



7 September 2026

ASIA THEMATIC INSIGHTS

WRITTEN BY



Alicia GARCIA HERRERO

Tel. +85239008680

alicia.garciaherrero@natixis.com

Soumitra DUTTA¹

This is a reprint of Bruegel Working Paper 25/2026.



<https://www.bruegel.org/working-paper/programmable-sovereignty-coordinated-specialisation-response-n-2-problem-artificial>

Programmable sovereignty: coordinated specialisation as a response to the N-2 problem in artificial intelligence

The dominant academic and political framing of power in artificial intelligence treats sovereignty as a function of production: a state is sovereign over AI to the extent that it builds, owns and can reproduce the AI stack. By this definition, meaningful AI sovereignty has been achieved only by the United States and, to a lesser extent, China. We term this the N-2 problem: all other states face structural dependence on one or both of these AI superpowers.

A separate and largely disconnected line of analysis treats digital sovereignty as a function of governance: in other words, the capacity to shape technology through rules, standards and market leverage, regardless of who builds the underlying systems.

We combine insights from these two literatures to make two contributions. First, we introduce the notion of 'programmable sovereignty': soft governance anchored by enough hard-stack control (especially pooled compute capacity) to make governance commitments credible and dependence reversible, without requiring full-stack AI autonomy. Second, we propose 'coordinated specialisation' as an institutional mechanism through which states facing the N-2 dilemma can pool capacities into a credible shared AI stack. We outline criteria for membership of this coalition, with reference to three institutional precedents (EuroHPC, India Stack and the ASEAN Digital Economy Framework). We conclude with policy implications.

Discover more on our [Website](#)

¹ Soumitra Dutta is director at Portulans Institute.

1. Introduction

Artificial intelligence is commonly framed as a geostrategic race to build the largest models, accumulate the most patents and announce the boldest strategies. The central analytical task is to measure progress and predict a winner. This framing is misleading. What we see is not a race in which many credible entrants compete on equal terms, but a consolidation of technological power in the United States and China. These two states are approaching near-full control over AI production systems, while every other economy faces dependence on one or both of the US and China.

This dependence is being converted into leverage. For example, in April and October 2025, China extended existing export controls on rare earth elements to cover medium and heavy elements, associated processing technologies and equipment¹. It also introduced licensing requirements for any foreign-made product that contains even trace quantities of Chinese-origin rare earth elements. In April 2026, Chinese authorities ordered the unwinding of Meta's \$2 billion acquisition of Manus, a Singapore-incorporated but China-origin AI startup. China ordered the transaction to be unwound on national-security grounds related to the transfer of autonomous-agent technology and personnel².

Then, in mid-June 2026, the US Department of Commerce gave AI firm Anthropic a reported ninety minutes³ to restrict access to two of its most capable models, Fable 5 and Mythos 5, to US citizens only. In response, Anthropic cut all access to both models⁴. The US administration lifted the order at the end of June.

These three episodes, occurring over a little more than a year, fit a pattern of regulatory actions. They demonstrate what this paper terms the 'N-2 problem': in a world of N states, only two are currently capable of pursuing full-stack sovereignty in AI. The rest – N-2 – face dependence on the US and/or China, which can be leveraged by either duopoly member, with little or no warning.

The N-2 problem mirrors the international political economy analysis on structural dependence in a context of monetary hegemony: the costs to non-reserve-currency states of operating within a dollar-centred international monetary system (ie N-1; Eichengreen, 2011; Cohen, 2018; Helleiner, 2014). We argue that AI dependence is more consequential than this N-1 situation for three reasons: it penetrates public administration, healthcare and security more deeply than macroeconomic policy; the costs of switching AI stack are categorically higher than the costs of switching reserve currencies; and the window during which credible alternatives can still be constructed is closing as supply chains and regulatory norms are becoming settled. These three factors are discussed further in section 2.

Research on AI sovereignty offers two responses to the N-2 problem (section 3). One branch of literature, which we term the 'hard-sovereignty tradition', treats AI sovereignty as a function of productive capacity, requiring ownership and control of the full technology stack, from critical minerals and energy inputs through to fabrication, foundation models and capital (Lee, 2018; Buchanan, 2020; Miller, 2022; Scharre, 2023; Ding, 2024). This results in a pessimistic conclusion for the great majority of states: that meaningful sovereignty requires fabrication capacity, frontier models and compute infrastructure that no economy outside the US and China can plausibly muster today.

The second literature, which we term the 'soft-sovereignty tradition', defines AI sovereignty as contingent instead on governance capacity: the ability to shape how technology is deployed, through laws, standards and market access, regardless of who builds the underlying systems

¹ See Ministry of Commerce of the People's Republic of China, 'Announcement No.18 of 2025 of the decision to implement export control on some medium and heavy rare earth related items', 4 April 2025, https://english.mofcom.gov.cn/Policies/AnnouncementsOrders/art/2025/art_Odd87cbee7b045bf93fabe6ab2faceee.html and Xin Liu, Nabeel Yousef, Stephanie Brown Cripps and Pauline Peng, 'China Intensifies Export Controls over Rare Earths and Related Technologies', *Freshfields*, 14 October 2025, <https://www.freshfields.com/en/our-thinking/blogs/risk-and-compliance/china-intensifiesexport-controls-over-rare-earths-and-related-technologies-102lpyn>

² Liv McMahon, 'China blocks Meta's \$2bn acquisition of AI start-up Manus', *BBC*, 27 April 2026, <https://www.bbc.com/news/articles/cj0v0gr2yz7o>.

³ Miranda Nazzaro, 'Trump administration lifts limits on Anthropic's new AI models', *The Hill*, 30 June 2026, <https://thehill.com/homenews/5948719-anthropic-fable-mythos-models-restored/>.

⁴ Alicia García-Herrero and Soumitra Dutta, 'Europe and its allies must plan to offset the AI dependency risk', *First Glance*, 18 June 2026, Bruegel, <https://www.bruegel.org/first-glance/europe-and-its-allies-must-planoffset-ai-dependency-risk>.

(Floridi, 2020; Bradford, 2020, 2023; O'Hara and Hall, 2018). This literature offers a more attractive scenario for states outside the duopoly. Nevertheless, we argue that this, too, is insufficient if pursued alone: governance capacity without any material compute capability is insufficient to prevent the kind of unilateral action that the June 2026 Anthropic episode illustrates.

This paper attempts to harmonise these two strands of literature by introducing the concept of 'programmable sovereignty', which we define as soft governance anchored by enough hard-stack control to make governance commitments credible and mitigate dependence. Programmable sovereignty is thus a hybrid position, creating a minimal sovereign fallback so that unilateral withdrawal is too costly for the N-2 countries and other states retain leverage for negotiation.

A natural question follows directly from the N-2 diagnosis: if infrastructural dependence creates such vulnerabilities, how can pooling governance of AI resources across several non-owning states not simply distribute the same underlying vulnerability across a larger group, rather than resolving it? The answer is nuanced. Pooled compute does not eliminate dependence on the duopoly; it only makes that dependence reversible rather than absolute. A coalition that jointly governs even a modest, shared compute base – sufficient to migrate critical workloads, sustain a minimum viable evaluation and deployment capacity and continue operating during a disruption – faces a different choice than a state with no fallback at all: the coalition can absorb a unilateral withdrawal and negotiate its terms of re-entry, rather than losing access outright with no recourse, as occurred in the Anthropic episode⁵.

Section 4 shows that the economies strongest in potential programmable-sovereignty governance (the European Union, India and, in certain aspects, Japan and South Korea) are among the weakest in hard-stack production, while the economy with the strongest hard-stack capability (Taiwan, dominant in semiconductor fabrication) is among the weakest in governance capacity. This informs the paper's second concept, 'coordinated specialisation': a coalition architecture through which complementary N-2 states pool their comparative advantages in a shared governance framework, rather than each attempting to replicate the full duopoly stack. Two scorecards are constructed for seven major economies (the US, China, the EU, Japan, South Korea, Taiwan and India): a hard-sovereignty scorecard spanning eleven layers of the AI production stack; and a programmable-sovereignty scorecard spanning seven governance levers.

Based on these scores, we argue that coordinated specialisation is the only feasible option available to the N-2 states. Section 5 expands on the concept of coordinated specialisation and outlines criteria for membership of a coordinated-specialisation partnership. Section 6 examines three existing institutional precedents: the EU's EuroHPC computing initiative, India Stack's digital public infrastructure and the ASEAN Digital Economy Framework Agreement. Each of these demonstrates a necessary component of coordinated specialisation while falling short of the fully integrated architecture this paper proposes. Section 7 develops the resulting policy implications and section 8 concludes.

2. The AI duopoly and the N-2 problem

2.1 Two full-stack sovereignty strategies

The US and China pursue different but equally exclusionary AI sovereignty strategies. The US treats AI infrastructure as an instrument of national strategic competition: the Biden administration's 2025 AI Diffusion Framework sought to restrict the international diffusion of advanced chips and model weights. The second Trump administration rescinded the framework⁶, but preserved its underlying geopolitical logic in the subsequent America's AI Action Plan, which emphasises domestic infrastructure build-out and efforts to control strategic competitors⁷ (The White House, 2025). The overall message is that compute, frontier

⁵ For a critique of the soft-sovereignty literature, see Ashraf and Veneziano (2026), whose critique of 'strategic agency' frameworks for AI sovereignty applies most directly to governance claims unaccompanied by any material fallback.

⁶ See U.S. Department of Commerce press release of 13 May 2025, 'Department of Commerce Announces Rescission of Biden-Era Artificial Intelligence Diffusion Rule, Strengthens Chip-Related Export Controls', <https://www.bis.gov/press-release/departement-commerce-announces-rescission-biden-era-artificialintelligence-diffusion-rule-strengthens>.

⁷ See The White House press release of 23 July 2025, 'White House Unveils America's AI Action Plan', <https://www.whitehouse.gov/releases/2025/07/white-house-unveils-americas-ai-action-plan/>.

models and large-scale data infrastructure are instruments of state power, rather than purely commercial assets.

The Chinese strategy is different but equally full-stack. It involves state-directed investment in domestic platform champions, control over critical upstream inputs such as rare earth elements and refined critical minerals, comprehensive data governance under the 2021 Data Security Law and the Personal Information Protection Law⁸ and political oversight of generative AI applications (Sheehan, 2023). The objective, beyond developing capable AI systems, is to ensure that AI development remains aligned with state priorities and insulated from external chokepoints.

The difference between chokepoint capacity and stack-level autonomy is crucial. The Netherlands, through the lithography monopoly held by the firm ASML, controls a critical chokepoint in advanced semiconductor manufacturing. But this does not give the Dutch hard AI sovereignty because the Netherlands has no credible programme to develop domestic alternatives to the core building blocks of the AI stack – which would require both a critical position in one (or more) layer(s) and a credible programme to maintain independence across every other layer under adversarial conditions. Only the US and China have assembled the institutional, financial and industrial infrastructure to pursue this programme across the full stack.

The EU, India and other large economies occupy structurally different positions. The EU possesses regulatory weight, institutional capacity and an established legal tradition in technology governance, but depends heavily for its AI deployment on non-European cloud infrastructure, frontier foundation models and semiconductor supply chains. It thus has rules without infrastructure. Emerging and middle-income economies face a related problem: their populations will generate vast quantities of data and their public and private sectors will adopt AI at scale, but without a clear strategy, the economic value thus generated will principally accrue to firms and jurisdictions outside their control.

2.2 Why AI dependence exceeds monetary dependence

The structural position of N-2 states arguably has a parallel in the dependence on a hegemonic currency, in conditions in which alternatives are costly to construct and coordination between dependent states is difficult. This is an established problem in international political economy (Eichengreen, 2011; Cohen, 2018).

However, we argue that AI dependence is more consequential than currency dependence for three reasons:

- i) First, AI dependence is infrastructural in a stronger sense than monetary dependence. Cross-border payments require a reference currency, but this single function does not compare to the breadth of the embedding of AI in the routine operations of public administration, healthcare, logistics and defence. By the time AI dependence becomes visible to policymakers, it is typically already embedded in code, contracts, cloud architecture and procurement systems (Farrell and Newman, 2019, 2023; Drezner *et al*, 2021).
- ii) Second, switching costs are categorically higher. A state can redenominate trade contracts in an alternative currency at comparatively low cost, but it cannot rebuild, within any comparable timescale, the model, cloud and semiconductor stack on which its hospitals, courts and tax administration depend, especially if AI take-up is associated with a loss of human knowledge on how systems work.

⁸ See Supreme People's Procuratorate of the People's Republic of China, 'Data Security Law of the People's Republic of China', https://en.spp.gov.cn/2021-06/10/c_948426.htm, and 'Personal Information Protection Law of the People's Republic of China', https://en.spp.gov.cn/2021-12/29/c_948419.htm.

- iii) Third, the window for constructing credible alternatives is closing more rapidly than was the case in the monetary domain, because supply-chain commitments, governance norms and training-data pipelines are being locked in now, rather than over the decades during which the post-Bretton Woods monetary order consolidated.

A further distinction between currency dependence and AI dependence is relevant here. Dollar-denominated financial exclusion requires multilateral coordination through banking systems and cannot be targeted with precision at a single firm, institution or individual without broad collateral effects. AI access withdrawal can, by contrast, be implemented unilaterally by a single company and targeted with the precision of a software access control. A nation, company or individual user can be cut off within minutes, without the collateral effects that multilateral financial exclusion entails. Section 4 quantifies the magnitude of AI dependence.

3. AI sovereignty: two literatures and a synthesis

3.1 The hard-sovereignty literature: sovereignty as productive autonomy

The dominant scholarly treatment of strategic power from AI treats sovereignty as a function of state-building. Lee (2018) set the terms for this literature by identifying AI development as a two-power race between the US and China, anticipating the duopoly structure underlying our N-2 argument. Buchanan (2020) decomposed AI capability into compute, algorithms and data, which is arguably the conceptual origin of every subsequent AI 'stack' framework. Miller (2022) isolated semiconductor hardware as the single binding constraint on the entire stack, and recast advanced chips as a strategic chokepoint comparable to oil.

Scharre (2023) widened Buchanan's triad by adding talent and institutions as a fourth domain. Ding (2024) then argued that diffusion capacity, or a society's capacity to adopt and absorb a general-purpose technology at scale, determines long-run technological power more reliably than where the initial invention comes from. This shifts the analytical question from who builds first to who deploys and governs most effectively (Ding, 2024). Sastry et al (2024) argued that compute is privileged in the analysis because it is detectable, excludable and quantifiable in a way other inputs are not. In this respect, the hardware layer is not just an input to be built but a lever that can be governed even by states that do not own it. This anticipated the soft-sovereignty literature discussed in section 3.2, and is the conceptual hinge for this paper's proposals. Within the hardsovereignty literature, the Stanford HAI AI Index (Stanford HAI, 2025, 2026) and the Epoch AI datasets⁹ supply the empirical infrastructure used in section 4 to measure national capability in compute, models, investment and talent.

Overall, the hard-sovereignty literature reaches a pessimistic conclusion for every N-2 economy: sovereignty over AI, defined as productive autonomy, requires fabrication capacity, frontier models and compute infrastructure that no N-2 state alone can realistically assemble.

3.2 The soft-sovereignty literature: sovereignty as governance capacity

A second and largely separate body of scholarship locates AI and digital sovereignty in governance rather than production. Floridi (2020) reframed the question as one of who shapes the digital environment and argued for a hybrid model of control in place of full-stack ownership. Bradford (2023) demonstrated empirically that the EU exercises substantial regulatory power through a rights-based regulatory model whose standards diffuse extraterritorially, rather than by out-building its technological rivals. Diffusion of standards happens because global firms find it more efficient to comply with the most stringent jurisdiction than to fragment products by market; regulation thus functions as a form of reach rather than a substitute for production. O'Hara and Hall (2018) focused on values instead of

⁹ See epoch.ai, 'Data on AI', <https://epoch.ai/data>.

specific rules, identifying four competing governance philosophies: a Silicon Valley open internet, a Brussels rights-based internet, a Washington commercial internet and a Beijing paternal internet. They and showed that the resulting contours of the global digital order are determined by these governance choices as much as by underlying engineering capability. For the EU, the General Data Protection Regulation (GDPR, Regulation (EU) 2016/679), for example, has extraterritorial reach operates because global firms find compliance cheaper than losing access to the EU market. The soft-sovereignty logic of the 'Brussels effect' rests on this assumption.

This literature offers N-2 states a better starting point than the hard-sovereignty tradition, because governance capacity, unlike fabrication capacity, does not require the absolute scale of duopoly investment to be matched. However, soft sovereignty is an incomplete foundation when pursued in isolation from any material capability (see section 3.3).

Among the recent literature relevant for our proposals, Fratini et al (2024) provided a systematic comparative review of competing digital-sovereignty models, against which the hard- and soft-sovereignty distinction can be situated. Hummel et al (2021) supplied the standard review of data sovereignty, the layer of the stack on which much of the soft-sovereignty literature's governance claims are built. Srivastava and Bullock (2024) underlined the claim that sovereignty has long been a contested rather than a settled category, even in its classical, territorial form.

Most directly relevant to our argument, however, is Ashraf and Veneziano's (2026) critical review of "strategic agency" in relation to AI sovereignty. They argue that frameworks, structurally similar to the soft-sovereignty literature reviewed here, locate sovereignty in a state's capacity to act strategically even when technological interdependence is irreversible. According to Ashraf and Veneziano (2026), such frameworks underspecify how infrastructural power structurally binds the strategic choices available to a dependent state. Such frameworks also fail to account for coercion pathways that become visible only under disruption (the dynamic illustrated by the Anthropoc episode). The minimal hard-stack component built into our definition of programmable sovereignty is intended specifically to address the disruptions emerging from coercion, rather than to repeat the vulnerability it identifies in governance-only accounts of sovereignty.

3.3 Programmable sovereignty: a bridging definition

The hard- and soft-sovereignty literatures rarely overlap: the former specifies what N-2 states cannot build, while the latter specifies what they might govern, without addressing whether governance alone is sufficient when the governed systems are owned by external and potentially adversarial parties. We seek to integrate the two traditions through a single concept: programmable sovereignty.

We define programmable sovereignty as soft governance anchored by a minimal degree of hard-stack control, above all pooled compute capacity. The hard-stack element is sufficient to make governance commitments credible and dependence reversible. Several elements of this definition require emphasis.

Programmable sovereignty is a spectrum between full dependence and full autonomy. It is distinct both from hard sovereignty (productive autonomy across the stack, currently approximated only by the US and China) and from soft sovereignty exercised alone (governance capacity unsupported by any hard-stack fallback). A state with governance levers but no minimal hard-stack capacity (no pooled compute, no sovereign evaluation infrastructure, no fallback routing option) thus does not possess programmable sovereignty; it possesses governance capacity that remains vulnerable to unilateral withdrawal by the party that owns the underlying systems. For example, when the US ordered the withdrawal of two Anthropoc frontier models (section 1) from the EU and elsewhere, the EU regulatory toolkit, including the

AI Act (Regulation (EU) 2024/1689) and the Digital Markets Act (Regulation (EU) 2022/1925), were of little use. This was not because these instruments are ineffective in their intended domain, but because they are market-access instruments designed to govern firm behaviour, rather than to constrain a foreign government's unilateral security decisions about its own companies.

The logic of programmable sovereignty is to remain embedded within global AI supply chains¹⁰ while maintaining a sufficient minimal fallback capacity, so that the dependent relationship becomes negotiated and reversible rather than unilateral and hostage-like. The minimal hard-stack component is thus like a spare wheel: a state that can reach the next point of exit should a current supplier default can afford to remain dependent in ordinary circumstances, can negotiate from a credible floor rather than from a cliff edge and can sustain a genuine, mutually beneficial co-dependence with the duopoly, rather than subordination to it.

3.4 Seven mechanisms

Programmable sovereignty can be implemented through a set of seven governance and operational mechanisms that allow a state or a coalition of states to direct, constrain and assure the behaviour of AI systems within their jurisdictions, regardless of where the underlying components were produced or are owned.

The mechanisms are:

- 1) Policy as code: the embedding of compliance requirements, such as data-residency or data-localisation rules, audit trails and safety thresholds, directly into AI systems and public procurement processes, rather than relying on after-the-fact litigation (Veale and Borgesius, 2021; Mariniello, 2026).
- 2) Routing control: the capacity to specify which models may process which categories of data, under which jurisdictional conditions and with which fallback options, so that sensitive workloads in defence, healthcare, taxation, identity management and the judiciary can be confined to trusted environments in which the compute provider functions as an operational intermediary subject to jurisdictional rules (Heim et al, 2024).
- 3) Evaluation capacity: the independent technical ability to test AI systems for safety, security, bias and reliability. This has started to be institutionalised through national AI Safety Institutes in several jurisdictions (Anderljung et al, 2023).
- 4) Cross-border data compacts: trusted bilateral or multilateral arrangements that permit data to flow across borders without the originating jurisdiction surrendering its rights to, or value derived from, that data. Examples include Japan's Data Free Flow with Trust framework and India's Data Empowerment and Protection Architecture (DEPA) consent architecture (Nilekani and Shah, 2015).
- 5) Reliability governance: the monitoring of what we term 'reliability debt', which is the accumulation of undetected operational risk that arises when AI systems are deployed faster than they can be tested, audited and corrected.
- 6) Pooled compute: the minimal hard-stack substrate (section 3.3). This could involve the joint procurement and shared governance of computing infrastructure, as done in the EU's EuroHPC initiative (see section 6.1).
- 7) Market leverage: the use of combined market size to require global AI providers to meet coalition-determined standards as a condition of market access. This mechanism is

¹⁰ In other words, continuing to procure advanced chips, deploy foreign-built foundation models and exchange data and services across borders, because the efficiency gains from interdependence exceed the costs of autarky in most cases.

negligibly impactful when exercised by any single N-2 economy (as shown by Bradford (2020, 2023), in relation to the 'Brussels effect') but potentially decisive when exercised at coalition scale. This mechanism gives force to the preceding six: rules that cannot be enforced against a non-compliant global provider are recommendations rather than governance.

Consistent with the definition in section 3.3, points 1 to 7 outline the minimum viable form of programmable sovereignty. It would preserve meaningful strategic optionality, but would not be sufficient to eliminate every vulnerability that follows from non-ownership of the full stack. Section 5 develops 'coordinated specialisation' as the mechanism through which N-2 states can move beyond this minimum viable form towards a more robust collective capability.


4. Measuring hard and programmable sovereignty

This section provides data to support the conceptual distinction developed in section 3.

4.1 State of play


To better gauge how far the rest of the world is from the US or China when it comes to the AI stack, we provide some key indicators. On innovation, the US and China clearly dominate the number of patents filed in AI-related fields to AI, with Europe, Japan and South Korea trailing far behind (Figure 1). On AI compute, the US clearly dominates (Figure 2) but China is making great strides. China also dominates the upstream part of the AI stack, including critical raw materials (Figure 3). In the same way, China has access to more and cheaper energy, which is crucial for faster AI adoption (Figure 4).

Figure 1: Numbers of AI related patents




Source: Bruegel based on WIPO.

Figure 2: Share of global AI GPU-cluster performance (% of total), 2026




Source: Bruegel based on Epoch AI. Note: AI GPU-cluster performance measures the total theoretical computing capacity of a hardware facility (a supercomputer or high-performance computing cluster) equipped with AI accelerators such as NVIDIA GPUs or Google TPUs.

Figure 3: Critical mineral processing concentration, 2020 and 2024



Source: García Herrero et al (2023).

Figure 4: Total installed electricity capacity (GW)



Source: García Herrero and Mu (2025).

Table 1 measures hard sovereignty across eleven layers of the AI technology stack (throughout this paper, we use the terms hard sovereignty and productive autonomy interchangeably, both denoting the degree to which a state builds, owns and can reproduce the stack, rather than governs it). Table 2 measures programmable sovereignty's constituent governance levers, in line with the seven mechanisms set out in section 3.4. Both tables classify seven major economies (the US, China, EU, Japan, South Korea, Taiwan and India) as 'leader', 'strong', 'emerging' or 'minimal' against the criteria. Classifications are a qualitative synthesis of the primary sources documented in the notes to each table, rather than a single composite index; we report them at this level of precision because the underlying evidence base, while extensive, does not support finer cardinal measurement across such a heterogeneous set of indicators.

Tables 1 and 2, read together show that hard-stack production capacity and programmable-sovereignty governance capacity are distributed in a complementary fashion across the seven economies examined.

The US and China score highly across most rows of Table 2, but, consistent with the definition in section 3.3, their position does not constitute programmable sovereignty as we define it: their governance capacities derive from owning the stack outright (the US) or directing it comprehensively through the state (China), not from governing a stack they do not own. Both economies, in other words, possess hard sovereignty. The scores in Table 2 for the US and China should be read as a description of how full-stack owners exercise comprehensive control, not as evidence that they have constructed a programmable-sovereignty solution to a dependence problem they do not face.

The EU is the clearest instance of programmable sovereignty's target condition and its central limitation. The EU is classified as a 'leader' or 'strong' for six of seven levers in Table 2, grounded in verifiable institutional facts: the AI Act embeds conformity assessment and post-market monitoring by design into high-risk systems; the EU AI Office, established in February 2024, provides institutionalised evaluation capacity for general-purpose models; the EU NIS2 Directive (network and information systems; Directive (EU) 2022/2555) imposes resilience requirements on critical digital infrastructure; and near-universal global compliance with the GDPR provides the clearest evidence of Brussels effect market leverage in the digital domain (Bradford, 2020). Yet, for nine of 11 layers in Table 1, the EU is classified as having 'minimal' or 'emerging' status. It is a 'leader' only for equipment and lithography, and this rests almost entirely on a single firm, ASML. This combination is precisely what our definition predicts will be vulnerable: governance leadership unsupported by a sufficient hard-stack fallback. The EU at present therefore exercises soft sovereignty extensively but programmable sovereignty only

Table 1: Hard AI sovereignty (who builds and owns the stack)

■ Leader ■ Strong ■ Emerging / catching up ■ Minimal

Layer of the stack	US	China	EU	Japan	Korea	Taiwan	India
Critical materials ¹¹	Emerging	Leader	Emerging	Emerging	Minimal	Minimal	Emerging
Energy / power ¹²	Strong	Leader	Emerging	Minimal	Minimal	Minimal	Emerging
Equipment & lithography ¹³	Strong	Emerging	Leader*	Strong	Minimal	Minimal	Minimal
Fabrication (foundry) ¹⁴	Emerging	Emerging	Minimal	Emerging	Strong	Leader	Emerging
Memory / HBM ¹⁵	Strong	Emerging	Minimal	Minimal	Leader	Minimal	Minimal
AI accelerators (design) ¹⁶	Leader	Emerging	Minimal	Minimal	Minimal	Minimal	Minimal
Compute clusters ¹⁷	Leader	Strong	Minimal	Emerging	Minimal	Minimal	Emerging
Software ecosystem ¹⁸	Leader	Emerging	Minimal	Minimal	Minimal	Minimal	Minimal
Foundation models ¹⁹	Leader	Strong	Emerging	Minimal	Minimal	Minimal	Emerging
Private AI investment ²⁰	Leader	Strong	Emerging	Emerging	Emerging	Minimal	Emerging
Research & talent ²¹	Leader	Leader	Strong	Emerging	Emerging	Minimal	Strong

partially, pending the build-out of its pooled-compute capacity under EuroHPC and the AI Factories initiative (see section 6).

Table 2: Programmable sovereignty (who can govern AI without building it)

Lever	US	China	EU	Japan	Korea	Taiwan	India
Build rules into the system ²² <i>policy as code</i>	Emerging	Strong	Strong	Emerging	Strong	Emerging	Strong
Control where data & decisions live ²³ <i>routing control</i>	Strong	Leader	Leader	Emerging	Emerging	Minimal	Strong
Check what vendors claim ²⁴ <i>evaluation capacity</i>	Strong	Strong	Strong	Emerging	Emerging	Minimal	Emerging
Share data without losing control ²⁵ <i>data compacts</i>	Emerging	Emerging	Leader	Strong	Emerging	Minimal	Leader
Stop hidden risk piling up ²⁶ <i>reliability governance</i>	Emerging	Strong	Strong	Emerging	Emerging	Minimal	Emerging
Pool computing as a shared utility ²⁷	Strong	Strong	Emerging	Emerging	Strong	Emerging	Emerging
Use the size of your market ²⁸ <i>Brussels Effect</i>	Leader	Strong	Leader	Emerging	Emerging	Minimal	Strong

India occupies a comparable but smaller-scale position: broadly emerging-to-minimal across the hard-stack layers in Table 1, but a leader on data compacts (via the DEPA consent architecture and Account Aggregator framework; see Prajapati et al, 2020) and strong across several other governance levers. Japan and South Korea occupy more specialised niches within the governance scorecard: Japan as the principal architect of the Data Free Flow with Trust

¹¹ IEA (2025); Powers-Riggs et al (2025); Bruno Venditti, 'Charted: China's Grip on Critical Mineral Refining', Visual Capitalist, 19 May 2026, <https://www.visualcapitalist.com/charted-chinas-grip-on-critical-mineral-refining/>. China refines ~99 percent of gallium and ~83 percent of germanium and leads refined output for 19 of the 20 minerals tracked by the IEA.

¹² Haag (2025); Alan Weissberger, 'IEEE ComSoc Technology Blog: February 2026', IEEE, 25 February 2026, <https://techblog.comsoc.org/2026/02/>. China ~3,200 GW installed capacity vs ~1,293 GW (US) and ~1,125 GW (EU); +429 GW added in 2024.

¹³ See eg MarketsandMarkets press release of 21 November 2025, 'Semiconductor Manufacturing Equipment Industry worth \$344.36 billion by 2032', <https://www.marketsandmarkets.com/PressReleases/semiconductor-manufacturingequipment.asp>. ASML has >95 percent of EUV lithography; the top five vendors hold 56 percent to 66 percent of all equipment revenue.

¹⁴ TrendForce foundry rankings 2025; see TrendForce press release of 12 March 2026, 'AI Demand Drives 4Q25 Global Top 10 Foundries Revenue Up 2.6% QoQ; Samsung Gains Share and Tower Moves Up in Rankings, Says TrendForce', <https://www.trendforce.com/presscenter/news/20260312-12965.html>. TSMC ~70 percent share, Samsung ~7 percent, SMIC ~5 percent, GlobalFoundries ~4 percent.

¹⁵ TrendForce and general industry reporting. SK Hynix and Samsung (Korea) lead high bandwidth memory; Micron (US) is third.

¹⁶ Kyle Chan, 'Competing AI strategies for the US and China', *Commentary*, 16 April 2026, Brookings, <https://www.brookings.edu/articles/competing-ai-strategies-for-the-us-and-china/>; Chang et al (2025). Nvidia, AMD and Google lead; Huawei Ascend is ~one generation behind.

¹⁷ Epoch AI, AI supercomputers dataset (see footnote 9); Stanford HAI (2025). The US hosts ~75 percent of benchmarked GPU-cluster performance vs China's ~14 percent, and has ~9× China's AI compute.

¹⁸ Chang et al (2025); industry analysis. Nvidia's CUDA remains the dominant developer moat; Huawei's CANN is China's main alternative.

¹⁹ Stanford HAI (2025). The US produced 40 notable models in 2024 vs China's 15 and Europe's three, with Chinese models reaching near benchmark parity.

²⁰ Stanford HAI (2025). US private AI investment reached \$109.1 billion in 2024 vs China's \$9.3 billion and the UK's \$4.5 billion.

²¹ Stanford HAI, (2025). China leads in publication and patent volume; the US leads in highly-cited research and top-talent concentration.

²² EU AI Act and European Commission, 'The General-Purpose AI Code of Practice', <https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai>; India DPDP Act & Rules 2025 with the DEPA consent rails; China Generative AI Services Management Measures; Korea AI Basic Act.

²³ EU Data Act (Regulation (EU) 2023/2854, applicable 12 Sep 2025) and sovereign-cloud requirements (Ingemarsson, 2026); China data localisation under PIPL and the Data Security Law (see footnote 8); India DPDP data-residency rules.

²⁴ EU AI Office and Testing & Experimentation Facilities (European Commission); the US, Korean and Japanese AI Safety Institutes; China's model registration and security review.

²⁵ EU Data Governance Act (Regulation (EU) 2022/868); European Commission, 'Common European data spaces', <https://digital-strategy.ec.europa.eu/en/policies/data-spaces>; India's Account Aggregator/DEPA framework; Government of Japan Digital Agency, 'Data Free Flow with Trust (DFFT)', <https://www.digital.go.jp/en/policies/dfft>.

²⁶ EU AI Act high-risk obligations; India DPDP breach reporting and the Data Protection Board; Korea AI Basic Act high-impact risk-management duties.

²⁷ EU EuroHPC AI Factories and InvestAI gigafactories (see section 6.1); India's IndiaAI Mission shared GPUs (34,000 now, 100,000 targeted by end-2026); Korea's sovereign-cloud and GPU build-out.

²⁸ GDPR, the Digital Markets Act (Regulation (EU) 2022/1925) and Digital Services Act (Regulation (EU) 2022/2065), with EU AI Act fines up to €35 million or 7 percent of global turnover; India Stack's populationscale rails (UPI, Aadhaar) disciplining private firms.

framework, agreed at the 2019 G20 Osaka summit, though its domestic AI Promotion Act remains deliberately soft law with no associated financial penalties; South Korea through its AI Basic Act, which took effect in January 2026 as, after the EU AI Act, the world's second comprehensive national AI statute with explicit extraterritorial reach, paired with a substantial sovereign-cloud investment programme. Taiwan is the empirical mirror image of the EU: in a leader position on a single, globally indispensable hard-stack layer (semiconductor fabrication), combined with minimal status across every governance lever in Table 2²⁹.

The substantive implication of this non-overlap is that no single N-2 economy here possesses the combination of governance capacity and minimal hard-stack fallback that the definition of programmable sovereignty requires. The EU, India, Japan and South Korea collectively possess the governance capacity; Taiwan and South Korea collectively possess critical fabrication and memory capacity that none of the governance leaders possesses individually. This is not a coincidental empirical pattern. It is the structural condition that coordinated specialisation, as an institutional design, is built to leverage.

5. Coordinated specialisation: from individual insufficiency to coalition sufficiency

Section 4's empirical finding that governance capacity and hard-stack production capacity are concentrated largely in different, non-overlapping sets of economies supports this paper's core proposal. If no single N-2 economy can independently satisfy the definition of programmable sovereignty, as developed in section 3.3, the relevant policy question is whether a coalition of complementary economies can satisfy it collectively. We argue that it could, through a coordinated-specialisation architecture in which each participating state contributes its comparative advantage under a shared governance framework.

The intellectual lineage of this argument runs through the industrial-policy literature on coordinated public action in strategic sectors characterised by high fixed costs, network effects and first-mover advantages that markets alone do not overcome efficiently (Mazzucato, 2013; Rodrik, 2022; Aiginger and Rodrik, 2020; Cherif and Hasanov, 2019). One historical precedent is the European aerospace consortium Airbus, which did not emerge from any single European state's attempt to construct a full national aerospace industry capable of matching Boeing. Rather, it emerged from national specialisation – French-built airframes, British-built wings, German-built fuselage sections, Spanish-built tail assemblies – coordinated within a shared governance structure and underwritten by a combined market scale that no single participating state could have provided alone (García-Herrero and Martens, 2026). Three features of the Airbus case could transfer directly to AI: deep functional specialisation rather than duplicated effort, governance arrangements agreed prior to, and conditioning, joint investment, and patient public capital sustained over decades, rather than calibrated to short-run political cycles.

We note, without resolving the question here, that coordinated industrial strategies of this kind have a mixed empirical record. The literature on industrial policy identifies factors in systematic failure – binding state-capacity constraints, weak enforcement of conditionality on participating firms or states, and rent-seeking among coalition members – that are not unique to AI (Bulfone, 2023; Juhász et al, 2024). On the available comparative evidence, the Airbus success was not a representative outcome among coordinated multi-state industrial ventures. The coalition design we propose should accordingly be read as institutionally feasible rather than institutionally guaranteed. The emphasis on sequencing governance before infrastructure

²⁹ A note on Japan's hard-stack position is warranted here. Table 1 classifies Japan as 'Strong' on equipment and lithography, primarily on the strength of Tokyo Electron (TEL), its leading semiconductor equipment manufacturer. This classification understates Japan's chokepoint position because it omits photoresist chemicals: Japan accounts for approximately 90 percent of global supply of the photosensitive coatings without which lithography cannot function. Photoresists represent a separate and equally critical material dependency in the semiconductor supply chain, which Japan could in principle leverage independently of its equipment position. We note this as both a factual addition to Japan's hard-stack profile and as an illustration of the more general point that the stack's critical dependencies extend beyond the eleven layers captured in Table 1: the layer taxonomy is analytically useful but should not be read as exhaustive.

is intended specifically to mitigate common causes of coordinated industrial projects underperforming their design (see section 7.4).

Our call for a partnership when it comes to AI sovereignty is similar to calls made by others. Abecassis et al (2025) proposed a coalition of “AI middle powers” that would combine complementary national capabilities with shared governance.

We go further by specifying some desirable characteristics of the coalition’s architecture – in other words, why each prospective member would join such a coalition, rather than seek bilateral arrangements with the duopoly or try to achieve the full AI stack on its own. The incentive to join differs by member type. For the EU, the coalition would fill a structural gap that no EU policy instrument alone can fill. It would provide hard-stack capacity (compute, fabrication and memory) that Brussels effect regulation cannot replace. When it comes to attracting others to the partnership, the EU offers a \$20 trillion market for AI deployment.

For manufacturing specialists such as Taiwan and South Korea, coalition membership would provide governance standing and rules-layer credibility that they lack in their current positions as indispensable production nodes. For India, the coalition would provide access to pooled compute that India alone cannot build. At the same time, India would contribute the data-governance architecture and population-scale digital infrastructure to make the coalition’s data-compact mechanism credible at the required scale. For potential Gulf-state participants, the coalition would offer a multilateral governance framework that would convert patient capital into a recognised strategic asset, rather than a commercial investment subject to duopoly discretion.

Coalition membership would be governed by four conditions:

- 1) A prospective member must bring a demonstrated comparative advantage in at least one layer of the stack – hard or governance – that the coalition needs and does not already hold, so that it adds capacity rather than duplicating it.
- 2) That contribution must be complementary rather than overlapping: the systematic non-overlap documented in section 4 is the selection criterion, not an incidental by-product of it.
- 3) The member must accept the shared governance framework, involving data residency, evaluation and audit protocols, and intellectual-property and benefit-sharing rules, as the condition of access to pooled compute and other joint capacity.
- 4) Participation must be politically feasible: if a state’s security commitments make formal membership untenable, as with Taiwan, it can be integrated through market and open-fab arrangements, rather than excluded outright.

Applied to the present case, the coordinated-specialisation architecture would aggregate the EU’s governance leadership (Table 2) with the complementary hard-stack specialisations of Taiwan (fabrication), South Korea (memory and, increasingly, sovereign compute), Japan (data-governance architecture) and India (population-scale digital public infrastructure and a rapidly expanding sovereign compute base), with sovereign capital as a plausible source of the patient, multi-decade financing that the Airbus precedent indicates is necessary. None of these economies alone, according to our scorecards, could independently satisfy the definition of programmable sovereignty developed in section 3. Pooled within a shared governance architecture, however, the coalition would collectively possess both the governance levers and the minimal hard-stack fallback – above all, jointly-governed compute capacity – that the definition requires, at a scale that no individual member can achieve.

This proposed coordinated-specialisation architecture raises a further question that the comparative-advantage logic alone does not resolve: what sustains coalition membership once formed, and what prevents defection once an individual member's own capability matures? A state that builds sufficient domestic compute capacity, for instance, might have an incentive to withhold it from coalition-wide pooling once its own minimal fallback is secure, free-riding on the market leverage that the coalition's combined scale continues to provide. We do not offer a full game-theoretic treatment of coalition stability here; this is properly a question for the collective-action and alliance literatures (the conditionality-enforcement problem identified in Bulfone et al, 2024). We flag it as an open question for future work, rather than treating coalition formation as self-sustaining once comparative advantages are identified. The governance before- infrastructure sequencing recommended in section 7.4, and the shared-metrics proposal in section 7.5, would be partial but insufficient institutional responses to this defection risk: binding commitments and observable compliance are the standard tools for sustaining cooperation when there are incentives to free-ride.

The political feasibility of coalition participation would vary significantly among prospective members. Taiwan's case is the most acute: its security architecture depends on the US, making it politically difficult – possibly impossible in the near term – to participate in mechanisms that are framed explicitly as coordinating leverage against US AI export decisions. A realistic coalition architecture may therefore need to begin without Taiwan as a formal member, treating Taiwan's fabrication capacity as accessible through market relationships and open-fab provisions, while leaving deeper integration for when the geopolitical context allows.

Japan and South Korea face related but less-extreme constraints: both have strong US security ties that counsel against framing coalition participation in adversarial terms. This reinforces our recommendation (section 7) that the coalition be presented as a framework for structured interdependence, capability diversification and mutual governance, rather than as an AI non-alignment movement against the US (or China). Its initial membership may begin with a nucleus of EU countries and signatories of the Comprehensive and Progressive Agreement for Trans-Pacific Partnership (CPTPP), before extending to Korea and India.

6. Institutional precedents and their limits

This section examines three institutional arrangements that each demonstrate a necessary component of the coordinated-specialisation architecture proposed in section 5. Read together, these three cases show the institutional feasibility of coordinated specialisation – but also the gap that remains to get there. EuroHPC is an example of pooled infrastructure and India Stack of programmable architecture, while the ASEAN framework is a model of regional governance coordination.

6.1 EuroHPC: pooled infrastructure under shared governance

EuroHPC is the clearest existing proof of concept for the minimal hard-stack substrate that programmable sovereignty requires, according to our definition. The EuroHPC Joint Undertaking³⁰ pools EU, EU member state and private resources to procure and operate world-class computing systems under jointly negotiated access rules rather than national ownership. The systems include LUMI (Finland), Leonardo (Italy) and JUPITER (Germany; Europe's first exascale system). Building on this base, the European Commission's AI Factories initiative, announced as part of the 2024 AI Innovation Package, federates and upgrades these systems into a network of AI-optimised hubs offering compute access to start-ups, small and medium enterprises, and researchers. EuroHPC establishes that infrastructure pooling under shared governance is operationally feasible at continental scale. Its frontier-capable tier, however, remains prospective rather than operational, and its governance has not, at time of writing,

³⁰ The European High Performance Computing Joint Undertaking. A joint undertaking is a form of partnership defined in EU law, established to work on a high-priority common goal. See https://www.eurohpc-ju.europa.eu/index_en.

been extended to the full set of mechanisms, including evaluation, routing and reliability governance, that our definition of programmable sovereignty presupposes.

The current compute facilities under EuroHPC in its current form, fall well short of the threshold at which these mechanisms become operationally significant. EuroHPC represents from 1 percent to 5 percent of EU AI compute demand (Ingemarsson, 2026); this is clearly insufficient to sustain even a fraction of critical public-sector workloads during a disruption scenario. The AI Factories and Gigafactories trajectory³¹, if completed at the planned scale³², would move towards a more meaningful threshold over three to five years. In the interim, the EuroHPC precedent should be read as establishing governance feasibility, showing that pooled infrastructure with shared multinational governance can be maintained operationally, rather than as current capability sufficiency.

6.2 India Stack: programmable architecture and its limits

India Stack, a digital infrastructure that comprises Aadhaar, the Unified Payments Interface, the Account Aggregator framework and the Open Network for Digital Commerce, demonstrates that population-scale digital public infrastructure, designed in and governed from the Global South, can operate at a level of technical sophistication and reliability comparable to any equivalent system globally. Its architectural significance for the present argument is that a public digital framework, on which private firms, financial institutions and public agencies can build, can be constructed without requiring state monopoly over the underlying digital services. Public priorities, including identity inclusion, low-cost payment access, consent-based data sharing and cross-platform interoperability, are encoded directly into the architecture. India Stack is accordingly a powerful institutional precedent for the kind of programmable architecture that coordinated specialisation requires. India's 'leader' classification on data compacts (Table 2) reflects this operational reality.

The move from digital public infrastructure to AI sovereignty proves harder. The IndiaAI Mission, targeting 100,000 GPUs of sovereign compute capacity³³ by the end of 2026, with approximately 34,000 already deployed, shows both what is achievable and what remains structurally difficult. India possesses substantial talent, with the highest AI skill-penetration rate globally, at 2.8 times the global average (Stanford HAI, 2024), alongside proven digital-infrastructure design capability and a domestic market of considerable scale, but it remains dependent on foreign-manufactured chips, foreign cloud infrastructure and foreign-built foundation models for the most capable tier of AI systems. AI sovereignty is thus structurally more demanding than digital-services sovereignty: India Stack demonstrates what programmable architecture looks like when correctly designed, while the IndiaAI Mission shows why a coalition-level strategy, rather than a purely national one, is needed.

6.3 The ASEAN Digital Economy Framework: regional governance as a partial building block

³¹ The AI Factories and Gigafactories comprise the EU's two-stage strategy to build sovereign, highperformance computing infrastructure and cut dependence on foreign cloud giants. See European Commission, 'AI Factories', <https://digital-strategy.ec.europa.eu/en/policies/ai-factories> and 'AI Gigafactories', https://commission.europa.eu/topics/competitiveness/competitiveness-coordination-tool/projects/ai-gigafactories_en.

³² EuroHPC is backed by a €7 billion joint budget for 2021-2027; is intended to reach exascale capacity (exceeding 1 ExaFLOP/s, 1 quintillion calculations per second), and aims to establish a federated European network of 9+ operational supercomputers (spanning petascale, pre-exascale and exascale systems) integrated with quantum processors and specialised AI Factories.

³³ For the IndiaAI Mission, target compute capacity is measured using three core criteria: 1) hardware count: the primary baseline metric measures the total number of specialised AI accelerators (such as NVIDIA H100, H200 or A100 GPUs) aggregated through public-private partnerships; 2) computational performance: capacity is evaluated in FLOPS (Floating-Point Operations Per Second); 3) service allocation: for researchers and startups, compute is operationalised and billed in GPU-hours, provided through a subsidised cloud marketplace.

The ASEAN Digital Economy Framework Agreement, negotiated and concluded during 2025 and 2026, aims to establish shared rules on digital trade, cross-border data flows and AI deployment across Southeast Asia. This framework would combine Singapore's regulatory sophistication, Indonesia's demographic scale, Vietnam's manufacturing base and Malaysia's expanding data-centre capacity. The agreement stops well short of the operational mechanisms required for coordinated specialisation, as per the architecture proposed in section 5. For example, the agreement does not pool compute capacity, establish joint model evaluation or impose collective procurement requirements at coalition scale. However, as an institutional building block for a more comprehensive future partnership, the agreement is significant.

7. Policy implications

The analysis developed in sections 3 to 6 supports a change in strategic posture for N-2 countries, rather than a catalogue of measures. The central implication is that no individual N-2 economy, the EU included, can secure programmable sovereignty alone. Attempting to do so by reconstructing the full hard stack would be both prohibitively costly and, based on the evidence set out in section 4, unnecessary, since the constituent capabilities exist already across a prospective coalition. The policy task is thus to convert dispersed and individually insufficient capabilities into a pooled and jointly-governed whole. In this respect, we set out five broad principles.

7.1 Lead with governance, not production

N-2 states retain genuine and immediately exercisable agency over the mechanisms that would underpin programmable sovereignty (section 3.4: policy as code, routing control, evaluation capacity, data compacts, reliability governance, pooled compute and market leverage). Policy should accordingly prioritise building and harmonising these levers, rather than subsidising full-stack production efforts. For the EU specifically, this implies deploying its regulatory leadership as coalition-scale leverage, rather than exercising it in isolation: the Brussels effect compounds in force when exercised through combined market and procurement weight across a coalition, and weakens when the EU regulates alone against systems it cannot itself replace. Evidence that this weakening is happening, with instances of major technology firms withdrawing AI models from the EU rather than adapting them for compliance³⁴, underscores the urgency of a partnership that can coordinate regulation.

7.2 Anchor governance with a minimal hard fallback

The June 2026 Anthropic episode (section 1) showed that AI governance commitments unsupported by any material capacity can be nullified within minutes. The implication, consistent with the definition developed in section 3.3, is not that the future coalition should attempt to construct a rival full stack, but that it should secure sufficient shared, jointly governed capacity, especially pooled compute and assured access to a small number of irreplaceable chokepoint assets. This would make its governance commitments credible and its dependence reversible. The objective is strategic optionality: the demonstrated ability to route critical workloads to trusted alternative infrastructure and to negotiate from a credible floor, rather than from crisis, while remaining embedded within global AI supply chains, rather than pursuing autarky.

³⁴ Major tech firms have chosen in some cases to withhold multimodal models rather than adapt to EU standards. Meta in 2024 opted out of deploying multimodal LLaMA models in Europe, while Apple during 2024-2026 withheld Apple Intelligence features from the EU market over conflicts between system-level AI privacy and the EU Digital Markets Act's interoperability mandates. See Katie Collins, 'Meta Follows in Apple's Footsteps by Restricting AI Releases in EU Countries', CNET, 18 July 2024, <https://www.cnet.com/tech/services-and-software/meta-follows-in-apples-footsteps-by-restricting-ai-releases-in-eu-countries/>.

7.3 Organise around demonstrated comparative advantage

The Airbus case (section 5) is instructive because no participant attempted by itself to construct the complete aircraft. In relation to AI, the EU's regulatory and lithography-equipment strengths, Taiwan's and South Korea's fabrication and memory leadership, Japan's data-governance architecture, India's population-scale digital public infrastructure and sovereign capital are complementary capabilities rather than competing claims to the same function. Each prospective coalition member should deepen and make available, under shared rules, its area of demonstrated comparative advantage, rather than dissipating scarce resources across stack layers over which it has no realistic prospect of leadership. Coordinated specialisation, not duplication, would convert individually insufficient capabilities into a collectively sufficient stack.

7.4 Put governance in place before infrastructure

Specialisation generates programmable sovereignty only when it is conducted in line with governance terms agreed in advance. These should cover data residency, evaluation and audit protocols, intellectual property and benefit-sharing arrangements, reliability obligations and dispute resolution. Without this sequencing, pooling capability would simply relocate dependence rather than resolving it: a state that supplies compute capacity under another coalition member's unilateral governance terms has exchanged one external master for another. Both the Airbus and EuroHPC precedents (sections 5 and 6.1) indicate that a shared governance framework, open to partners from the outset, should precede and help shape any joint infrastructure investment. The founding act of the proposed partnership should accordingly be an international agreement establishing this governance framework.

7.5 Make programmable sovereignty empirically measurable

The N-2 problem is dangerous in part because dependence accumulates off the radar screen of policymakers until it is activated. It is embedded in code, contracts and procurement decisions made well before the dependence becomes visible. The future coalition should accordingly commit to a small set of shared indicators. These could include the speed of trustworthy AI deployment within regulated public-sector applications, the share of critical workloads operating without completed independent evaluation or any sovereign fallback option, and the share of critical workloads for which no alternative internal to the coalition exists. With such shared metrics, programmable sovereignty would become empirically observable. Coalition states would have an evidentiary basis from which to assess whether their collective autonomy is increasing or eroding over time, and from that a research agenda for future quantitative work.

These five policy principles amount to a strategic posture rather than an institutional design. The instruments through which the partnership is realised – which body convenes it, how compute access is reciprocated among members, how evaluation capacity is pooled, how patient capital is mobilised – are matters for negotiation among the governments concerned. Various institutional designs could satisfy the underlying logic we establish in this paper: that programmable sovereignty is attainable for N-2 states hard-stack fallback, not soft governance alone; and that it must be constructed within a closing window because supply chains, governance norms and data-infrastructure commitments are presently being locked in.

8. Conclusion

This paper has argued that the organisation of AI production and governance has so far resulted in an N-2 world, in which only the US and China can credibly pursue full-stack technological sovereignty. For the rest of the world – the N-2 states – attempting to reconstruct the full stack is neither affordable nor achievable within any policy-relevant timeframe.

Our proposed alternative, programmable sovereignty, rests on a definition of governance anchored by a deliberately minimal degree of hard-stack control, above all pooled compute capacity. This should be sufficient to make governance commitments credible and dependence reversible. Programmable sovereignty is neither hard sovereignty (productive autonomy, available only to the US and China) nor soft sovereignty pursued alone (governance unsupported by any material fallback). It is a hybrid and is the paper's principal conceptual contribution to the literatures reviewed in section 3.

Our two scorecards (section 4) demonstrate that the capabilities required to achieve programmable sovereignty exist already, in aggregate, across a plausible coalition of N-2 economies. We also show that the economies strongest in governance capacity are weakest in hard-stack production capacity and vice versa. This non-overlap is the empirical foundation for our coordinated-specialisation proposal. It would pool, within a shared governance framework modelled on the European aerospace consortium Airbus, the EU's regulatory leadership, the fabrication and memory strength of Taiwan and South Korea, Japan's data-governance architecture, India's digital public infrastructure and sovereign capital.

Several qualifications to the our analysis should be stated. First, the mechanism by which pooled compute confers programmable sovereignty is, on closer inspection, dual rather than single. The first mechanism is the fallback: a coalition with jointly governed compute capacity can migrate open-weight and fine-tuned models to trusted infrastructure during a disruption, sustain minimum viable deployment for critical publicsector applications and negotiate from a credible exit option rather than from the edge of a cliff. This fallback is genuine but limited for closed-weight frontier models, which under the provider's control regardless of where inference infrastructure is located.

A second, and often more powerful mechanism, is indispensability (Schenk and Poitiers, 2026). Indispensability occurs when an individual economy (in this case a member of the coalition) supplies compute, energy or other critical inputs to the US and Chinese AI firms that the coalition also depends on. Such inputs can potentially be a chokepoint and therefore offer leverage. The longer-run strategic objective of the coalition's compute investment should therefore be understood from two angles: building the fallback capacity that makes dependence reversible, and reaching indispensability in some key aspects so that it can potentially use coercion against the existing duopoly. This can dilute the N-2 problem.

Caution is also warranted about the current effectiveness of market leverage as a governance mechanism. There is evidence that in for AI, the Brussels effect is showing signs of reversal rather than diffusion. Major technology firms have begun withholding AI models from the EU market rather than modifying them for EU compliance, and the EU has itself retreated from several elements of the AI Act under competitive pressure. If this trend continues, the premise underlying our seventh mechanism (section 3.4) – that firms will comply with coalition standards rather than exit the market – may not hold for the EU alone. Such a reversal would strengthen, rather than undermine, the case for coordination.

The scorecards in section 4 measure governance and production capacity but not the conditions under which a coalition forms and persists. We have flagged this collectiveaction question in section 5 as a matter for more formally modelled work. We do not address the normative question of whether a given exercise of governance capacity, by a coalition or by any of its individual members, is legitimate in the eyes of the populations whose data and behaviour it governs.

References

- Abecassis, A., J. Barry, I. Bello, Y. Bengio, A. Bergeaud, Y. Bonnet ... K. VijayRaghavan (2025) A Blueprint for Multinational Advanced AI Development, AI Governance Institute, University of Oxford, available at https://aigi.ox.ac.uk/wp-content/uploads/2025/11/blueprint_for_multinational_ai_development.pdf
- Aiginger, K. and D. Rodrik (2020) 'Rebirth of Industrial Policy and an Agenda for the Twenty-First Century', *Journal of Industry, Competition and Trade* 20: 189–207, available at <https://doi.org/10.1007/s10842-019-00322-3>
- Anderljung, M., J. Barnhart, A. Korinek, J. Leung, C. O'Keefe, J. Whittlestone (2023) 'Frontier AI Regulation: Managing Emerging Risks to Public Safety', mimeo, available at <https://arxiv.org/abs/2307.03718>
- Ashraf, A. and V. Veneziano (2026) 'AI Sovereignty Without Power? Reviewing Strategic Agency, Infrastructural Dependence and Digital Imperialism', *AI & Society*, available at <https://doi.org/10.1007/s00146-026-03116-4>
- Bradford, A. (2020) *The Brussels Effect: How the European Union Rules the World*, Oxford University Press, available at <https://doi.org/10.1093/oso/9780190088583.001.0001>
- Bradford, A. (2023) *Digital Empires: The Global Battle to Regulate Technology*, Oxford University Press, available at <https://doi.org/10.1093/oso/9780197649268.001.0001>
- Buchanan, B. (2020) *The AI Triad and What It Means for National Security Strategy*, Center for Security and Emerging Technology, available at <https://doi.org/10.51593/20200021>
- Bulfone, F. (2023) 'Industrial Policy and Comparative Political Economy: A Literature Review and Research Agenda', *Competition & Change* 27(1): 22–43, available at <https://doi.org/10.1177/10245294221076225>
- Bulfone, F., T. Ergen and E. Maggor (2024) 'The Political Economy of Conditionality and the New Industrial Policy', MPIfG Discussion Paper 24/6, Max Planck Institute for the Study of Societies, available at https://pure.mpg.de/rest/items/item_3612897_2/component/file_3612917/content
- Chang, W., R. Arcesati and A. Hmadi (2025) China's drive toward self-reliance in artificial intelligence: from chips to large language models, MERICS, available at https://merics.org/sites/default/files/2025-07/MERICS%20Report-AI_Stack_final.pdf
- Cherif, R. and F. Hasanov (2019) 'The Return of the Policy That Shall Not Be Named', Working Paper 2019/074, International Monetary Fund, available at <https://www.imf.org/en/publications/wp/issues/2019/03/26/the-return-of-the-policy-that-shall-not-be-named-principles-of-industrialpolicy-46710>
- Cohen, B.J. (2018) *Currency Statecraft*, University of Chicago Press, available at <https://press.uchicago.edu/ucp/books/book/chicago/C/bo29141097.html>
- Ding, J. (2024) *Technology and the Rise of Great Powers*, Princeton University Press, available at <https://press.princeton.edu/books/paperback/9780691260341/technology-and-the-rise-of-great-powers>
- Drezner, D.W., H. Farrell and A.L. Newman (eds) (2021) *The Uses and Abuses of Weaponized Interdependence*, Brookings Institution Press, available at <https://www.brookings.edu/books/the-uses-and-abuses-of-weaponized-interdependence/>

Eichengreen, B. (2011) *Exorbitant Privilege*, Oxford University Press

Farrell, H. and A.L. Newman (2019) 'Weaponized Interdependence', *International Security* 44(1): 42–79, available at https://doi.org/10.1162/isec_a_00351

Farrell, H. and A.L. Newman (2023) *Underground Empire*, Henry Holt

Floridi, L. (2020) 'The Fight for Digital Sovereignty: What It Is, and Why It Matters, Especially for the EU', *Philosophy & Technology* 33: 369–378, available at <https://doi.org/10.1007/s13347-020-00423-6>

Fratini, S., E. Hine, C. Novelli, H. Roberts and L. Floridi (2024) 'Digital Sovereignty: A Descriptive Analysis and a Critical Evaluation of Existing Models', *Digital Society* 3(3): 59, available at <https://doi.org/10.1007/s44206-024-00146-7>

García-Herrero, A., H. Grabbe and A. Kaellenius (2023) 'De-risking and decarbonising: a green tech partnership to reduce reliance on China', *Policy Brief* 19/2023, Bruegel, available at <https://www.bruegel.org/policy-brief/de-risking-and-decarbonising-green-tech-partnership-reducereliance-china>

García-Herrero, A. and B. Martens (2026) 'Europe needs a strategy to close the artificial intelligence compute gap', *Analysis* 13/2026, Bruegel, available at <https://doi.org/10.64153/HFQM9739>

García-Herrero, A. and H. Mu (2025) 'China can decarbonise the world – but even that won't fix its overcapacity problem', *Analysis* 34/2025, Bruegel, available at <https://www.bruegel.org/analysis/china-can-decarbonise-world-even-wont-fix-its-overcapacity-problem>

Haag, A. (2025) 'The State of AI Competition in Advanced Economies', *FEDS Notes*, Board of Governors of the Federal Reserve System, available at <https://doi.org/10.17016/2380-7172.3930>

Heim, L., T. Fist, J. Egan, S. Huang, S. Zekany, R. Trager, M.A. Osborne and N. Zilberman (2024) 'Governing Through the Cloud: The Intermediary Role of Compute Providers in AI Regulation', *White Paper*, Oxford Martin School, available at <https://www.oxfordmartin.ox.ac.uk/publications/governing-through-the-cloud-the-intermediary-role-of-compute-providers-in-ai-regulation>

Helleiner, E. (2014) *The Status Quo Crisis*, Oxford University Press, available at <https://doi.org/10.1093/acprof:oso/9780199973637.001.0001>

Hummel, P., M. Braun, M. Tretter and P. Dabrock (2021) 'Data Sovereignty: A Review', *Big Data & Society* 8(1), available at <https://doi.org/10.1177/2053951720982012>

IEA (2025) *Global Critical Minerals Outlook 2025*, International Energy Agency, available at <https://www.iea.org/reports/global-critical-minerals-outlook-2025>

Ingemarsson, L. (2026) *EU AI Infrastructure and Compute Capacity Report 2026*, Alice Labs, available at <https://alicelabs.ai/reports/eu-ai-infrastructure-compute-capacity-2026>

Juhász, R., N. Lane and D. Rodrik (2024) 'The New Economics of Industrial Policy', *Annual Review of Economics* 16: 213–242, available at <https://doi.org/10.1146/annureveconomics-081023-024638>

Lee, K.-F. (2018) *AI Superpowers: China, Silicon Valley, and the New World Order*, Houghton Mifflin

- Mariniello, M. (2026) 'The right balance: how to fix European Union artificial intelligence regulation', Policy Brief 12/2026, Bruegel, available at <https://doi.org/10.64153/BVAQ6479>
- Mazzucato, M. (2013) *The Entrepreneurial State*, Anthem Press, available at <https://marianamazzucato.com/books/the-entrepreneurial-state/>
- Miller, C. (2022) *Chip War*, Scribner, available at <https://www.simonandschuster.com/books/Chip-War/Chris-Miller/9781982172015>
- Nilekani, N. and V. Shah (2015) *Rebooting India*, Penguin Random House India, available at <https://www.penguin.co.in/book/rebooting-india-2/>
- O'Hara, K. and W. Hall (2018) 'Four Internets', CIGI Paper 206, Centre for International Governance Innovation, available at <https://www.cigionline.org/documents/1563/Paper%20no.206web.pdf>
- Powers-Riggs, A., B. Hart, M.P. Funaiole, S. Thomsett, S. Zieleniewski and T. Rose (2025) 'Beyond Rare Earths: China's Growing Threat to Gallium Supply Chains', CSIS Briefs, Center for Strategic and International Studies, available at <https://www.csis.org/analysis/beyond-rare-earths-chinasgrowing-threat-gallium-supply-chains>
- Prajapati, D., J. Mali and F. Mehta (2020) *Data Empowerment And Protection Architecture*, NITI Aayog, available at <https://www.niti.gov.in/node/709>
- Rodrik, D. (2022) 'An Industrial Policy for Good Jobs', Commentary, 28 September, Brookings, available at <https://www.brookings.edu/articles/an-industrial-policy-for-good-jobs/>
- Sastry, G., L. Heim, H. Belfield, M. Anderljung, M. Brundage, J. Hazell ... D. Coyle (2024) 'Computing Power and the Governance of Artificial Intelligence', mimeo, available at <https://doi.org/10.48550/arXiv.2402.08797>
- Scharre, P. (2023) *Four Battlegrounds: Power in the Age of Artificial Intelligence*, W.W. Norton, available at <https://wwnorton.com/books/9780393866865>
- Poitiers, N. and T. Schenk (2026) 'Revamping Europe's chips strategy: indispensability, not selfsufficiency', Analysis 12/2026, Bruegel, available at <https://doi.org/10.64153/DADR3364>
- Sheehan, M. (2023) 'China's AI Regulations and How They Get Made', Working Paper, Carnegie Endowment, available at https://assets.carnegieendowment.org/static/files/202307-Sheehan_Chinese%20AI%20gov-1.pdf
- Srivastava, S. and J. Bullock (2024) 'AI, Global Governance, and Digital Sovereignty', mimeo, available at <https://doi.org/10.48550/arXiv.2410.17481>
- Stanford HAI (2024) *The 2024 AI Index Report*, Stanford Institute for Human-Centered AI, available at <https://hai.stanford.edu/ai-index/2024-ai-index-report>
- Stanford HAI (2025) *The 2025 AI Index Report*, Stanford Institute for Human-Centered AI, available at <https://hai.stanford.edu/ai-index/2025-ai-index-report>
- Stanford HAI (2026) *The 2026 AI Index Report*, Stanford Institute for Human-Centered AI, available at <https://hai.stanford.edu/ai-index/2026-ai-index-report>
- Veale, M. and F. Zuiderveen Borgesius (2021) 'Demystifying the Draft EU Artificial Intelligence Act', *Computer Law Review International* 22(4): 97–112, available at <https://ssrn.com/abstract=3896852>

Appendix: Coalition feasibility, supporting data

This appendix reports supporting quantitative data referenced in section 7 on the financial feasibility of coalition-scale coordinated specialisation.

Table A1: Principal coalition sovereign wealth funds (2025)

Fund	Country	Assets 2025 (\$bns)	Source
Government Pension Fund Global	Norway	~2,200	Multiple confirmed (end-2025)
Abu Dhabi Investment Authority	UAE	~1,110	SWF Institute 2025
Public Investment Fund	Saudi Arabia	~1,000–1,150	SWF Institute 2025
GIC Private Limited	Singapore	~936	GIC Annual Report, Aug. 2025
Temasek	Singapore	~320	Temasek Annual Report 2025
Four-fund total		~5,246–5,396	
2% mandate, annual		~105–108	Comparable to current EU+China private AI investment combined

Table A2: Cumulative AI investment by selected coalition members to 2024

Country	Cumulative AI investment, 2013–2024 (\$bns)	Estimated 2025 annual (\$bns)	Source
India	11.1	~3.5	Stanford HAI (2025)
Japan	5.9	~1.5	Stanford HAI (2025)
South Korea	8.96	~1.0	Stanford HAI (2025)
Australia	3.88	~0.8	Stanford HAI (2025)

Natixis CIB Research

Head of CIB Research



Jean-François Robin

+33 1 58 55 13 09

jean-francois.robin@natixis.com

Chief Economist, Asia Pacific



Alicia Garcia Herrero

+852 3900 8680

alicia.garciaherrero@natixis.com

Emerging Asia



Trinh Nguyen

+852 3900 8726

trinh.nguyen@natixis.com

Greater China



Jianwei Xu

+852 3900 8034

jianwei.xu@natixis.com

Japan, Pacific



Kohei Iwahara

+813 4519 2144

kohei.iwahara@natixis.com

Asia Pacific, Thematic Research



Gary Ng

+852 3915 1242

gary.ng@natixis.com



Haoxin Mu

+852 3900 8067

haoxin.mu@natixis.com

In compliance with Market Abuse Regulation (MAR) that came into force on 3 July 2016, please find relevant information regarding investment recommendations. As a staff member of Natixis, the producer of the document is considered as the author of an investment recommendation as defined by MAR. Policy for managing conflicts of interest: https://natixis.groupebpce.com/wp-content/uploads/2022/11/summary_policy_management_of_conflicts_of_interest_s_eng_220826.pdf

Information about potential conflicts of interest within Natixis/Groupe BPCE: <https://home.cib.natixis.com/regulation>

This document is a communication for discussion and information purposes only and is exclusively intended for and may only be transmitted to eligible counterparties, professional clients and qualified investors, as such terms (or their equivalents) are understood and applied under all relevant laws and regulations applicable to you and us, for purposes of this communication. This document is not intended for retail client use. It is not for further distribution or distribution to retail clients. It should not be reproduced (in whole or in part), forwarded, transmitted, or summarized to any person other than the original addressee(s) without the prior written consent of Natixis. This communication is highly confidential and is the property of Natixis. If you are not the intended recipient of this document and/or any related attachments, or if you do not agree to or cannot comply fully with the requirements and terms of the disclaimers which follow, please delete it (and/or them) and immediately notify the sender.

This document does not constitute a personalized recommendation. It is intended for general information and the products or services described herein do not take into account any specific investment objective, financial situation or particular need of any recipient. Natixis does not provide any advice herein, including in particular in case of investment services. Under no circumstances should this document be considered as an official confirmation of a transaction to any person or entity.

The distribution, possession or delivery of this document in, to or from certain jurisdictions may be restricted or prohibited by law. Recipients of this document are therefore required to ensure that they are aware of, and comply with, such restrictions or prohibitions. Neither Natixis, nor any of its affiliates, directors, employees, agents or advisers nor any other person may be deemed liable to anyone in relation to the distribution, possession or delivery of this document in, to or from any jurisdiction.

Further, it is important to note that not all of the products, services and information are available in all jurisdictions. The information does not constitute an offer or solicitation to buy or sell any security, currency, derivative, investment fund or other product, service or information to anyone in any jurisdiction in which an offer or solicitation is not authorized or cannot legally be made or to any person to whom it is unlawful to make an offer or solicitation. For information concerning any specific jurisdiction, please contact your local sales representative. In addition, this material is not intended to provide specific financial, investment, tax, legal or accounting advice for you in any jurisdiction, and should not be relied upon in that regard. You should not act or rely on the information without seeking the advice of a professional in the relevant jurisdiction.

Securities or financial instruments mentioned herein may not be suitable or appropriate for all investors or in all geographical areas. Every state in the US, and most countries throughout the world, have their own laws regulating the types of securities and other investment products which may be offered to their residents, as well as the process for doing so. As a result, the products and services discussed in this material may not be eligible for sale or otherwise in some jurisdictions. This material is not, and under no circumstances should be construed as, a solicitation to act as securities broker or dealer or otherwise in any jurisdiction by any person or company that is not legally permitted to carry on the business of a securities broker or dealer or otherwise in that jurisdiction.

Any guarantee, funding, interest or currency swap, underwriting or more generally any undertaking (if mentioned in this document) should be treated as preliminary only and is subject to a formal approval and written confirmation in accordance with Natixis' current internal procedures. Natixis has neither verified nor conducted an independent analysis of the information contained in this document. Therefore, Natixis makes no statement or warranty and makes no commitment to the recipients of this document in any way (express or implied) as to its relevance, the accuracy or completeness of the information contained therein or the appropriateness of the assumptions to which it refers. Indeed, the information contained herein does not take into account the specific accounting or tax rules that would apply to Natixis' counterparties, customers or potential customers. Natixis cannot therefore be held liable for any differences in valuation between its own data and that of third parties, these differences being due in particular to considerations on the application of accounting rules, tax or valuation models. The statements, assumptions and opinions contained in this document may be changed or may be withdrawn by Natixis at any time without notice.

Natixis shall not be liable for any financial loss, damages, or any decision taken on the basis of the information contained in this document and Natixis does not hold itself out as providing any advice, particularly in relation to investment services.

In any event, you should independently retain any internal and/or external advice that you consider necessary or desirable to obtain, including any financial, legal, tax or accounting advice, or any other specialist advice, in order to verify in particular that the investment(s), transaction(s), structure(s) or arrangement(s) described in this document meets your investment or financial objectives and constraints, and to obtain an independent valuation of such investment(s)/transaction(s), its risks factors and rewards. It should not be assumed that the information contained in this document will have been updated subsequent to the date stated on the front page of this document. In addition, the delivery of this document does not imply in any way an obligation on anyone to update the information contained herein at any time.

Prices and margins are deemed to be indicative only and are subject to change at any time depending on, inter alia, market conditions. **Past performances and simulations of past performances are not a reliable indicator and therefore do not anticipate future results.** The information contained in this document may include results of analyses from a quantitative model, which represent potential future events, that may or may not be realized, and is not a complete analysis of every material fact representing any product. Information may be

changed or may be withdrawn by Natixis at any time without notice. Natixis makes no representation or warranty, express or implied, as to the accomplishment of or reasonableness of, nor should any reliance be placed on any projections, targets, estimates or forecasts, or on the statements, assumptions and opinions expressed in this document. Nothing in this document should be relied on as a promise or guarantee as to the future.

Natixis may have conflicts of interest: Natixis may from time to time, as principal or agent, be involved in a wide range of activities globally, have positions in, or may buy or sell, or make a market in any securities, currencies, financial instruments or other assets underlying the instruments to which the Information relates, including possibly a position that was accumulated on the basis of this analysis material prior to its dissemination. Trading desks may also have or take positions inconsistent with this material. Natixis activities related to those instruments may have an impact on the price of the relevant underlying asset and may give rise to conflicting interests or duties.

Natixis may provide services to any member of the same group as the recipient of the information or any other entity or person (a Third Party), engage in any transaction (on its own account or otherwise) with respect to the recipient of the information or a Third Party, or act in relation to any matter for itself or any Third Party, notwithstanding that such services, transactions or actions may be adverse to the recipient of the information or any member of its group, and Natixis may retain for its own benefit any related remuneration or profit. In addition, Natixis may, whether by virtue of the types of relationships described in this paragraph or otherwise, from time to time be in possession of information in relation to a particular instrument that is or may be material in the context of that instrument and that may or may not be publicly available or known to you. Our providing you with an indicative price quotation or other information with respect to any such instrument does not create any obligation on the part of Natixis to disclose to you any such information (whether or not confidential).

Natixis reserves the right to record, exchange, analyse and use relevant information about you and your relationships with any affiliates of Natixis (including the nature of your transactions) for combating money laundering and terrorists financing as well as for credit assessment, market research, insurance and administrative purposes including back-office management. This may include information provided by you, or someone acting on your behalf. Relevant information may also be exchanged with members of the Natixis group and others, for audit or compliance purposes and if required by appropriate governmental and non-governmental regulators as well as in accordance with applicable laws and regulations.

No part of author compensation was, is or will be, directly or indirectly related to the specific recommendations or views expressed in this report. The personal views of authors may differ from one another. Natixis, its subsidiaries and affiliates may have issued or may issue reports that are inconsistent with, and/or reach different conclusions from, the information presented herein.

The stocks/companies mentioned might be subject to specific disclaimers. Please click on the following link to consult them: <https://www.research.natixis.com/Site/en/disclaimers/3>.

France: Natixis is supervised by the European Central bank (ECB). Natixis is authorized in France by the Autorité de Contrôle Prudentiel et de Régulation (ACPR) as a Bank-Investment Services Provider and subject to its supervision. Natixis is regulated by the Autorité des Marchés Financiers in respect of its investment services activities.

This material is not, and is not intended to be, a "research report", "investment research" or "independent research" as may be defined in applicable laws and regulations worldwide. However, it may constitute "research" as defined in Recital 28 of the Commission Delegated Directive (EU) 2017/593 and in the Financial Conduct Authority's Handbook.

EU: In compliance with Market Abuse Regulation (MAR), please find relevant information regarding investment recommendations. If this material is an investment recommendation as that term is defined in Article 3(1)(35) of the EU and UK MAR, distribution of this material is subject to the relevant provisions of MAR. Policy for managing conflicts of interest: [Summary-Policy-management-of-conflicts-of-interests ENG 20241212.pdf](https://natixis.groupebpce.com/wp-content/uploads/2022/11/summary_policy_management_of_conflicts_of_interest_s_eng_20241212.pdf) Information about potential conflicts of interest within Natixis/Groupe BPCE: <https://home.cib.natixis.com/regulation>.

USA: Natixis, a foreign bank and broker-dealer, makes this report available solely for distribution in the United States to major U.S. institutional investors as defined in Rule 15a-6 under the U.S. securities Exchange Act of 1934. This document shall not be distributed to any other persons in the United States. All major U.S. institutional investors receiving this document shall not distribute the original nor a copy thereof to any other person in the United States.

Natixis is a National Futures Association (NFA) member in the United States and has filed its Swap Dealer registration per the Commodity Futures Trading Commission (CFTC) regulations, and is registered as a Security-Based Swap Dealer with the Securities and Exchange Commission (SEC). Please refer to the Natixis website and applicable final legal documentation relating to the relevant product for additional risk disclosures. Investors are exposed to the default risk of the issuer and/or the guarantor as applicable. For any swap products, please also refer to Dodd-Frank Act - Compliance - Natixis https://www.natixis.com/natixis/en/dodd-frank-act-rpz5_115246.html.

Structured securities, options, futures and other derivatives are complex instruments and may involve a high degree of risk. Before entering into any such transactions, please ensure that you have read and understood the Options Clearing Corporation's Characteristics and Risks of Standardized Options, also known as the options disclosure document (ODD) available at: <http://www.theocc.com/about/publications/character-risks.jsp>.

THIS COMMUNICATION IS NOT A "RESEARCH REPORT" AS DEFINED IN CFTC REGULATIONS 1.71 and 23.605 AND IS TO BE DISTRIBUTED ONLY TO INSTITUTIONAL CLIENTS. THIS COMMUNICATION HAS BEEN GENERATED BY ONE OR MORE EMPLOYEES OF NATIXIS AND IS BEING CONVEYED TO YOU EITHER (I) COMMENTARY ON ECONOMIC, POLITICAL, OR MARKET CONDITIONS AND STATISTICAL SUMMARIES OF MULTIPLE COMPANIES' FINANCIAL DATA, WHICH MAY INCLUDE LISTS OF CURRENT RATINGS, [research.natixis.com](https://www.research.natixis.com)

AND ON THE BASIS OF WHICH NO INVESTMENT DECISION SHOULD BE MADE AS THIS MATERIAL IS CONDENSED, NOT COMPREHENSIVE, AND DOES NOT INCLUDE ALL RISK FACTORS OR OTHER MATTERS THAT MAY BE MATERIAL, OR (II) A SOLICITATION FOR ENTERING INTO A DERIVATIVES TRANSACTION, BUT IS NOT A BINDING OFFER TO BUY OR SELL ANY FINANCIAL INSTRUMENT. THE VIEWS EXPRESSED HEREIN ARE THOSE OF THE AUTHOR AS OF THE DATE OF THIS COMMUNICATION AND ARE NOT TO BE RELIED ON AS AUTHORITATIVE, AND DO NOT NECESSARILY REFLECT THE VIEWS OF NATIXIS OR ANY OF ITS AFFILIATES OR SUBSIDIARIES, OR THEIR RESPECTIVE OFFICERS, DIRECTORS, EMPLOYEES, CONTRACTORS, AGENTS, REPRESENTATIVES, SUCCESSORS AND ASSIGNS. THIS IS NOT A RECOMMENDATION OR INVESTMENT ADVICE. ANY VIEWS EXPRESSED HEREIN MAY NOT BE OBJECTIVE OR INDEPENDENT OF THE AUTHOR OR OTHER NATIXIS (AND ITS AFFILIATES) SALES AND TRADING DESKS, WHO ARE ACTIVE PARTICIPANTS IN THE MARKETS, INVESTMENTS, OR STRATEGIES REFERRED TO IN THIS COMMUNICATION. THEY ARE SUBJECT TO CHANGE WITHOUT NOTICE. THIS COMMUNICATION IS BASED ON PUBLIC INFORMATION AND SHALL NOT BE USED NOR CONSIDERED AS AN UNDERTAKING FROM NATIXIS. NATIXIS TAKES NO RESPONSIBILITY FOR ANY ERRORS OR OMISSIONS AND DOES NOT GUARANTEE THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THE COMMUNICATION. NEITHER THE AUTHORS NOR NATIXIS HAVE ANY OBLIGATION TO UPDATE, MODIFY OR OTHERWISE NOTIFY THE RECIPIENT IN THE EVENT ANY INFORMATION CONTAINED HEREIN, INCLUDING ANY OPINION OR VIEW, CHANGES OR BECOMES INACCURATE. NATIXIS SHALL NOT BE LIABLE FOR ANY FINANCIAL LOSS, DAMAGES, OR ANY DECISION TAKEN ON THE BASIS OF THE INFORMATION DISCLOSED IN THIS COMMUNICATION AND NATIXIS DOES NOT PROVIDE ANY ADVICE, INCLUDING IN CASE OF INVESTMENT SERVICES. THIS COMMUNICATION IS COMMUNICATED TO EACH RECIPIENT FOR INFORMATIONAL PURPOSES ONLY.

This document does not constitute a financial analysis or research report and has not been developed in accordance with legal requirements designed to promote the independence of investment research. Accordingly, there are no prohibitions on dealing ahead of its dissemination. Opinions expressed may differ from the opinions expressed by other business units of Natixis.

Natixis and each of its respective affiliates accepts no liability whatsoever for any direct, indirect or consequential losses (in contract, tort or otherwise) arising from the use of this material or reliance on the information contained herein. However, this shall not restrict, exclude or limit any duty or liability to any person under any applicable laws or regulations of any jurisdiction which may not lawfully be disclaimed.

No representation or assurance is provided by Natixis with respect to whether any product or financial instrument referenced in this document meets requirements or criteria established under any regulatory, industry or other rules and standards related to climate, sustainability or environmental matters.

Canada: In Canada, this document may only be transmitted to "permitted clients" as defined in National Instrument 31-103. No further distribution is permitted. Distribution to retail investors is prohibited. In Canada, Natixis operates through Natixis Canada Branch, a Foreign Bank Branch (Schedule 3 – Lending only). Natixis Securities Americas LLC is not a registered broker-dealer in any Canadian province but operates under the International Dealer Exemption in each province. Natixis S.A. is not a registered broker-dealer in any Canadian province and operates under the International Dealer Exemption in Quebec and Ontario. Any securities transaction or other offer, sale, or solicitation will be structured to be in compliance with such exemption.

Mexico: As per your request, this information is delivered to you, which strictly pertains to activities performed by Natixis, a Foreign Financial Entity in its country of origin, such institution is herein represented by its Representation Office in Mexico. This information is being delivered to you for information purposes only. Please note that the Representation Office in Mexico is prevented from negotiating, executing, or participating in (i) operations related to listed derivatives; and (ii) securing funds from the general public in Mexico.

Brazil: This communication and the matters discussed therein have not been and will not be issued, placed, distributed or offered in the Brazilian market. Documents relating to this communication, as well as information contained therein may not be supplied to the public in Brazil nor be used in connection with any offer for subscription or sale of financial products or services to the public in Brazil. This communication will not be offered in Brazil, except in circumstances which do not constitute any breach or noncompliance with the Brazilian applicable legislation. Persons wishing to acquire or invest in any debt or equity interest in Brazil should consult with their own counsel as to the applicability of registration requirements or any exemption therefrom. This material is only intended for persons who are eligible investors residing in Brazil, which are considered to be Professional Investors (Investidores Profissionais), as designated by the applicable regulation, mainly the CVM Resolution No. 30 from the 11th of May 2021.

Argentina: Products/services included in the presentation have not been and will not be publicly issued, placed, distributed, offered, or registered in Argentine Securities Commission (Comisión Nacional de Valores (CNV). Neither this nor any other material may be utilized in connection with any general offering to the public within Argentina. Any Argentine resident who acquires any products or services will do so under his/her own responsibility under the terms of a private offering to him/her, from outside of Argentina.

Panama: This document and its content are for information purposes and shall not be interpreted as banking or financial intermediation, business solicitation and/or public offering of any kind. Natixis is duly authorized by the relevant authorities to carry out banking activities in the United States. Any and all services provided by Natixis are provided on a cross border basis outside of Panama. Natixis is not authorized by the Superintendencia de Banks nor the Superintendencia of the Securities Market to carry out any activities in Panama. The information contained herein is not intended to be a source of advice or credit analysis with respect to the material presented, and the information and/or documents contained in this material do not constitute investment advice or an invitation to offer.

Peru: The products or information discussed here have not been registered under the Peruvian Securities Market Law (Supreme Decree N° 093-2002-EF) or before the Superintendencia del Mercado de Valores (SMV). There will be no public offering of the shares in Peru. The SMV has not reviewed the information provided to the investor. This material is only for the exclusive use of institutional investors in Peru and is not for public distribution.

Colombia: By receiving this document you confirm that you have requested us to provide you this document and information on the matters described in it. You also confirm that your interest in requesting such documents and information arose out of your own private interest and was not the result of any direct or indirect solicitation, promotion or offering of services or products by Natixis or by any of its representatives. Accordingly, the information contained in this document is not intended and should not be construed as constituting onshore promotion, marketing or solicitation of securities-market services or products. The distribution of the information contained in this document is restricted by law and persons who access it are required to comply with all such restrictions. The information is not intended to be published or made available to any person in any jurisdiction where doing so would contravene any applicable laws or regulations. By receiving this document, you confirm that you are aware of the laws in Colombia relating to the promotion and marketing of financial services products and you warrant and represent that you will not pass on or utilise the information contained in this document in a manner that could constitute a breach of such laws by Natixis, its affiliates or any other person. Any securities mentioned in this document have not been and will not be registered with the National Register of Securities ("Registro Nacional de Valores y Emisores") maintained by the Colombian Financial Superintendency ("Superintendencia Financiera de Colombia") and may not be publicly offered or sold in Colombia. This information does not constitute and should not be construed as an offer to enter into any agreement, or to purchase or subscribe any securities. The information on the following pages is provided for information purposes only and does not constitute investment, legal, tax or other advice or any recommendation to buy, sell or otherwise transact in any of the funds or securities mentioned. Prospective investors should take appropriate professional advice before making any investment decision.

Bolivia: Natixis is not authorized to conduct in Bolivia activities reserved by law for financial intermediation entities and/or brokerage firms. The products/services offered will be provided from Natixis's offices outside of Bolivia.

Chile: This communication and any accompanying information (the "Materials") are intended solely for informational purposes and do not constitute (and should not be interpreted to constitute) the selling, or conducting of business with respect to such products or services in Chile (this "Jurisdiction"), or the conducting of any brokerage, banking or other similarly regulated activities in this Jurisdiction. Natixis ("Bank") is not registered (or intended to be registered) in this Jurisdiction. The Materials are private, confidential and are sent by the Bank only for the exclusive use of the addressee. The Materials must not be publicly distributed and any use of the Materials by anyone other than the addressee is not authorized. The addressee is required to comply with all applicable laws in this Jurisdiction, including, without limitation, tax laws and exchange control regulations, if any. The services do not include the offer or intermediation of any securities. Additionally, the Bank will not be liable for investment decisions made by the client based on the information derived or obtained from the Bank's materials.

UK: In the UK, Natixis SA is authorised by L'Autorité de contrôle prudentiel et de résolution (ACPR), supervised by the European Central Bank (ECB) and regulated by the Autorité des Marchés Financiers (AMF). Authorised by the Prudential Regulation Authority. Subject to regulation by the Financial Conduct Authority and limited regulation by the Prudential Regulation Authority. Details about the extent of our regulation by the Prudential Regulation Authority are available from us on request.

Germany: Natixis Zweigniederlassung Deutschland is subject to a limited form of regulation by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) with regards to the conduct of its business in Germany under the right of establishment there. The transfer / distribution of this document in Germany is performed by / under the responsibility of Natixis Zweigniederlassung Deutschland.

Spain: Natixis is regulated by Bank of Spain and the CNMV (Comisión Nacional del Mercado de Valores) for the conduct of its business under the right of establishment in Spain.

Italy: Natixis is regulated by Bank of Italy and the CONSOB (Commissione Nazionale per le Società e la Borsa) for the conduct of its business under the right of establishment in Italy.

Japan: Natixis Tokyo Branch is authorized in Japan by the Financial Service Agency as a Bank. Natixis Japan Securities Co., Ltd. (NJS) is a Financial Instruments Business Operator (Director General of Kanto Local Finance Bureau (Kinsho) No. 2527 and a Money Lender (Tokyo Governor (4) No. 31548), supervised by the Financial Services Agency. NJS is a member of the Japan Securities Dealers Association and Type II Financial Instruments Firms Association. This document is intended solely for distribution to Professional Investors as defined in Article 2.31 of the Financial Instruments and Exchange Act.

Hong Kong: In Hong Kong, Natixis is regulated by the Hong Kong Monetary Authority and registered with the Securities and Futures Commission (the "SFC") to carry on Type 1 (dealing in securities), Type 2 (dealing in futures contracts), Type 4 (advising on securities), Type 5 (advising on futures contracts) and Type 6 (advising on Corporate Finance) regulated activities under the Securities and Futures Ordinance (Cap. 571 of the Laws of Hong Kong) ("SFO"). This document has not been approved by the SFC nor has a copy of this document been registered by the Registrar of Companies in Hong Kong, unless specified otherwise. The investments contained herein may or may not be authorised by the SFC. You are advised to exercise caution in relation to the investments contained herein. If you are in any doubt about any of the contents of this document, you should obtain independent professional advice. The investments may not be offered or sold in Hong Kong, by means of any document, other than (i) to "professional investors" as defined in the SFO and any rules made under the SFO, or (ii) in other circumstances which do not result in the document being a "prospectus" as defined in the Companies (Winding Up and Miscellaneous Provisions) Ordinance (Cap. 32 of the Laws of Hong Kong) (the "CWUMPO") or which do not constitute an offer to the public within the meaning of the CWUMPO. No person shall issue or possess for the purposes of issue, whether in Hong Kong or elsewhere, any advertisement, invitation or document relating to the investments, which to his knowledge is or contains an invitation to the public of Hong Kong (except if permitted to do so under the applicable laws and rules of Hong Kong) other than with respect to investments which are or are intended to

be disposed of only to persons outside Hong Kong or only to "professional investors" as defined in the SFO and any rules made under the SFO. This document is not and should not be construed as Natixis (or any affiliate or anyone else) offering or holding itself out as being willing or able to provide any service or to carry on any activity which it is not licensed to provide or carry on in Hong Kong. Natixis Asia Limited is registered with the Securities and Futures Commission to carry on Type 1 (dealing in securities), Type 2 (dealing in futures contracts), Type 4 (advising on securities) and Type 5 (advising on futures contracts) regulated activities under the Securities and Futures Ordinance.

Singapore: Natixis Singapore branch is regulated by the Monetary Authority of Singapore. If this document is distributed in Singapore, this document is intended solely for distribution to Institutional Investors, Accredited Investors and Expert Investors as defined in Section 4A of the Securities and Futures Act 2001 of Singapore.

People's Republic of China: In Mainland China, Natixis Shanghai branch and Natixis Beijing branch are regulated in the People's Republic of China (the "PRC", which, for such purposes, does not include the Hong Kong or Macau Special Administrative Regions or Taiwan) by National Financial Regulatory Administration, the People's Bank of China, and the State Administration of Foreign Exchange for the engagement of banking business. This document is intended solely for distribution to investors of the PRC who are qualified/authorized under the laws of the People's Republic of China to engage in the purchase of Product of the type being offered or sold. PRC investors are responsible for verifying their eligibility to purchase the Product, obtaining all relevant approvals/licenses, verification and/or registrations themselves from relevant governmental authorities, and complying with all relevant applicable regulatory requirements under the PRC legal regime, including, but not limited to, all relevant foreign exchange regulations and/or outbound investment regulations.

Taiwan: In Taiwan, Natixis Taipei Branch is authorized by the Financial Supervisory Commission as a commercial bank and is subject to its supervision. The Product may not be sold or offered to Taiwan resident investors or in Taiwan unless they are made available, (i) outside Taiwan for purchase by such investors outside Taiwan and/or (ii) investors are authorized to engage in the purchase of product in accordance with applicable laws.

South Korea: In Korea, Natixis Asia Limited Seoul Branch is registered as a licensed financial investment business entity under the Financial Investment Services and Capital Markets Act ("FSCMA") and regulated by the Financial Supervisory Services and Financial Services Commission. This document/communication is distributed by Natixis Asia Limited Seoul Branch in the republic of Korea to Professional Investor only and is not intended for investors in Korea who are not Professional Investor within the meaning of the Article 9(5) of the FSCMA and should not be passed on to any such persons. Some products or transactions described in this document/communication may not be authorized in Korea and may not be available to Korean investors. This document/communication is being provided in accordance with procedures required under applicable laws and regulations as well as internal control standards of Natixis Asia Limited, Seoul Branch.

India, Indonesia, and Thailand: Natixis has set-up representative offices in India, Indonesia and Thailand. The activities of the representative offices are supervised by the Reserve Bank of India, the Otoritas Jasa Keuangan (Financial Services Authority) and the Bank of Thailand respectively.

Australia: Natixis Asia Limited is registered with the Securities and Futures Commission to carry on Type 1 (dealing in securities), Type 2 (dealing in futures contracts), Type 4 (advising on securities) and Type 5 (advising on futures contracts) regulated activities under the SFO. If this document is distributed in Hong Kong, this document is only to Professional Investors as defined in the SFO of Hong Kong and any rules made under the SFO.

In Australia, Natixis has a wholly owned subsidiary, Natixis Australia Pty Limited ("NAPL"). NAPL is registered with the Australian Securities and Investments Commission and holds an Australian Financial Services License (No. 317114) which enables NAPL to conduct its financial services business in Australia with "wholesale" clients. Details of the AFSL are available upon request Natixis S.A. is not an Authorised Deposit-Taking Institution under the Australian Banking Act 1959, nor is it regulated by the Australian Prudential Regulation Authority. Any references made to banking in the document refer to Natixis activities outside of Australia.

United Arab Emirates: Natixis is authorized by the ACPR and regulated by the Dubai Financial Services Authority (DFSA) for the conduct of its business in and from the Dubai International Financial Centre (DIFC). The document is being made available to the recipient with the understanding that it meets the DFSA definition of a Market Counterparty or Professional Client; the recipient is otherwise required to inform Natixis if this is not the case and return the document. In accordance with the Decision No (13/R.M) of 2021 on the Financial Activities Rulebook and Mechanisms of Adjustment of the Securities and Commodities Authority (SCA) (the SCA Rulebook), the products to which this document relates may only be promoted in the UAE (excluding the Dubai International Financial Centre (DIFC) and the Abu Dhabi Global Market (ADGM)) as follows: (i) to a person who meets the definition of a "professional investor" under the SCA Rulebook; or (ii) following a 'reverse' (i.e., unsolicited) enquiry by an investor. Further, this document does not constitute a public offer of securities in the United Arab Emirates (including the DIFC and the ADGM) and is not intended to be a public offer. The SCA has not verified this document or other documents in connection with the products. Accordingly, the SCA has not approved this document or any other associated documents nor taken any steps to verify the information set out in this document. The SCA may not be held liable for the accuracy or completeness of the information in this document. The products to which this document relates may be illiquid or subject to restrictions on their resale. Prospective investors should conduct their own due diligence on the products. If you do not understand the contents of this document you should consult an authorised financial advisor.

The recipient also acknowledges and understands that neither the document nor its contents have been approved, licensed by or registered with any regulatory body or governmental agency in the GCC or Lebanon.

Oman: In Oman, Natixis neither has a registered business presence nor a representative office and does not undertake banking business or provide financial services in Oman. Consequently, Natixis is not regulated by either the Central Bank of Oman or Oman's Financial Services Authority. This document has been prepared by Natixis. The information contained in this document is for discussion purposes only and neither constitutes an offer of securities in Oman as contemplated by the Commercial Companies Law of Oman (Royal Decree 18/2019) or the Securities Law (Royal Decree 46/2022), nor does it constitute an offer to sell, or the solicitation of any offer to buy non-Omani securities in Oman as contemplated by Article 139 of the Executive Regulations to the Capital Market Law (issued vide CMA Decision 1/2009). Additionally, this document is not intended to lead to the conclusion of a contract for the sale or purchase of

securities. Natixis does not solicit business in Oman and the only circumstances in which Natixis sends information or material describing financial products or financial services to recipients in Oman, is where such information or material has been requested from Natixis and by receiving this document, the person or entity to whom it has been dispatched by Natixis understands, acknowledges and agrees that this document has not been approved by the CBO, the FSA or any other regulatory body or authority in Oman. Natixis does not market, offer, sell or distribute any financial or investment products or services in Oman and no subscription to any securities, products or financial services may or will be consummated within Oman. This document has not been approved by the FSA or any other regulatory body or authority in Oman, and no authorization, licence or approval has been received by Natixis from the FSA or any other regulatory authority in Oman, to market, offer, sell, or distribute the securities within Oman. Natixis does not advise persons or entities resident or based in Oman as to the appropriateness of investing in or purchasing or selling securities or other financial products. Nothing contained in this document is intended to constitute Omani investment, legal, tax, accounting or other professional advice. The recipient of this communication represents that it is a financial institution and/or is a sophisticated investor (as described in Article 139 of the Executive Regulations of the Capital Market Law) and that its officers/employees have such experience in business and financial matters that they are capable of evaluating the merits and risks of investments.

Qatar: The investments described in this document have not been, and will not be, offered, sold or delivered at any time, directly or indirectly, in the State of Qatar in a manner that would constitute a public offering. This document has not been, and will not be, filed with, reviewed by or approved by the Qatar Central Bank, the Qatar Financial Markets Authority or any other relevant Qatari authority. This document is intended for the original recipient only and should not be provided to any other person. It is not for general circulation in the State of Qatar and should not be reproduced or used for any other purpose.

Saudi Arabia: Natixis Saudi Arabia Investment Company, licensed and regulated by the Capital Market Authority (Number: 19205-31), located in Alfaisaliah Tower, 16th Floor, King Fahad Road, Riyadh – KSA.

Turkey: Natixis is not licensed in Turkey to perform any banking/capital market activity. No information in this document is provided for the purpose of offering, marketing and sale by any means of any capital market instruments in the Republic of Turkey. Therefore, this document may not be considered as an offer made or to be made to residents of the Republic of Turkey. Accordingly, neither this document nor any other material may be utilized in connection with any offering to the public within the Republic of Turkey without the prior approval of the Turkish Capital Markets Board. However, according to Article 15 (d) (ii) of the Decree No. 32 there is no restriction on the purchase or sale of products in this document by residents of the Republic of Turkey, provided that: they purchase or sell such products in the financial markets outside of the Republic of Turkey; and such sale and purchase is made through banks, and/or licensed brokerage firms in the Republic of Turkey.

Links from or to web sites outside Natixis sites are meant for convenience only. Natixis does not review, endorse, approve or control, and is not responsible for any sites linked from or to Natixis sites, the content of those sites, the third parties named therein, or their products and services. Linking to any other site is at your sole risk and Natixis will not be responsible or liable for any damages in connection with linking. No endorsement or approval of any third parties or their advice, opinions, information, products or services is expressed or implied by these materials.