional Differentiation): Organizations and governments should differentiate their dependency management strategies according to the specific risk dimension creating the greatest exposure, recognizing that effective interventions for financial dependency (multi-vendor contracts, usage caps, open-source alternatives) differ fundamentally from those for legal-contractual dependency (contingency protocols, domestic hosting requirements) and geopolitical dependency (multilateral engagement, non-aligned provider diversification).

Proposition 3 (Dynamic Adequacy): The optimal level of autonomy at each layer changes over time as the technology landscape evolves, requiring periodic strategic reassessment rather than a fixed framework endpoint. In particular, the commoditization of Layer 3 capabilities through open-source model releases - an ongoing trend observable in the 2023-2026 period - creates layer transition opportunities that enable transition economies to upgrade their strategic position incrementally without requiring step-change resource investments.

4. VIETNAM AS EMPIRICAL CONTEXT

4.1 AI Adoption Landscape

Vietnam presents an instructive empirical context for the LAAT framework application. As a lower-middle-income economy with approximately 98 million inhabitants, a rapidly growing digital economy, and an ambitious national AI strategy, Vietnam exemplifies the transition economy dilemma: sufficiently advanced to have significant AI adoption and AI-related national interests, but insufficiently resourced to pursue full technological sovereignty at the frontier.

Vietnam's AI adoption has accelerated substantially in recent years. According to the UNDP's AI Landscape and Assessment Report (2025), approximately 18% of Vietnam's approximately 170,000 registered enterprises had adopted AI systems in operational capacities by 2024, representing a 39% year-on-year increase. Adoption is sectorally concentrated: information technology firms account for 31% of AI-active enterprises, followed by financial services (22%) and education (17%), with growing activity in retail, logistics, and manufacturing (UNDP, 2025). Critically, the UNDP report notes that 55% of AI-active firms identify talent shortages as the primary constraint on deeper adoption - a structural gap that shapes the type of AI dependency these organizations typically incur (reliance on external expertise alongside reliance on external platforms).

Vietnam’s national AI governance framework is anchored by Decision No. 127/QD-TTg (2021), which sets targets for Vietnam to rank among the top 5 AI nations in ASEAN and top 60 globally by 2025, and top 4 in ASEAN and top 50 globally by 2030, alongside goals for at least three national AI data center facilities by 2030 (Government of Vietnam, 2021). The strategy was designed before the generative AI diffusion wave of 2022-2026. Vietnam’s first dedicated AI law - the Law on Artificial Intelligence No. 134/2025/QH15, adopted by the National Assembly on 10 December 2025 and in force from 1 March 2026 - has since been enacted, providing the statutory governance architecture the strategy anticipated (Tilleke & Gibbins, 2025). The gap between the strategy's aspirational architecture and the operational realities of Vietnam's AI ecosystem is a central theme of the policy analysis that follows.

4.2 Infrastructure: Investment Momentum and Present Constraints

Vietnam's data center and AI infrastructure landscape is in rapid transition. The country has attracted more than USD 7 billion in announced AI data center investments as of mid-2026, including a planned 200-megawatt campus in Kinh Bac City designed to host approximately 100,000 GPU units, and significant expansions by major state-owned telecommunications operators (VietnamPlus, 2026). NVIDIA has established a research and development presence in Vietnam, and the regulatory framework was amended in 2024 to permit full foreign ownership of data center facilities, removing a significant barrier to international infrastructure investment (NVIDIA, 2025).

Despite this momentum, the gap between announced infrastructure plans and operational AI compute capacity is substantial. Vietnamese organizations currently rely overwhelmingly on foreign cloud AI services for their AI compute requirements - a dependency that persists as domestic infrastructure develops through a necessarily time-delayed buildout cycle. This gap is characteristic of transition economies at early stages of AI infrastructure development: the political and economic will for strategic AI independence precedes the technical and financial means of achieving it, creating a vulnerability window during which organizations deepen their platform dependencies before domestic alternatives are available.

4.3 Three Cases: Navigating the LAAT Layers

Three Vietnamese AI cases illustrate how organizations within the country's ecosystem are navigating - often implicitly and without explicit reference to an autonomy framework - the strategic choices that the LAAT model makes analytically explicit.

MoMo, Vietnam’s leading financial super-app with over 40 million registered users and 1.6 billion annual transactions (Amy, 2025), represents an advanced example of Layer 1 implementation with significant Layer 2 characteristics. MoMo's AI applications - including credit scoring, automated customer service through its Moni AI assistant, electronic Know Your Customer (eKYC) automation, and real-time behavioral analytics - are deeply integrated across its core product suite. While MoMo's foundation model infrastructure relies substantially on international providers, its proprietary training data (derived from 1.6 billion annual transactions spanning Vietnam's distinct financial behavior patterns) and domain-specific fine-tuning represent non-replicable Layer 2 assets. MoMo's credit scoring models, trained on Vietnamese consumer financial

License: This is an open access article under the CC BY 4.0 license: https://creativecommons.org/licenses/by/4.0/

data, would retain independent strategic value even in a scenario of forced platform transition - an implicit form of switching optionality that illustrates how Layer 2 investment can mitigate Layer 1 dependency without requiring Layer 3 capability.

GreenMind, developed by GreenNode (a VNG-affiliated company) and introduced at NVIDIA AI Day in Ho Chi Minh City in 2025, represents Vietnam's most advanced current example of Layer 2 autonomy (Thai Khang, 2025). GreenMind-Medium-14B-R1 is a Vietnamese large language model with 14.7 billion parameters, fine-tuned from the open-source Qwen2.5-14B-Instruct architecture on a corpus of more than one million curated Vietnamese-language instruction-response pairs, and deployed as the first Vietnamese open-source LLM optimized for NVIDIA's NIM inference framework (GreenNode, 2025). Because GreenMind is built through fine-tuning of an existing open-source foundation model rather than through independent pre-training of a new architecture, it is appropriately classified at Layer 2; attribution of Layer 3 characteristics would require demonstrated independent foundation model capability, which has not been established for this release. Its technical architecture illustrates LAAT Proposition 3 in practice: by building on an open-source foundation model rather than proprietary API access, GreenNode has created a Layer 2 asset that is substantially independent of foreign provider licensing decisions and that can be operated on domestically or regionally owned infrastructure. The open-source release of GreenMind further creates public goods that reduce the Layer 2 entry costs for other Vietnamese organizations.

The GreenMind case illustrates Layer 2's strategic advantages with precision, but it equally exposes a structural limitation that the framework must acknowledge explicitly: fine-tuned models carry what the machine learning engineering literature terms hidden technical debt (Sculley et al., 2015) - accumulated compatibility obligations that materialize each time the upstream base model architecture is updated. GreenMind-Medium-14B-R1 is built on Qwen2.5-14B-Instruct; when Alibaba releases a successor architecture, as has occurred across every major open-source model family on roughly 6-12 month release cycles, GreenNode faces a decision among three costly paths: (a) re-fine-tuning on the new base at full compute cost to access capability improvements; (b) maintaining the existing model without upstream updates, accepting progressive capability divergence from frontier performance; or (c) transitioning to a different base architecture entirely, which invalidates the compatibility of existing fine-tuning pipelines, evaluation benchmarks, and deployment infrastructure. None of these options is costless, and for organizations with constrained AI engineering capacity - the typical condition across a transition economy's enterprise population - the maintenance cycle of Layer 2 can consume the engineering bandwidth that would otherwise support broader capability development. The switching optionality attributed to Layer 2 assets is real but asymmetric: training datasets and fine-tuning methodology transfer across base model transitions; the fine-tuned model artifact itself - the primary operational asset - does not. Layer 2 investment strategies must therefore budget explicitly for recurring maintenance costs, not only initial fine-tuning capital, and should prioritize fine-tuning infrastructure and evaluation tooling - the genuinely portable components - over model artifacts, which are architecture-perishable. Layer 2 development is correctly identified as more strategic than Layer 1 for transition economies, but it is better understood as risk diversification - specifically, mitigation of the legal-contractual and geopolitical dimensions of dependency - than as a stable endpoint that eliminates upstream technical dependency. The residual maintenance dependency on base model release cycles is the structural cost of Layer 2's relative autonomy, and this cost must be made analytically explicit in any policy framework that recommends Layer 2 investment as a near-term strategic priority.

Masan Group's WinCommerce division, operating Vietnam's largest modern retail network, illustrates Layer 1 deployment with attention to operational continuity risk. The WiNARE AI-powered ordering system and Digital Twin supply chain modeling tools, deployed across WinMart and WinMart+ retail locations, have achieved measurable efficiency gains, including a 15% reduction in staff operational workload as reported by the company itself (Masan Group, 2025). Masan's AI integration pattern is concentrated in operational optimization applications where switching costs, while present, are less severe than in customer-facing or core decision-making AI applications. This portfolio characteristic implicitly reduces its legal-contractual and technical dependency exposure without requiring explicit autonomy strategy - though it also represents a ceiling on the competitive differentiation Masan can derive from AI, since operational optimization tools are more easily replicated by competitors than proprietary model assets.

5. APPLYING LAAT: VIETNAM'S AI DEPENDENCY RISK PROFILE

5.1 Financial Dependency Risk: High and Growing

Vietnam's AI-active enterprises are primarily dependent on API-based consumption of foreign AI services, exposing them to the pricing dynamics of a market characterized by high provider concentration and accelerating compute costs. The structural economics of frontier AI provision illustrate why this risk is not transient: Industry estimates project OpenAI's cash consumption at approximately USD 27 billion in 2026, sustained by venture capital and strategic investment rather than operational profitability, suggesting that current API pricing may reflect a subsidized phase (FutureSearch, 2026). As major AI providers eventually seek to reduce capital subsidies and achieve sustainable gross margins, organizations in price-sensitive markets like Vietnam face significant financial exposure to repricing events.

For Vietnamese small and medium-sized enterprises - which constitute the overwhelming majority of the country's enterprise population - this risk is particularly acute. Lacking the scale to negotiate volume discounts or enterprise contracts, these organizations consume AI services at retail pricing and may have integrated AI deeply into operational workflows before the full cost trajectory of their dependency becomes visible. The assessment for this dimension is HIGH, with an upward trajectory as adoption deepens.

5.2 Technical Dependency Risk: Medium-High with Emerging Mitigation

Technical dependency risk is moderate in aggregate but highly variable across Vietnam's AI landscape. Organizations that have developed Layer 2 assets - proprietary training datasets, fine-tuned models, domain-specific evaluation frameworks - possess meaningful switching options at the technical level. Those that have integrated AI purely through direct API calls to proprietary models, without architectural abstraction layers, face substantial switching costs.

License: This is an open access article under the CC BY 4.0 license: https://creativecommons.org/licenses/by/4.0/

The emergence of GreenMind as an open-source Vietnamese foundation model meaningfully reduces technical dependency risk for organizations operating in Vietnamese-language contexts. Specifically, it demonstrates the viability of Layer 2 development for Vietnamese applications and provides a non-proprietary alternative for Vietnamese-language tasks that reduces the switching cost of departing from foreign foundation models. This is a material improvement in Vietnam's technical dependency risk profile, though the gap between GreenMind's capabilities and frontier foreign models remains significant for complex reasoning tasks. A further qualification is analytically necessary: GreenMind's risk mitigation value is not static. Open-source base model release cycles create recurring compatibility costs - re-validation, re-fine-tuning, and deployment re-integration - that constitute an ongoing technical liability. Organizations relying on GreenMind or comparable fine-tuned assets must maintain the engineering capacity to track and respond to upstream architectural changes; those that cannot will experience progressive capability degradation relative to organizations able to sustain the maintenance cycle. Layer 2 autonomy is thus best characterized as a risk profile that reduces legal-contractual and geopolitical exposure while sustaining a distinct form of technical dependency - one more tractable and less acute than platform API lock-in, but not absent. The assessment for this dimension is MEDIUM-HIGH, with a moderately improving trajectory contingent on sustained maintenance investment in Layer 2 assets.

5.3 Legal-Contractual Dependency Risk: Very High

The June 12, 2026 Anthropic incident provides the definitive empirical case for legal-contractual AI dependency risk, and its implications for Vietnam are severe. The incident demonstrated that U.S. export control authorities can compel immediate, global service termination without notice, without compensation, and without regard for the operational dependencies of affected organizations (Koren et al., 2026). Vietnamese organizations using Anthropic's models - including those in financial services, healthcare, and logistics where operational continuity is critical - had limited immediate operational recourse against this outcome, even where subsequent legal challenges may have been available; and the incident was the first application of the Export Control Reform Act (2018) to a broadly deployed AI software service.

This risk dimension is particularly significant for Vietnam given the country's strategic positioning between the United States and China in the ongoing technology competition. Vietnamese organizations using American AI services face exposure to U.S. export control authority; those using Chinese AI services face analogous exposure from the Chinese state and from potential U.S. secondary sanctions that could restrict access to Chinese AI platforms. Legal-contractual dependency risk is, in this context, fundamentally a function of geopolitical exposure rather than of commercial contractual relationships alone. The assessment for this dimension is VERY HIGH, reflecting the demonstrated willingness of major powers to use AI service control as a geopolitical instrument.

5.4 Geopolitical Dependency Risk: High

Vietnam's AI ecosystem reflects the country's broader 'bamboo diplomacy' strategic posture - maintaining relationships with multiple major powers while avoiding formal alignment. This posture creates distinctive challenges in the AI domain, where the major foundation model providers are concentrated in the United States (OpenAI, Anthropic, Google, Meta) and China (Baidu, Alibaba, Huawei), and where each provider relationship carries implicit geopolitical alignment implications. The expansion of U.S. AI export controls - first to China-destined AI hardware through the Bureau of Industry and Security and now, through the Anthropic incident, to AI software accessible from abroad - suggests a trajectory in which AI services may increasingly function as instruments of foreign policy rather than purely commercial products.

The development of open-source foundation models by non-U.S., non-Chinese actors - including AMI Labs in Paris (Patel, 2026) and various European and Japanese academic and commercial initiatives - represents a potential partial mitigation for geopolitical dependency risk, as it creates provider options outside the primary geopolitical competition. However, the capability gap between these models and U.S. frontier systems remains substantial in most domains, limiting their near-term utility as dependency alternatives. The assessment for this dimension is HIGH.

6. POLICY IMPLICATIONS: A THREE-PHASE FRAMEWORK FOR TRANSITION ECONOMIES

The LAAT framework analysis generates a three-phase policy roadmap applicable to transition economies broadly. The roadmap is organized around the framework's central distinction between dependency that is strategically acceptable and dependency that poses systemic risk. Each phase addresses a different temporal horizon and a different combination of dependency risk dimensions. While the propositions are drawn from the Vietnam analysis, they are generalizable to other developing countries at comparable stages of AI adoption - particularly those in Southeast Asia, South Asia, Sub-Saharan Africa, and Latin America that share the structural characteristics of rapid AI adoption, constrained public resources, and limited Layer 3 infrastructure.

6.1 Phase 1 (Years 1-2): Dependency Mapping and Contingency Institutionalization

The immediate priority across transition economies is systematic knowledge of existing dependency profiles. In most developing countries, no comprehensive official mapping of organizational AI dependencies - by provider, by layer, by risk dimension - currently exists. Governments should mandate AI dependency disclosure requirements for organizations in critical sectors (financial services, healthcare, telecommunications, energy), analogous to supply chain risk disclosure requirements now standard in manufacturing policy. This disclosure requirement serves a dual purpose: it generates the policy-relevant data absent from most developing-country AI governance frameworks, and it creates organizational awareness of dependency risks that most AI-adopting enterprises have not explicitly assessed.

The June 12, 2026 Anthropic incident provides a concrete and universally applicable policy rationale for the second key Phase 1 instrument: mandatory AI Operational Continuity Standards for critical sectors. These standards should specify minimum requirements for tested fallback capabilities (alternative providers, cached model deployments, manual process alternatives) in the event of sudden provider service termination. The incident demonstrated that export control authority can be exercised without notice, without compensation, and without transition periods - a scenario for which most organizations in developing countries are operationally unprepared. In Vietnam, the Law on Artificial Intelligence No. 134/2025/QH15 (in force from 1 March 2026) provides a direct statutory vehicle through which to incorporate such requirements; comparable opportunities exist in ongoing AI governance processes across ASEAN, the African Union, and South Asian regional bodies.

License: This is an open access article under the CC BY 4.0 license: https://creativecommons.org/licenses/by/4.0/

Phase 1 also requires the establishment of a national AI infrastructure baseline - a systematic inventory of compute capacity, network connectivity, and technical talent - to provide the factual foundation for strategic investment decisions in subsequent phases. Singapore's NAIS 2.0 explicitly incorporated such an infrastructure audit as a precondition for strategy formulation. For Vietnam specifically, this baseline should quantify the gap between announced data center investment plans (USD 7 billion in announced commitments as of mid-2026) and operational AI compute capacity currently available to domestic organizations, a gap that is likely substantial and that defines the near-term dependency management imperative.

6.2 Phase 2 (Years 2-5): Selective Capability Building and Public Goods Provision

The medium-term priority for transition economies is strategic investment in Layer 2 capability assets that reduce dependency vulnerability without requiring unrealistic commitments to frontier foundation model development. The LAAT framework's Proposition 1 (Layered Adequacy) implies that the highest return on AI public investment in developing country contexts typically lies at the Layer 2 level - adaptation, fine-tuning, domain customization, and local evaluation infrastructure - rather than at the frontier Layer 3 level where resource requirements are prohibitive.

Three investment priorities are broadly generalizable across developing country contexts. First, local-language and domain-specific AI infrastructure as a national public good: government-sponsored development of curated, high-quality language datasets, evaluation benchmarks, and fine-tuning tooling reduces the marginal cost of Layer 2 development for all domestic organizations. India's National AI Stack - which treats dataset development as national infrastructure analogous to roads or electricity grids - represents the most architecturally sophisticated example of this approach currently in operation. Vietnam's experience with GreenMind illustrates that commercially viable Layer 2 development is achievable with substantially smaller resources than frontier model training: GreenMind-Medium-14B-R1 was built by fine-tuning an existing open-source architecture on a curated Vietnamese corpus, a model that is replicable for other developing-country languages and domains.

Second, AI adaptation talent as a strategic priority: targeted investment in Layer 2 skills - data curation, fine-tuning engineering, domain-specific evaluation, and AI system auditing - generates widely distributed dependency-mitigation capacity across the enterprise ecosystem. This is a substantially higher-return public investment than equivalent spending on frontier AI research for economies where the talent pipeline and compute infrastructure for frontier work are not yet established. The skills required for effective Layer 2 AI work are learnable within shorter training cycles than frontier research skills, making them tractable targets for vocational and tertiary education programs within the Phase 2 timeframe.

Third, sector-specific model co-investment in domains of genuine comparative advantage: developing countries typically possess unique data assets in domains where foreign foundation models - trained predominantly on English-language and Western-institutional corpora - perform poorly. Local-language medical documentation, agricultural data from climate conditions not represented in global training datasets, legal and regulatory corpora in non-dominant languages, and governmental administrative records represent domains where locally trained specialized models would outperform foreign general-purpose alternatives. Government co-investment in such models creates Layer 2 assets with both commercial value and strategic dependency-mitigation properties.

6.3 Phase 3 (Years 5-10): Multilateral Engagement and Collective Architecture

The long-term priority for transition economies is participation in multilateral AI governance frameworks and regional infrastructure initiatives that convert bilateral dependency exposure into collective bargaining leverage. Individual developing countries have limited negotiating power with major AI providers; collectively, they represent a substantial and growing share of global AI adoption. Regional bodies - ASEAN, the African Union, CARICOM, SAARC - provide existing institutional frameworks through which collective AI governance standards and shared infrastructure investments can be developed.

The EU's CADA model offers a directly relevant template: by establishing formally recognized sovereignty assurance levels for AI workloads, it enables governments and organizations to specify dependency risk thresholds in procurement contracts and regulatory requirements, creating market pressure on providers to offer more resilient service architectures. A developing-country regional equivalent - whether at the ASEAN, African Union, or broader Global South level - would generate analogous market effects while providing developing nations with a multilateral governance platform through which to engage major AI-producing nations on equitable access and operational continuity guarantees.

For Vietnam specifically, its current strategic positioning - active diplomatic relationships with both the United States and China, growing ASEAN leadership role, and demonstrated AI ecosystem dynamism - makes it well-placed to be an early advocate and convener for such multilateral frameworks. Vietnam’s ongoing national AI strategy update - now proceeding alongside implementation of the Law on Artificial Intelligence No. 134/2025/QH15 (in force from 1 March 2026) - provides an opportunity to embed multilateral AI governance engagement as an explicit strategic objective rather than a peripheral diplomatic consideration. More broadly, the June 2026 Anthropic incident has created a window of heightened developing-country attention to AI dependency risks that multilateral governance initiatives can leverage; the political salience of the issue is unlikely to be higher in the near term.

7. CONCLUSION

This paper has developed and applied the Layered AI Autonomy for Transition Economies (LAAT) framework to address a significant gap in the AI strategy and governance literature. Existing frameworks for AI sovereignty provide insufficient guidance for the strategic realities facing transition economies: they are too demanding in their requirements for full technological self-sufficiency, and too poorly calibrated for the institutional, financial, and technical constraints of rapidly developing nations that must navigate AI dependency risks with limited resources.

The LAAT framework contributes three analytical instruments. A three-layer model of AI value chain activity - distinguishing Use, Adapt, and Build layers with different dependency dynamics and resource requirements - enables strategic differentiation of where transition economy investment generates the highest autonomy returns. A four-dimension dependency risk model - distinguishing Financial, Technical, Legal-Contractual, and Geopolitical risks - provides systematic diagnostic capability that moves beyond general concerns about AI dependency to

License: This is an open access article under the CC BY 4.0 license: https://creativecommons.org/licenses/by/4.0/

actionable risk-specific management. A three-phase policy roadmap - from dependency auditing through selective capability building to multilateral engagement - provides a sequenced implementation architecture calibrated to transition economy resource constraints.

Applied to Vietnam, the framework reveals a dependency risk profile that is high across all four dimensions, with particularly acute vulnerability in the legal-contractual dimension following the June 2026 Anthropic incident. Yet the analysis also identifies meaningful mitigation pathways available within Vietnam's current resource envelope, most notably through Layer 2 capability development along the lines demonstrated by GreenMind, and through governance interventions that create organizational awareness and contingency preparedness for legal-contractual disruption scenarios.

Several limitations of this analysis warrant acknowledgment. Most significantly, the framework has been developed without the benefit of primary empirical data from Vietnamese policymakers, AI practitioners, or enterprise technology leaders. The dependency risk assessments, while grounded in available secondary evidence, would benefit from qualitative empirical validation through expert interviews and structured case study research. Additionally, the technology landscape is evolving rapidly: the commoditization of AI capabilities through open-source model releases and the emergence of architecturally novel AI approaches - including the World Model paradigm being pursued by initiatives such as AMI Labs (Patel, 2026) - may shift the strategic significance of framework layers faster than policy implementation cycles can accommodate.

Future research should extend the LAAT framework through comparative application across Southeast Asian transition economies - particularly the Philippines, Indonesia, and Thailand - to assess the generalizability of its propositions across varied institutional contexts. Quantitative operationalization of the four dependency dimensions, enabling cross-national comparison and longitudinal tracking, would substantially increase the framework's empirical contribution. Longitudinal case study research tracking the evolution of GreenMind and comparable Layer 2 initiatives would generate valuable evidence on the conditions under which Layer 2 investment successfully mitigates Layer 1 and legal-contractual dependency risks in practice.

REFERENCES

1) Amy Adelaide (2025). Vietnam's fintech power play: MoMo's profit milestone signals a new era for Southeast Asia. 1950.ai. https://www.1950.ai/post/vietnam-s-fintech-power-play-momo-s-profit-milestone-signals-a-new-era-for-southeast-asia
2) Avgerou, C. (2002). Information systems and global diversity. Oxford University Press.
3) Baldwin, C. Y., & Clark, K. B. (2000). Design rules: The power of modularity. MIT Press.
4) Bradford, A. (2023). Digital empires: The global battle to regulate technology. Oxford University Press.
5) Bughin, J., Seong, J., Manyika, J., Chui, M., & Joshi, R. (2018). Notes from the AI frontier: Modeling the impact of AI on the global economy. McKinsey Global Institute.
6) CloudZero. (2026). AI vendor lock-in: How AI is creating a new dependency problem. CloudZero. https://www.cloudzero.com/blog/ai-vendor-lock-in/
7) Flanigan, R. (2026, April 1). Nearly 3 in 4 enterprises say losing AI vendors would disrupt core business operations. Zapier. https://zapier.com/blog/ai-vendor-lock-in-survey/
8) Couldry, N., & Mejias, U. A. (2019). The costs of connection: How data is colonizing human life and appropriating it for capitalism. Stanford University Press.
9) Couture, S., & Toupin, S. (2019). What does the notion of sovereignty mean when referring to the digital? New Media & Society, 21(10), 2305-2322. https://doi.org/10.1177/1461444819865984
10) Csernatoni, R. (2024). Charting the geopolitics and European governance of artificial intelligence. Carnegie Endowment for International Peace. https://carnegieendowment.org/research/2024/03/charting-the-geopolitics-and-european-governance-of-artificial-intelligence
11) Floridi, L. (2020). The fight for digital sovereignty: What it is, and why it matters, especially for the EU. Philosophy & Technology, 33(3), 369-378. https://doi.org/10.1007/s13347-020-00423-6
12) FutureSearch. (2026). OpenAI revenue, losses, and profitability in 2026: Full financial breakdown. FutureSearch AI. https://futuresearch.ai/openai-revenue-forecast/
13) Government of Vietnam. (2021). Quyet dinh so 127/QD-TTg ngay 26/01/2021 ban hanh Chien luoc quoc gia ve nghien cuu, phat trien va ung dung Tri tue nhan tao den nam 2030 [Decision No. 127/QD-TTg: National strategy for AI research, development and application to 2030]. Office of the Prime Minister.
14) GreenNode. (2025). GreenMind: A next-generation Vietnamese large language model for structured and logical reasoning. arXiv preprint. https://arxiv.org/pdf/2504.16832
15) Gregor, S. (2006). The nature of theory in information systems. MIS Quarterly, 30(3), 611-642. https://doi.org/10.2307/25148742
16) Heeks, R. (2010). Do information and communication technologies (ICTs) contribute to development? Journal of International Development, 22(5), 625-640. https://doi.org/10.1002/jid.1716
17) Inside Global Tech. (2026, June 4). EU tech sovereignty package. Inside Global Tech. https://www.insideglobaltech.com/2026/06/04/eu-tech-sovereignty-package/
18) Jaakkola, E. (2020). Designing conceptual articles: Four approaches. AMS Review, 10(1-2), 18-26. https://doi.org/10.1007/s13162-020-00161-0
19) Kleine, D., & Unwin, T. (2009). Technological revolution, evolution and new dependencies: What is new about ICT4D? Third World Quarterly, 30(5), 1045-1067. https://doi.org/10.1080/01436590902959339
20) Koren, K., Kurland, K., & Mehta, A. (2026, June 16). The Department of Commerce restricted access to Anthropic's latest models: What comes next? Center for Strategic and International Studies. https://www.csis.org/analysis/department-commerce-restricted-access-anthropics-latest-models-what-comes-next

License: This is an open access article under the CC BY 4.0 license: https://creativecommons.org/licenses/by/4.0/

21) Markowitz, H. (1952). Portfolio selection. The Journal of Finance, 7(1), 77-91. https://doi.org/10.2307/2975974
22) Masan Group. (2025, September 18). Technology applications optimize supply chain and reduce losses in modern retail. Masan Group News. https://www.masangroup.com/news/masan-news/Technology-Applications-Optimize-Supply-Chain-and-Reduce-Losses-in-Modern-Retail.html
23) Ministry of Electronics and Information Technology. (2024). IndiaAI Mission. Government of India. https://indiaai.gov.in/indiaai-mission
24) NVIDIA. (2025). Thailand, Vietnam embrace sovereign AI to drive economic growth. NVIDIA Blog. https://blogs.nvidia.com/blog/thailand-vietnam-sovereign-ai/
25) Parker, G. G., Van Alstyne, M. W., & Choudary, S. P. (2016). Platform revolution: How networked markets are transforming the economy and how to make them work for you. W. W. Norton & Company.
26) Patel, N. (2026, March 9). Yann LeCun's AMI Labs raises $1.03 billion to build world models. TechCrunch. https://techcrunch.com/2026/03/09/yann-lecuns-ami-labs-raises-1-03-billion-to-build-world-models/
27) Perez, C. (2002). Technological revolutions and financial capital: The dynamics of bubbles and golden ages. Edward Elgar Publishing.
28) Pohle, J., & Thiel, T. (2020). Digital sovereignty. Internet Policy Review, 9(4). https://doi.org/10.14763/2020.4.1532
29) Roberts, H., Cowls, J., Morley, J., Taddeo, M., Wang, V., & Floridi, L. (2021). The Chinese approach to artificial intelligence: An analysis of policy, ethics, and regulation. AI & Society, 36(1), 59-77. https://doi.org/10.1007/s00146-020-00992-2
30) Rochet, J.-C., & Tirole, J. (2003). Platform competition in two-sided markets. Journal of the European Economic Association, 1(4), 990-1029. https://doi.org/10.1162/154247603322493212
31) Roth, A. E. (2012). The theory and practice of market design. Nobel Prize lecture. Nobel Foundation. https://www.nobelprize.org/prizes/economic-sciences/2012/roth/lecture/
32) Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. Advances in Neural Information Processing Systems, 28. https://proceedings.neurips.cc/paper/2015/hash/86df7dcfd896fcaf2674f757a2463eba-Abstract.html
33) Sevilla, J., Heim, L., Ho, A., Besiroglu, T., Hobbhahn, M., & Villalobos, P. (2022). Compute trends across three eras of machine learning. arXiv preprint arXiv:2202.05924. https://arxiv.org/abs/2202.05924
34) Shapiro, C., & Varian, H. R. (1999). Information rules: A strategic guide to the network economy. Harvard Business School Press.
35) Smart Nation and Digital Government Office. (2023). National artificial intelligence strategy 2.0: AI for the public good, for Singapore and the world. Singapore Government. https://file.go.gov.sg/nais2023.pdf
36) Tilleke & Gibbins. (2025). Vietnam's AI push: Updated national strategy and first AI law by end of 2025. Tilleke & Gibbins Insights. https://www.tilleke.com/insights/vietnams-ai-push-updated-national-strategy-and-first-ai-law-by-end-of-2025/
37) UNDP. (2025). AI landscape and assessment report: Viet Nam 2025. United Nations Development Programme. https://www.undp.org/sites/g/files/zskgke326/files/2025-04/undp_aila_eng_report_2025_final.pdf
38) Thai Khang (2025). Vietnam unveils first Vietnamese LLM on NVIDIA NIM. VietnamNet. https://vietnamnet.vn/en/vietnam-unveils-first-vietnamese-llm-on-nvidia-nim-2445851.html
39) VietnamPlus. (2026). Vietnam emerges as new hotspot for AI data centre investment. Vietnam News Agency. https://en.vietnamplus.vn/vietnam-emerges-as-new-hotspot-for-ai-data-centre-investment-post332260.vnp

License: This is an open access article under the CC BY 4.0 license: https://creativecommons.org/licenses/by/4.0/

Corresponding Author: Pham Duy Minh