Strategic and Financial Requirements for a Sovereign AI Infrastructure

Gilles CHEMLA[a,b], Tahmid QUDDUS ISLAM[c] , and Simeng LI[a]
[a] *Imperial Business School, Centre for Financial Technology*
[b] *Harvard University, CNRS, and CEPR*
[c] *Consultant & Independent Advisor*

ORCiD ID: Gilles Chemla https://orcid.org/0000-0002-7369-9241,
Tahmid Quddus Islam https://orcid.org/0009-0000-9855-6418,
Simeng Li https://orcid.org/0009-0005-9966-4636

Abstract. Sovereign AI refers to the ability to develop, operate, secure, control, and sustain critical AI systems under geopolitical or operational stress - ensuring continuity and growth of AI services even if energy supplies are disrupted, global supply chains fragment, or cloud access is restricted. As artificial intelligence (AI) becomes integral to business, government, and society, questions on sovereign and defense AI infrastructure have moved to the forefront. This paper points out that achieving AI sovereignty requires sustainable control over key dependencies in the AI stack (such as energy, commodities, computer hardware, cloud platforms, data, algorithms, and information systems). To achieve this, countries need to own an increasingly integrated web of technology, innovation, supply chain, financial infrastructure, information systems, education, and to have adequate governance and political systems in place. We propose a framework of requirements for sovereign AI, outlining the essential measures needed to make AI capabilities resilient against external shocks. This perspective requires an in-depth transformation of economic systems to ensure technological sovereignty for NATO and its allies.

Keywords. Sovereign AI; technology; supply chain; commodities; governance; finance; information systems; politics.

1. Introduction

Sovereignty is multidimensional. It involves control over an increasingly integrated web of technology, innovation, talent, supply chain, finance, information systems, and governance and politics [1,2]. Sovereign economies need to own developed technological and financial infrastructures [3–6] and control their industries [7–9] while sustaining political sovereignty in a world where economic crises and multinational information systems can generate polarization [10–15] and be subject to foreign influence [16–19]. Countries must also be prepared to cope with the cost of potential conflicts both on their land and in neighboring countries [20–22].

Technology, and in particular artificial intelligence (AI) are fast transforming not only business [23], but also the challenges above [24–27]. This in turn poses new

questions on countries' ability to control their AI infrastructure. This paper examines the infrastructure needed to ensure the structural resilience of AI capabilities. We introduce the concept of sovereign AI as the ability to develop, operate, secure, control and sustain AI systems independently under geopolitical or operational stress.

While securing systems against traditional cybersecurity attacks and addressing AI ethical concerns are necessary, we point out that sovereignty is about ensuring those systems are developed, controlled, sustained and that they remain available and effective when external support is cut off or hostile pressure is applied.

Sovereign AI thus hinges on managing a stack of interdependent resources and technologies - from energy and hardware up through cloud platforms, data, and algorithms as well as finance, governance, talent, information systems, and politics. For example, advanced AI applications are only possible if you have electricity to power data centers, chips to perform computations, networks to move data, and so on. We point out that the entire chain is only as strong as its weakest link. If any one layer (say, access to high-end chips or reliable cloud servers) is compromised, the AI capability as a whole is jeopardized. This also implies that not every AI workload demands a fully sovereign approach; many "commodity" AI uses can rely on standard global services. However, mission-critical AI - the applications core to national security or essential societal functions - do require a sovereign-grade treatment, with special measures to guarantee their functionality under duress.

2. From "AI Opportunities" to "Sovereign AI"

For NATO allies in particular, ensuring AI is developed and operational when needed - despite disruptions - has become as important as developing the AI applications themselves.

Several factors are driving this sovereignty focus. First, there is significant concentration risk in the AI supply chain: for example, only a handful of companies and countries manufacture cutting-edge semiconductors or provide the largest cloud computing platforms. This means many AI projects worldwide depend on the same few external providers, which may raise concerns regarding supply insurance [28–31]. Second, policy and geopolitical risks are growing: export controls, sanctions, or even corporate decisions can suddenly restrict access to essential AI resources.

The move toward sovereign AI is also driven by a significant macroeconomic incentive. As AI emerges as a general-purpose technology, national prosperity will increasingly depend on a country's ability to harness the domestic economic multiplier of this innovation. By building sovereign capabilities, a nation keeps high-value roles in areas such as machine learning and data engineering within the border, creating indigenous employment opportunities, preventing brain drain, and retaining intellectual property within its borders. Furthermore, sovereign AI serves as a catalyst for development and innovation across other critical sovereign domains, such as healthcare and national defense, ultimately boosting long-term economic growth.

We have already seen challenges emerge - the more AI technology spreads and becomes commercially ubiquitous, the more governments find themselves dependent on globally distributed supply chains that could be disrupted. Pursuing sovereign AI often requires accepting some inefficiency or higher upfront costs (like duplicating infrastructure or using domestic alternatives). It is essentially a resilience vs. efficiency

trade-off: the extra investment is an insurance premium against worst-case scenarios where relying on the cheapest global option might leave one empty-handed.

More broadly, the growth of AI is also confronted with physical constraints. In particular, natural resources such as copper, graphite, and stainless steel are unlikely to be produced in sufficient quantities to meet demand for AI services. Further, energy, water, and environmental constraints will likely prove to be serious challenges to AI development. Access to scarce resources are likely to generate continuing geopolitical tensions.

3. Framework: Key Requirements for Sovereign AI

Sovereign AI is a vertically integrated ecosystem. There are six key areas that must be addressed, corresponding to different layers of the AI ecosystem. These range from physical infrastructure to intangible data and algorithms, alongside institutional enablers. The layers are interconnected and together form a cohesive ecosystem. Vulnerabilities at any stage can undermine the performance and resilience of the entire stack. The following sub-sections outline each of these six dimensions and specify the core objectives that must be achieved within each.

3.1. Physical & Infrastructure Prerequisites: Operating AI at Scale

The foundation of sovereign AI is a stable and independent power supply. The energy footprint begins with semiconductor manufacturing and ends with millions of daily interactive uses, not to mention that training AI models involves thousands of GPUs running at full capacity for months. To support the normal functions of AI and remain sovereign, a nation must be able to sustain this extreme electrical load. Such an architecture insulates AI operations from systemic grid instability and external geopolitical pressures on energy supplies. The nation should build robust national grids, together with backup power sources - including local generation - to ensure a reliable electricity supply for AI operations. It must also ensure that the power demand of AI does not crash the civilian power grid. More importantly, since most AI facilities are vulnerable to systemic failures, cyberattacks, or geopolitical sabotage, the state must maintain grid independence and continuity. For example, the nation can deploy dedicated microgrids and on-site generation at priority facilities as backup power, maintaining the basic functions of these facilities during extreme cases. Finally, in the event of a total system collapse, power plants should have a black start capability to restart and recover quickly.

Equally critical is the management of water resources and thermal loads. As AI compute clusters generate immense quantities of heat, continued reliance on traditional water-based cooling systems constitutes vulnerability in water-scarce or contested environments. Thus, states must redirect investment toward advanced alternative cooling technologies, including air cooling [32], immersion cooling [33], and waste-heat recovery [34].

In addition, resource constraints may altogether slow the development of AI and technological infrastructure. In the AI supply chain, copper is foundational for electricity delivery; nickel and graphite are important inputs for lithium-ion batteries used in data centers; and rare earth elements are essential components of disks and chips. The availability of these commodities in the coming years has become a serious concern. As

demand for these materials surges, supply is often physically constrained, since the mining and refining of metals are characterized by long and inflexible production cycles. Furthermore, the supply of these commodities is largely concentrated in specific regions, making it more vulnerable to geopolitical risks. Export controls, trade restrictions, and tariffs can therefore disrupt supply chains, readily turning these materials into choke points for other countries. Maintaining sovereignty over critical commodities begins with the strategic development and utilization of domestic resources wherever feasible. Where full self-sufficiency is unrealistic, sovereign control must be extended through engagement with global markets. Rather than relying on a narrow set of suppliers, countries should diversify across multiple producing regions and stages of the value chain. Long-term offtake agreements, strategic partnerships, and durable diplomatic relationships can serve as insurance mechanisms. When commodity prices fluctuate dramatically, governments should also consider using financial instruments, such as derivatives, to hedge and mitigate these risks, thereby stabilizing input costs.

Beyond the provisioning of raw inputs, AI sovereignty extends to the security and resilience of hardware and computing facilities themselves. Given that the production of advanced semiconductors is concentrated in a limited number of global jurisdictions, a sovereign strategy must prioritize reliable and resilient access to advanced AI accelerators and other critical components. Diversified and trusted supply chains are therefore essential to prevent AI development from being disrupted by trade disputes or diplomatic shifts. And all the hardware and facilities require an integrated security approach that addresses physical and cybersecurity risks in tandem. A unified cyber-physical defense framework is essential to protect against sabotage, hacking, and other forms of disruption [35]. These facilities should be treated as critical national infrastructure and demand continuous and coordinated protection across both domains. It should be noted that sovereign business models may evolve toward decentralized, smaller scale production technologies. For example, the war between Ukraine and Russia has shown how large (nuclear) plants and energy grids can become important targets. Smaller production technologies, such as hydrogen or other forms of renewable energies, may be more difficult to target in case of conflicts.

Further, sovereign AI systems must maintain sufficient buffers and retain the capacity to rapidly expand computational capability under conditions of stress. This requires the deliberate preservation of flexibility in AI infrastructure, including the maintenance of reserve computational headroom that can be activated as surge capacity during emergencies. Such buffers ensure that AI-enabled decision-support and intelligence systems can scale immediately in response to sudden operational demands. Operational endurance also depends on the ability to sustain and repair physical assets in a fully self-sufficient manner. To prevent capability degradation during periods of crisis or isolation, organizations must transition away from just-in-time supply chains toward a model centered on strategic reserves of critical spare parts and maintain a domestically trained technical workforce. This internal capacity guarantees immediate on-site maintenance and repair without reliance on foreign vendors, thereby providing immediate solutions during disruptions.

3.2. Data & Cloud Infrastructure: Facilitating Computational Resilience

The first requirement for secure intelligence is the establishment of a sovereign cloud architecture that ensures critical AI workloads are not dependent on foreign commercial providers. This requires deliberate sovereignty decisions in which each workload is

matched to an appropriate computing environment. For example, low-risk or non-sensitive tasks may operate on public cloud platforms, while high-sensitivity missions should be hosted on fully private clouds or on-premises data centers. Beyond physical deployment, a state must retain full control over identity and access governance, including the management of its own cryptographic keys and authorization policies. No external actor can revoke access, disrupt services, or disable platforms, thereby preserving continuous control over national digital assets.

Once the cloud architecture is secured, attention shifts to the data layer. Achieving genuine data sovereignty requires that sensitive information remain within national or trusted allied jurisdictions, shielding it from extraterritorial legal claims and unauthorized access. This necessitates rigorous oversight of data residency and data provenance, including the ability to verify the origin, integrity, and authenticity of datasets. In parallel, states must prepare for potential disruptions to cross-border data flows by ensuring that AI systems can operate effectively in isolated or degraded connectivity modes. To balance strategic autonomy with international cooperation, secure collaboration frameworks, such as federated learning and secure data enclaves, can be employed, enabling partner states to jointly train models and exchange insights without transferring raw data or relinquishing sovereign control.

The final component involves the implementation of robust resilience, continuity, and disaster recovery mechanisms. Because cloud-based systems remain exposed to cyber intrusions, physical disruptions, and abrupt service interruptions, AI infrastructure must incorporate redundant failover capabilities. This includes maintaining offline backups, mirrored datasets, and secondary computing sites that can be activated immediately in the event of primary system failure. The objective is that essential AI functions including decision-support and intelligence analysis continue operating in a degraded but functional state rather than experiencing a complete shutdown, even during severe crises.

3.3. Models & Algorithms: Achieving Sovereign-by-Design

Sovereignty in AI requires a shift from consuming external services to possessing the model. True credibility is established by developing its own model and owning the fundamental model artifacts that define a model's intelligence and behavior [36]. Specifically, the trained weights must be secured, as they constitute the knowledge stored within the neural network [37]. They are the important inputs for independent air-gapped operations. Also, the model architecture, which defines the structural blueprint and internal flow of information, enables both security audits and hardware optimization [37]. Beyond these backbones, a nation must maintain its own operational metadata and lifecycle infrastructure required for long-term self-sufficiency. This includes control over training datasets to provide transparency regarding provenance and to eliminate external biases [38]. Furthermore, the state must preserve the hyperparameters and the exact software environment used during training. They are all essential components for the reliable reproduction and maintenance of the model. Once the model is controlled domestically, the sovereign AI system does not suffer from operational risk when severed from global updates.

For these artifacts to remain secure, they must be governed by a Test, Evaluation, Validation, and Verification (TEVV) regime [39,40]. This includes absolute control over updates, which freezes the known-good versions in the country. Another example involves regular adversarial robustness testing that intentionally inputs malicious or

misleading data into the model to uncover vulnerabilities and strengthen its overall resilience. TEVV regimes also evaluate the effectiveness of human oversight and control mechanisms [39]. This includes validating clearly defined roles for operators within AI-enabled systems, such as human-in-the-loop or human-on-the-loop configurations for high-consequence decisions. Consequently, TEVV maintains accountability, mitigates automation risk, and ensures that sovereign AI systems remain aligned with human judgment under both normal and adversarial conditions.

By holding these artifacts within the state, a nation secures its core functions against capability suspension, where an external provider might abruptly restrict or disable access due to political pressure, commercial shifts, or "AI throttling." This internal control mitigates the risks of a provider suddenly raising prices exorbitantly or cutting off access during a geopolitical crisis. Moreover, this approach helps with principles-aligned design, where AI systems are built from the ground up to adhere to national safety, ethics, and governance standards. This is not merely for public image, but rather makes sure that AI behaves predictably even in novel situations. A country truly owns its AI outputs by integrating robust fault protection mechanisms, efficient decision-making logs, and strict alignment with command intent. Ultimately, this sovereign-by-design framework transforms AI from a vulnerable third-party service into a resilient and accountable governing infrastructure that remains functional precisely when national security is most at stake.

A comprehensive sovereign strategy must also manage the development of small-scale, efficient models designed for edge deployment. While massive cloud-based models offer high performance, they are often brittle in environments where constant cloud connectivity is impossible. In contrast, compact models capable of running locally provide a critical baseline of capability. For example, drones and ground vehicles use compressed models for real-time obstacle avoidance and path planning, even when GPS or communication links are jammed [41]. Similarly, handheld medical devices equipped with lightweight AI can assist in triage and diagnostics in remote field hospitals, operating entirely offline to keep patient data fully local and immediate [42]. Furthermore, smart sensors in power grids or water systems can use local inference to detect anomalies, such as leakages or surges, and trigger automatic shut-offs locally [43]. This can function well even if the central command center is offline. These deployable models are the safety nets that even if a nation is cut off from the global internet, its frontline systems remain operational and effective.

3.4. Financing & Investment Infrastructure: Building the Enabling Layer

Funding the development of sovereign AI requires a transition in fiscal logic, from cost optimization toward resilience optimization and sovereign investment models. Conventional financial models favor outsourcing several parts of supply chains to global providers on the basis of lower opportunity cost. A sovereign strategy, by contrast, requires acceptance of a resilience premium that is costly in the short run but beneficial in the long run. Building sovereign AI involves substantial upfront capital expenditure on domestic data centers, energy infrastructure, and secure supply chains. These investments should be treated not as routine procurement costs, but as strategic insurance against systemic dependency failures.

Given the scale of capital required, sovereign AI cannot be financed by any single entity in the economy. Long-term autonomy therefore depends on coordinated participation across government, private industry, and financial intermediaries. At the

core of this strategy, the government plays an anchoring role through instruments such as procurement policy, Sovereign Wealth Funds (SWF), and Development Financial Institutions (DFI). Government procurement can serve as a powerful policy lever by creating guaranteed demand for domestic AI solutions, thereby improving their investment profile and accelerating private capital participation [44]. SWFs and DFIs can act as an anchor to absorb early-stage risk by providing first-loss protection or subsidizing insurance to private partners. In addition, the intergenerational horizon of SWFs and DFIs makes them uniquely positioned to finance capital-intensive hard technologies, including secure energy grids, domestic semiconductor fabrication, and large-scale data infrastructure. Although these projects are strategically essential, they are commercially unattractive in the short run and thus are often deterred by private investment. The holdings by SWFs and DFIs could also be a positive signal of firm quality to private investors, attracting the participation of private sector funds [3,45].

On the other hand, the private sector players, in particular private equity (PE) and private credit, are major investors in innovation and technology and often boost the industrialization of AI [6]. Private equity firms typically bring expertise on complex capital structures, operational scaling, and asset management, enabling an efficient build-out of the physical AI layer. Venture capital (VC) typically fuels the algorithmic and application layers of the AI stack. VC serves as a discovery mechanism for high-potential domestic innovations that can be scaled rapidly [6,46]. Through sovereign–private consortiums, private investors can rely on state-backed guarantees to reduce the risk assessment of large infrastructure projects. This facilitates the construction of AI facilities, which are then operated with commercial discipline while remaining subject to sovereign security requirements and priority access provisions during national emergencies. Over time, VC-backed innovations can be scaled and integrated into larger platforms supported by private equity and sovereign capital, creating a continuous pipeline from early-stage innovation to national-scale deployment. In the US, the FIRRMA restricts strategic innovation to domestic investors [4]. In Europe, the private equity sector is substantially less developed. There appears to be a scarcity of European institutional investors. Instead, development financial institutions (DFIs) invest in VC funds but appear to be limited [3].

Alternatively, sovereign AI projects could choose to be funded by public-private-partnerships (PPP). PPPs involve collaboration between a government agency and a private-sector company to finance, build, and operate projects, allowing large-scale government projects to be completed with private funding [47]. While the public partner defines objectives and monitors compliance, the private partner designs, completes, implements, and funds the project. And the risks are shared across two parties. Structures such as Build-Operate-Transfer or Design-Build-Finance-Operate convert large upfront investments into predictable, performance-based payments. These contracts typically include step-in rights that the state can assume control of critical facilities during emergencies or supply chain disruptions. In this way, operational efficiency is achieved without sacrificing sovereign accountability.

To align these diverse capital flows, the financial framework must explicitly value resilience. A resilience dividend model provides one possibility [48]. It quantifies the economic and security losses associated with dependency failures, such as cloud service outages, hardware embargoes, or sudden access restrictions. With various sources of risk accounted for, policymakers can compare the cost of sovereign investment against the far larger costs of systemic failure. This reframes high-capital projects from inefficient

spending to prudent national insurance, helping sustain political commitment across budget cycles and electoral horizons.

In addition, long-term financial sustainability requires that sovereign AI generate near-term economic value. Rather than relying indefinitely on subsidies, the ecosystem should transit toward a demand-driven, purpose-fit market, by prioritizing high-demand domains such as civil administration, energy systems, healthcare, and precision agriculture. A general sovereign AI can be provided as a secure base model, and domestic firms can develop specialized applications and agents tailored to specific sectors. This can accelerate the adoption of AI and its application, meanwhile these deployments generate immediate returns that can be reinvested into future research and development. Such a reinforcing, self-sustaining innovation cycle can be attractive to both public and private capital while strengthening national AI capability.

Tokens, platforms, and cryptocurrencies are also increasingly used as an alternative means of financing and payment [49–52]. In particular, they appear to be used broadly during geopolitical conflicts. Alexakis, Anselmi, and Petrella (2024) find evidence that there was a strong outflow from Ukrainian Hryvnia to cryptocurrencies following the conflicts between Russia and Ukraine [53]. It is driven by the individuals seeking to protect against their savings from devaluation and using cryptocurrencies as an alternative payment system when the national banking industry was crippled. Alvarez, Argente, and Patten (2022) document El Salvador's attempt in making Bitcoin legal tender in September 2021 and investigate the potential of cryptocurrencies to become a sovereign medium of exchange [54]. Their finding shows that despite a significant initial push by the government, the use of Bitcoin for daily transactions remains low and is primarily concentrated among those banked, educated, young male. 20% of the firms accept Bitcoin as a payment method, but 88% of which convert Bitcoin to dollars immediately after receiving the payment. Policymakers are actively exploring stablecoin infrastructure and considering the possibility that stablecoins become sovereign cross-border infrastructure [51]. Recently, A7A5 token, a stablecoin pegged to Russian Rupee, serves as a tool for shadow payments, as it facilitates cross border financial flows into and out of Russia despite of Western sanctions [55]. Howcroft and Wilkes (2026) reveal that Iran's crypto activity surges during recent geopolitical tension [56]. The Iranian government and its military arms are using cryptocurrencies and stablecoin transferring money to bypass US sanctions. In the meantime, approximately 20% of the general population use or hold crypto assets to hedge against high inflation in the country. These alternative financial instruments, when integrated into a robust legal framework, could complement the traditional financing sectors.

3.5. Supply Chain Governance: Managing Full-Stack Sovereignty

True sovereignty is not a binary state but a function of end-to-end resilience. It is insufficient to secure a single link if the surrounding links remain vulnerable. To assess the credibility of the supply chain, the state can use holistic dependency mapping to conduct an exhaustive audit of all inputs, energy, hardware, software, data, and human capital, required to maintain AI capability. This can be paired with choke-point analysis to identify single points of failure within the supply chain. For example, only one or two companies globally may be capable of fabricating the most advanced chips used in AI, a critical rare mineral for chip production may be sourced primarily from a rival power, or a single undersea cable may carry most of the nation's data traffic. After clearly identifying such choke points, strategies can be developed to invest in redundancies,

substitutes, stockpiles, or diplomatic arrangements to mitigate these risks. Regular stress testing should be conducted to simulate extreme scenarios, such as the sudden loss of access to a semiconductor supplier, a cyber-kinetic attack on a data center, or the shutdown of AI facilities during natural disasters. By regularly assessing key risks, the state can prepare backup plans and response mechanisms to effectively handle emergencies.

Further, it should be noted that just-in-time manufacturing creates a number of vulnerabilities, in particular in times of market stress [61–64]. The safety and fit of all components of the supply chain with the production process are also critical. Software and data should be protected as rigorously as hardware. Vertical integration often provides all parts of the supply chain with superior incentives [65]. When this is not feasible, this requires tracking the provenance of open-source libraries and AI models, using secure build processes and code verification, and ensuring that updates and patches are authentic and safe. Clear jurisdictional safeguards are essential for the industry [57,58]. For example, the state should establish clear legislative guidance for all domestic participants in the supply chain, protecting user privacy and developer intellectual property [59,60]. At the same time, when services or suppliers are located in other countries, they may be subject to foreign legal regimes. In such cases, sovereign AI architectures should favor arrangements in which legal jurisdiction aligns with trust, so that AI operations are less likely to be disrupted, hijacked or shut down from abroad.

3.6. Human Capital & Talent Development: Securing the Intellectual Foundation

Finally, human capital is arguably the most critical component of the ecosystem beyond physical infrastructure. AI systems require adequate talent development. In many countries, the workforce is well educated but lacks the skills required to operate advanced technologies and AI systems. Traditional academic models often lag behind the rapid evolution of AI, leaving even well-educated graduates unable to meet the recruitment requirements of firms. Special attention should therefore be paid to education systems to develop these capabilities. Modern education systems should move away from rote memorization, fixed syllabi, and standardized answers, and instead focus on cultivating students' abilities in critical thinking, problem framing, and continuous reskilling. Governments should also emphasize the development of talent in fundamental disciplines such as mathematics and physics. In doing so, the state can create a self-sustaining pipeline of scientists, engineers, and system architects.

Beyond basic education and training, talent sovereignty depends on the creation of high-value career paths that provide domestic experts with the intellectual and financial incentives to remain within national borders. This is fundamentally achieved through the construction of a comprehensive, vertically integrated supply chain stack that fosters a dense ecosystem of private and public enterprises. In such a thick labor market, talent faces a lower risk of unemployment and is not overly exposed to the failure of any single firm. If a scientist leaves one project, there are ten other well-funded firms in the same supply chain ready to hire him. To sustain this workforce, the nation must also cultivate a premier research environment that treats the generation of new ideas as a primary industrial output. This requires streamlined funding mechanisms, the establishment of national research centers, and dedicated regulatory sandboxes for experimentation. In addition, a sufficient and readily accessible computational capacity, founded by the physical layer of the supply chain, must be available to support large-scale research and experimentation.

To achieve true independence, a state must ensure that its AI development is driven by a robust domestic workforce, rather than relying excessively on external personnel or foreign firms, and that its leading scientists and experts are not lost to foreign employers or lured abroad by superior opportunities. At the same time, the nation should remain open to global talent. Governments can offer attractive incentives, such as immediate leadership roles in national laboratories, fast-track research funding, and comprehensive family relocation support, to encourage top scientists to return to their home country. In addition, a nation must also actively attract elite talent from around the world. By offering specialized visa programs and streamlining application processes, experts in high-demand fields can be drawn into the country, where they contribute to the development of a stronger and more resilient AI ecosystem.

Overall, the government should cultivate a premier ecosystem that offers robust domestic training pipelines, world-class research environments, and amplified career opportunities. While the goal is self-sufficiency, a nation should also position itself as a premier global destination that attracts the best experts from all over the world.

4. Convergence: Emerging Technologies, Information Systems, and Sovereign AI

Sovereign AI also intersects with other emerging technologies, reinforcing the need for a holistic approach [24,27]. For example, integrating future quantum computing capabilities will require strategic choices - quantum resources will be limited, so they should be applied where they have the greatest impact under constraints. Similarly, the drive for edge computing and autonomy aligns with sovereignty: pushing AI to operate on local devices (with minimal reliance on connectivity) ensures that critical functions can persist even in jammed or disconnected environments. Cybersecurity is also vital: AI systems must be robust against cyber-attacks, since losing AI access through hacking can be as damaging as losing it through supply denial. Moreover, semiconductors form a common foundation for AI, quantum, and other advanced tech, so securing chip supply benefits multiple domains. Finally, communications resilience remains crucial - AI networks should be designed to gracefully handle outages or interference through distributed architectures and backup links.

Cyber risks and foreign bot–driven information manipulation may directly affect the sovereignty of AI systems. In particular, the ability of a state or organization to control how its AI is trained, operated, and trusted for decisions may be targeted because the most sovereignty-critical assets are often data, model behavior, and the information environment models draw from. Cyberattacks such as data poisoning, model tampering, and exploitation of system vulnerabilities can silently alter outputs or degrade reliability, meaning the “owner” of the system no longer fully governs what the AI does or why it does it. At the same time, foreign influence operations using coordinated bots and synthetic content can pollute training corpora, manipulate retrieval sources, and flood platforms with narratives that steer user prompts and downstream human decisions - effectively outsourcing parts of the AI’s “world model” and outputs to adversarial actors.

In addition, AI can be used to compromise security systems. It also makes it very easy to create considerable numbers of bots or to write on behalf of social media accounts. This can create new risks, e.g. when foreign countries want to influence electoral results or citizen behavior. This implies that with the current internet infrastructure, the large

number of people who obtain information primarily through social media can be easily manipulated.

5. Illustrative Cases

There are numerous examples and scenarios that help illustrate why sovereign AI measures are needed.

For instance, cyber-attacks on Ukraine's power grid (and the grid's remarkable resilience) highlight the importance of assured energy supplies for keeping AI systems online during conflict. Another scenario is the reliance on a single cloud service provider - if that provider suffers an outage or intentionally withdraws service, it can instantly paralyze numerous AI applications that lack an alternative. A third example is export controls on hardware: if a country is suddenly cut off from advanced AI chips or other key components, its ability to train new models or maintain current systems could quickly erode. Finally, incidents of data poisoning have shown that an AI system can be compromised not only by physical disruptions but by subtle corruption of its training data undermining its reliability.

These cases illustrate the diverse ways in which external events can disrupt AI, reinforcing why resilience and sovereign control are vital.

6. Policy recommendations

Overall, to build sovereign AI, a nation and allied nations need to effectively control the full stack of the AI supply chain, from physical inputs to institutional governance. Sovereignty can fail if any layer of the supply chain is broken or choked by foreign competitors. Policy design should therefore treat AI not as a standalone digital technology, but as an integrated industrial system, whose weakest link determines its overall resilience.

Policymakers should prioritize reliable, scalable, robust, and controllable energy and AI facilities. This includes producing sustainable energy, rebuilding electricity power grids, and setting up modern data centers accompanied by advanced cooling technologies. Countries should manage the supply of critical commodities such as copper, rare earths, semiconductors, and related inputs. These resources are often globally concentrated and physically constrained. It is important to fully exploit domestic resources where possible. When this is not feasible, countries could turn to global markets, maintaining diversified suppliers and stable political relationships with producing countries. Commodity price volatility should also be managed through appropriate risk-hedging instruments, such as derivatives.

Compute hardware, data-center facilities, and advanced chips sit at the core of the AI supply chain. While complete domestic manufacturing may be unrealistic, countries should aim to occupy strategically important positions within global value chains - for example, by holding critical patents or technological capabilities that are difficult to circumvent. This reduces the risk of external choke points. Policymakers should ensure assured access to key components and establish clear contingency plans for potential supply disruptions.

Data and cloud infrastructure form the connective tissue of the AI supply chain. Sovereignty at this layer does not require isolation from global data flows, but it does

require enforceable control over data storage, access restrictions, and cybersecurity risks. Policymakers should establish clear standards governing data usage, data security, and user privacy protection.

At the model and algorithmic layer, sovereignty depends on ownership and the replicability of critical model artifacts, ranging from training data pipelines and hyperparameters to model architectures. Equally important, model validity, robustness, and reliability must be continuously tested and monitored over time.

The development of sovereign AI should not occur in a vacuum. Instead, it should be aligned with essential societal needs and focus on solving critical problems that were previously intractable. This requires AI systems with broad application potential, which can attract market participants, generate economic returns, and help fund ongoing research and development. This flywheel creates a sustainable ecosystem that does not rely solely on government intervention, but also incentivizes private actors to move in the desired direction. Policymakers should therefore safeguard market integrity and promote fair competition.

Finally, finance and human capital flow across the entire AI supply chain. Governments should coordinate public funding, private investment, and public–private partnerships to support strategically important layers of the AI stack. While public funding provides direction and stability, private capital adds dynamism and efficiency. Policymakers should also monitor emerging financial instruments - including cryptocurrencies, stablecoins, and central bank digital currencies (CBDCs) - which may complement traditional financing channels when appropriately regulated. With respect to human capital, sovereign AI strategies should prioritize education, talent retention, secure research environments, and diverse career pathways to reduce brain drain, alongside targeted immigration policies that attract global talent into the country rather than away from it.

Ultimately, controlling the full AI supply chain is not an end in itself, but a means to achieve measurable resilience under adverse conditions. Systematic stress testing across the entire AI stack, comprehensive backup and reactivation plans, and clearly defined fallback strategies in the face of geopolitical uncertainty should become industry standards and be implemented on a regular basis. The greater the resilience, the stronger the sovereignty.

7. Conclusion

Sovereign AI requires credible control of supply chains, sovereign financing, and sovereign information systems. The framework presented in this paper illustrates that a nation's AI capability is only as robust as the weakest link in its dependency stack. Accordingly, success with AI will depend not only on having cutting-edge algorithms, but also on securing all the layers - commodities, energy, compute, data, talent, and more - that those algorithms rely on. In the end, the strategic advantage will belong to those who can keep their AI systems operational through disruption and crises.

References

[1] Arquilla J, Ronfeldt D. The emergence of noopolitik: toward an American information strategy. 1999 Jan. https://www.rand.org/pubs/monograph_reports/MR1033.html.

[2] European parliament resolution of 22 January 2026 on European technological sovereignty and digital infrastructure (2025/2007 (ini)). European Parliament; 2026 Jan. https://oeil.europarl.europa.eu/oeil/en/procedure-file?reference=2025/2007(INI).

[3] Andonov A, Li A, Smeets P. Do development financial institutions create impact through venture capital investments? SSRN Electron J 2025. doi: 10.2139/ssrn.5101408.

[4] Paine F, Townsend RR, Xu T. Should governments restrict foreign investments in startups? Entrepreneurship and Innovation Policy and the Economy, Volume 5. University of Chicago Press; 2025 Apr. https://www.nber.org/books-and-chapters/entrepreneurship-and-innovation-policy-and-economy-volume-5/should-governments-restrict-foreign-investments-startups.

[5] Quinet A, Jaravel X, Schularick M, Zettelmeyer J. Economic principles for European rearmament. Kiel Policy Brief 2025 Sept;195. https://www.kielinstitut.de/publications/economic-principles-for-european-rearmament-18868/.

[6] Chemla G, Kitten M. Private equity in healthcare and life sciences. John Wiley & Sons; 2026.

[7] Bolton P, Farrell J. Decentralization, duplication, and delay. J Polit Econ 1990;98(4):803–26. http://www.jstor.org/stable/2937769.

[8] Tirole J. The theory of corporate finance. Princeton university press; n.d.

[9] Bolton P, Dewatripont M. Authority in organizations
[the handbook of organizational economics]. Introd Chapters 2012. https://ideas.repec.org//h/pup/chapts/9889-9.html.

[10] Conover M, Ratkiewicz J, Francisco M, Goncalves B, Menczer F, Flammini A. Political polarization on twitter. Proc Int AAAI Conf Web Soc Media 2021 Aug;5(1):89–96. doi: 10.1609/icwsm.v5i1.14126.

[11] Bakshy E, Messing S, Adamic LA. Exposure to ideologically diverse news and opinion on Facebook. Science 2015 June;348(6239):1130–2. doi: 10.1126/science.aaa1160.

[12] Bail CA, Argyle LP, Brown TW, Bumpus JP, Chen H, Hunzaker MBF, et al. Exposure to opposing views on social media can increase political polarization. Proc Natl Acad Sci 2018 Sept;115(37):9216–21. doi: 10.1073/pnas.1804840115.

[13] Cinelli M, De Francisci Morales G, Galeazzi A, Quattrociocchi W, Starnini M. The echo chamber effect on social media. Proc Natl Acad Sci 2021 Mar;118(9):e2023301118. doi: 10.1073/pnas.2023301118.

[14] Doerr S, Gissler S, Peydró J, Voth H. Financial crises and political radicalization: how failing banks paved Hitler's path to power. J Finance 2022 Dec;77(6):3339–72. doi: 10.1111/jofi.13166.

[15] Funke M, Schularick M, Trebesch C. Populist leaders and the economy. Am Econ Rev 2023 Dec;113(12):3249–88. doi: 10.1257/aer.20202045.

[16] Moll B, Schularick M, Zachmann G. The power of substitution: the great German gas debate in retrospect 2023 Sept. https://www.brookings.edu/articles/the-power-of-substitution-the-great-german-gas-debate-in-retrospect/.

[17] Balan S, Vorotynskyy V, Patalakha V, Rybak I, Tarasiuk V. Information resilience of society and sustainable democratic development in the context of political transformations. Riv STUDI SULLA SOSTENIBILITA 2026 Jan;10(2):13. doi: 10.3280/riss2025oa21091.

[18] Gajewska M, Katulski J. Doppelgangers, oligarchs and distortions. research notes on Russian disinformation tactics in Israel and their influence on western partners. Ann Univ Paedagog Cracoviensis Stud Ad Bibl Sci Pertin 2026 Jan;23:610–34. doi: 10.24917/20811861.23.29.

[19] Vybranovska S. Hybrid warfare and Russian active measures: disinformation as a tool to undermine democratic processes 2025. https://ekmair.ukma.edu.ua/handle/123456789/36219.

[20] Hall GJ, Sargent TJ. Three world wars: fiscal–monetary consequences. Proc Natl Acad Sci 2022 May;119(18):e2200349119. doi: 10.1073/pnas.2200349119.

[21] Ettmeier S, Kriwoluzky A, Schularick M, Steege L ter. Fatal austerity. CERGE-EI Work Pap 2025 Oct. https://ideas.repec.org//p/cer/papers/wp801.html.

[22] Schularick M, Federle J. The price of war. Kiel Policy Brief 2024 Feb;171. https://www.kielinstitut.de/publications/the-price-of-war-17475/.

[23] Chemla G, Tena V. AI adoption, incentives, and firm value in general equilibrium * 2026. doi: 10.2139/ssrn.5997775.

[24] Quddus Islam T, Garlick R, Fei W, Chana D, Gogioso S, Bose S, et al. Quantum computing: moving quickly from theory to reality. 2023 July. https://ir.citi.com/gps/nh6PRNGut2dR%2BanVBfV8YtiV33RvBKQo3jNLwYvfjZJlTVKHYMGV3MicyGBcqqwlxVLhv7bO7aBa%2FLLNdXwpIA%3D%3D.

[25] Quddus Islam T, Garlick R. Quantum sensing: tech's new eyes and ears. Citigroup Inc.; 2024 Feb. https://www.citivelocity.com/t/r/eppublic/2zlpv.

[26] Garlick R, Fei W, Quddus Islam T, Odunsi A, Spielman A, Wilkie M, et al. The rise of ai robots: physical AI is coming for you. Citigroup Inc.; 2024 Dec. https://ir.citi.com/gps/H558-XNr_iTlGa7Qq7H9AYb5ZT2W851WZdFgPNEDsBtSeTgp7JcaTdS_uBfLVLpwfMQYeB5O5TwV9YcIDGuGOMjE2luzQprf.
[27] Quddus Islam T, Garlick R, Krause HH. Mapping the innovation landscape: comparing 100 emerging and disruptive technologies, business models, and ideas. 2024 Apr. http://www.citivelocity.com/t/eppublic/34TqS.
[28] Bolton P, Whinston MD. Incomplete contracts, vertical integration, and supply assurance. Rev Econ Stud 1993;60(1):121–48. doi: 10.2307/2297815.
[29] Chemla G. Downstream competition, foreclosure, and vertical integration. J Econ Manag Strategy 2003 June;12(2):261–89. doi: 10.1111/j.1430-9134.2003.00261.x.
[30] Rey P, Tirole J. Chapter 33 a primer on foreclosure. Handbook of Industrial Organization. Elsevier; 2007. p. 2145–220. doi: 10.1016/S1573-448X(06)03033-0.
[31] Aïd R, Chemla G, Porchet A, Touzi N. Hedging and vertical integration in electricity markets. Manag Sci 2011 Aug;57(8):1438–52. doi: 10.1287/mnsc.1110.1357.
[32] Chang Y-W, Chang C-C, Ke M-T, Chen S-L. Thermoelectric air-cooling module for electronic devices. Appl Therm Eng 2009 Sept;29(13):2731–7. doi: 10.1016/j.applthermaleng.2009.01.004.
[33] Kuncoro IW, Pambudi NA, Biddinika MK, Widiastuti I, Hijriawan M, Wibowo KM. Immersion cooling as the next technology for data center cooling: a review. J Phys Conf Ser 2019 Dec;1402(4):044057. doi: 10.1088/1742-6596/1402/4/044057.
[34] Jouhara H, Khordehgah N, Almahmoud S, Delpech B, Chauhan A, Tassou SA. Waste heat recovery technologies and applications. Therm Sci Eng Prog 2018 June;6:268–89. doi: 10.1016/j.tsep.2018.04.017.
[35] Zhang H, Liu B, Wu H. Smart grid cyber-physical attack and defense: a review. IEEE Access 2021;9:29641–59. doi: 10.1109/ACCESS.2021.3058628.
[36] Ferraris AF, Audrito D, Caro LD, Poncibò C. The architecture of language: understanding the mechanics behind llms. Camb Forum AI Law Gov 2025;1:e11. doi: 10.1017/cfl.2024.16.
[37] Mukhamediev RI, Popova Y, Kuchin Y, Zaitseva E, Kalimoldayev A, Symagulov A, et al. Review of artificial intelligence and machine learning technologies: classification, restrictions, opportunities and challenges. Mathematics 2022 July;10(15):2552. doi: 10.3390/math10152552.
[38] Liesenfeld A, Dingemanse M. Rethinking open source generative ai: open washing and the eu ai act. The 2024 ACM Conference on Fairness, Accountability, and Transparency. Rio de Janeiro Brazil: ACM; 2024 June. p. 1774–87. doi: 10.1145/3630106.3659005.
[39] Waters G. Testing, evaluation, verification and validation (tevv) of digital twins: a comprehensive framework 2025. doi: 10.48550/ARXIV.2507.04555.
[40] Tabassi E. Artificial intelligence risk management framework (ai rmf 1.0). Gaithersburg, MD: National Institute of Standards and Technology (U.S.); 2023 Jan. doi: 10.6028/NIST.AI.100-1.
[41] Ji J, Khajepour A, Melek WW, Huang Y. Path planning and tracking for vehicle collision avoidance based on model predictive control with multiconstraints. IEEE Trans Veh Technol 2017 Feb;66(2):952–64. doi: 10.1109/TVT.2016.2555853.
[42] Gerke S, Babic B, Evgeniou T, Cohen IG. The need for a system view to regulate artificial intelligence/machine learning-based software as medical device. Npj Digit Med 2020 Apr;3(1):53. doi: 10.1038/s41746-020-0262-2.
[43] Ranyal E, Sadhu A, Jain K. Road condition monitoring using smart sensing and artificial intelligence: a review. Sensors 2022 Apr;22(8):3044. doi: 10.3390/s22083044.
[44] Belenzon S, Cioaca L. Guaranteed demand and corporate r&d. Manag Sci 2025 Aug;:mnsc.2023.00087. doi: 10.1287/mnsc.2023.00087.
[45] Dewenter KL, Han X, Malatesta PH. Firm values and sovereign wealth fund investments. J Financ Econ 2010 Nov;98(2):256–78. doi: 10.1016/j.jfineco.2010.05.006.
[46] De Bettignies J-E, Chemla G. Corporate venturing, allocation of talent, and competition for star managers. Manag Sci 2008 Mar;54(3):505–21. doi: 10.1287/mnsc.1070.0758.
[47] Engel E, Fischer R, Galetovic A. The basic public finance of public-private partnerships. J Eur Econ Assoc 2013 Feb;11(1):83–111. doi: 10.1111/j.1542-4774.2012.01105.x.
[48] Surminski S, Tanner T, editors. Realising the “triple dividend of resilience.” Cham: Springer International Publishing; 2016. doi: 10.1007/978-3-319-40694-7.
[49] Chemla G, Tinn K. Learning through crowdfunding. SSRN Electron J 2017. doi: 10.2139/ssrn.2796435.
[50] Chemla G, Tinn K. How wise are crowds on crowdfunding platforms? SSRN Electron J 2019. doi: 10.2139/ssrn.3506377.
[51] Chemla G, Knottenbelt W, Li Z, Xiong X, Delgado De Molina Rius A, Burnett AL. Mind the gap: how stablecoins can secure the UK’s financial future * 2025. doi: 10.2139/ssrn.5486386.
[52] Li S. How do flash loans affect market liquidity? SSRN Electron J 2025. doi: 10.2139/ssrn.5932835.

[53] Alexakis C, Anselmi G, Petrella G. Flight to cryptos: evidence on the use of cryptocurrencies in times of geopolitical tensions. Int Rev Econ Finance 2024 Jan;89:498–523. doi: 10.1016/j.iref.2023.07.054.
[54] Alvarez F, Argente D, Van Patten D. Are cryptocurrencies currencies? bitcoin as legal tender in El Salvador. Cambridge, MA: National Bureau of Economic Research; 2022 Apr. doi: 10.3386/w29968.
[55] Hawkins O, Ivanova P, Hemingway E. Crypto coin for Russian shadow payments moves $9bn. Financ Times 2025 June.
[56] Howcroft E, Wilkes TR. Exclusive: Iran's surging crypto activity draws us scrutiny. Reuters 2026 Feb. https://www.reuters.com/business/finance/irans-surging-crypto-activity-draws-us-scrutiny-2026-02-03/.
[57] Chemla G, Hennessy CA. Government as borrower of first resort. J Monet Econ 2016 Dec;84:1–16. doi: 10.1016/j.jmoneco.2016.09.001.
[58] Chemla G, Hennessy CA. Skin in the game and moral hazard. J Finance 2014;69(4):1597–641. http://www.jstor.org/stable/43611198.
[59] Chemla G. Takeovers and the dynamics of information flows. Int J Ind Organ 2004 Apr;22(4):575–90. doi: 10.1016/j.ijindorg.2003.11.001.
[60] Chemla G. Hold-up, stakeholders and takeover threats. J Financ Intermediation 2005 July;14(3):376–97. doi: 10.1016/j.jfi.2004.08.004.
[61] Chemla G, Faure-Grimaud A. Dynamic adverse selection and debt. Eur Econ Rev 2001 Oct;45(9):1773–92. doi: 10.1016/S0014-2921(00)00080-5.
[62] Moran J, Romeijnders M, Doussal PL, Pijpers FP, Weitzel U, Panja D, et al. Timeliness criticality in complex systems. Nat Phys 2024 Aug;20(8):1352–8. doi: 10.1038/s41567-024-02525-w.
[63] Moran J, Pijpers FP, Weitzel U, Bouchaud J-P, Panja D. Critical fragility in sociotechnical systems. Proc Natl Acad Sci 2025 Mar;122(9):e2415139122. doi: 10.1073/pnas.2415139122.
[64] Martin D, Moran J, Panja D, Bouchaud J-P. Resilient-to-fragile transition and excess volatility in supply chain networks 2026. doi: 10.48550/ARXIV.2601.20450.
[65] Grossman SJ, Hart OD. The costs and benefits of ownership: a theory of vertical and lateral integration. J Polit Econ 1986 Aug;94(4):691–719. doi: 10.1086/261404.