

Policy Brief

Series Information:
EPINOVA Policy Brief
Series on Strategic
Competition,
Geopolitics, and
Emerging Security

Recommended
Citation:
Wu, S. (2026). *South
Korea's AI Power under
Uneven Interdependence:
Semiconductor Depth,
National Conversion,
and the KIC Model-
Selection Case* (Policy
Brief No. EPINOVA-
2026-PB-74). Global AI
Governance and Policy
Research Center,
EPINOVA LLC.
[https://doi.org/10.67037/
epinova.pb.2026.074](https://doi.org/10.67037/epinova.pb.2026.074)


Disclaimer:
The analysis in this
publication is based on
publicly available
information, official
statements, media
reporting, and the
author's independent
analysis. It does not
represent the official
position of any
government. This
publication is intended
solely for research and
policy discussion and
does not constitute legal,
military, operational, or
sanctions compliance
advice.

South Korea's AI Power Under Uneven Interdependence:
Semiconductor Depth, National Conversion, and the KIC Model-Selection Case

Author: Shaoyuan Wu
Affiliation: Global AI Governance and Policy Research Center, EPINOVA LLC
Date: October 19, 2026

Key Judgments

- **South Korea possesses substantial specialized AI capacity, but national conversion remains uneven.** Korean firms have reached commercial HBM4 supply milestones, domestic model developers have advanced through successive government evaluations, and selected industrial and administrative AI applications are operating or being piloted. However, the durability, national coverage, and disruption resilience of these capabilities remain insufficiently documented.
- **Strategic supply-chain importance does not automatically translate into AI Boundary Power.** Korean memory suppliers occupy consequential positions in globally interdependent semiconductor production networks. Their relational leverage depends on effective control rights, critical customer interfaces, functional substitutability, and the capacity to alter access conditions rather than supplier importance alone.
- **The KIC case demonstrates nonexclusive AI model adoption without establishing foreign control or operational dependence.** Reporting based on materials submitted by the Korea Investment Corporation (KIC) identifies Chinese-developed Qwen3 and BGE-M3 in core language and retrieval functions alongside two U.S.-developed models. However, the available evidence does not establish hosting arrangements, continuing control rights, switching costs, or measured organizational benefits.
- **Boundary Autonomy depends on credible functional substitution rather than supplier diversity alone.** The coexistence of domestic and foreign AI technologies does not necessarily reduce dependence when critical functions share underlying infrastructure, software, or access conditions. Meaningful autonomy requires independently verified replacement options, manageable switching costs, and the ability to sustain essential functions under disruption.
- **South Korea's AI strategy should prioritize verifiable national conversion and operational resilience.** Public institutions should audit model, data, software, and hosting control paths; evaluate domestic and foreign models against equivalent institutional tasks and total operating costs; and conduct function-preserving substitution and continuity tests. These measures would help distinguish technological availability from nationally usable capacity and credible autonomy.



GLOBAL AI
GOVERNANCE
RESEARCH CENTER

Policy Brief

Executive Summary

South Korea combines globally consequential AI-memory manufacturing with domestic models, public-private compute allocation, advanced cloud and telecommunications organizations, and manufacturing applications. These strengths do not comprise a nationally controlled end-to-end AI stack. Advanced accelerators, semiconductor equipment, packaging relationships, model software, and overseas production sites create external interfaces over which Korean actors exercise only partial authority. The analytical question is how Korean institutions convert heterogeneous resources into functions they can access and sustain, while retaining credible alternatives where consequential interfaces are controlled elsewhere.

KIC provides a diagnostic institutional case. Its reported use of Chinese-developed language and retrieval models is compatible with substantial domestic model-development capacity and close U.S.-linked semiconductor and security relationships. Yet reported use volumes, the timing of its GPU expansion contract, and unknown system architecture prevent strong claims about which models are operationally indispensable, whether foreign actors possess continuing technical control, or whether cross-border adoption has become systemic influence (Han, 2026). The practical implication is not model-origin neutrality at any cost: it is function-based procurement subject to verifiable security, legal, and continuity controls.

Why This Matters

South Korea occupies a distinctive position in the emerging global AI power order. Its advanced semiconductor and memory industries make it a critical contributor to global AI infrastructure, while domestic model development, industrial applications, and public-sector initiatives demonstrate growing capacity beyond hardware production. Yet these strengths coexist with substantial dependence on externally controlled technologies and services. Understanding Korea's AI power therefore requires distinguishing specialized technological depth from nationally usable capacity, relational leverage, and operational autonomy.

The KIC case makes these distinctions particularly relevant. The reported adoption of Chinese-developed language and retrieval models alongside U.S.-developed components illustrates how institutional technology choices can cross geopolitical and technological boundaries. Such adoption does not establish foreign control, strategic realignment, or durable dependence. It instead raises questions about who controls critical interfaces, how models are integrated into institutional workflows, and whether essential functions can be maintained when access conditions change.

These questions extend beyond South Korea. As AI ecosystems become increasingly interconnected, technological leadership, supply-chain importance, institutional adoption, and boundary control may be distributed across different actors. For policymakers, the central challenge is to convert specialized resources into sustainable national capabilities while preserving credible alternatives at critical interfaces. The Korean case thus provides a useful lens for understanding how AI power and strategic autonomy develop under uneven technological interdependence.

Policy Brief

Analytical Scope and Evidence Rules

The framework distinguishes four objects. **AI National Power (ANP)** is the usable capacity of a national system to mobilize, govern, and sustain AI-related resources for specified functions. **AI Boundary Power (ABP)** is an actor-level relational capacity to alter another actor's conditions of access, integration, or continued use at a specified interface. **Boundary Autonomy** concerns the target's credible ability to switch, bypass, substitute, diversify, or internalize while preserving the function. **Systemic Influence** is a later outcome requiring persistent, sufficiently broad, and credibly attributable changes to others' feasible options, network arrangements, or governing rules (Wu, 2026a, 2026b). None is a synonym for model performance or asset ownership.

For each boundary relation, three evidentiary states are kept separate: **available control capacity** (enforceable rights and practical control over a consequential interface); **exercised control** (an observed change to access or admissibility); and **observed target effects** (measurable responses attributable to that change). Evidence of the latter strengthens a causal account but is not automatically a definitional prerequisite for the existence of a capacity. Likewise, a lack of observed intervention does not demonstrate that no intervention path exists.

Maturity codes in this brief distinguish operating or repeatedly allocated capacity from procured resources, infrastructure under construction, announced partnerships, and policy objectives. Official assessments and corporate reports are attributed to their publishers, not silently treated as independent evaluations. The KIC evidence principally follows the chain **KIC-submitted material → legislator's disclosure/comment → Yonhap's report → this analysis**. The underlying selection scores, contracts, audit logs, data-flow maps, and investment-decision records have not been independently inspected. KIC is a diagnostic case, not a representative sample of Korean public-sector procurement.

1. Korea's National AI-Power Configuration

1.1 High-Bandwidth Memory Within a Distributed Production Architecture

South Korea's most visible global AI specialization is advanced memory. Samsung announced commercial HBM4 mass production and initial customer shipments in February 2026, while SK hynix reported the beginning of HBM4 mass shipments during the second quarter (Samsung Electronics, 2026; SK hynix, 2026). These developments establish product-specific commercial capabilities rather than guaranteed supply at every future volume or performance requirement. Their strategic significance derives from HBM's critical role in advanced accelerator systems, where performance requirements, customer qualification, packaging integration, and production schedules can constrain short-term substitutability.

These capabilities operate within a distributed production architecture. SK hynix's collaboration with TSMC on HBM4 base-die development and advanced packaging illustrates how Korean memory production depends on complementary technologies, manufacturing services, and customer specifications across jurisdictions (SK hynix, 2024). Korean suppliers may possess consequential contractual rights over production allocation, delivery schedules, and customer qualification processes, but such rights are interface-specific and do not necessarily permit unilateral changes to customer access.

Policy Brief

Three evidentiary distinctions are therefore necessary. Available boundary-control capacity depends on enforceable rights, interface criticality, and the availability of functional substitutes. Exercised control requires evidence that a controller actually changed specified access, allocation, or participation conditions. Observed effects concern identifiable target responses attributable to that intervention. The reviewed evidence establishes Korea's specialized supply capacity and strategic industrial importance, but does not demonstrate a Korean-initiated restriction against a named foreign customer. Commercial bargaining, temporary scarcity, and deliberate boundary intervention must consequently remain analytically distinct.

1.2 Compute: An Operating Baseline and a Separately Coded Expansion Pipeline

Korea has an operational public-access compute baseline that predates its largest expansion announcements. MSIT reported in February 2025 that 416 H100 GPUs secured by government within the Gwangju AI Data Center's 880-H100 installation were **then being used** to support industry, academia, and research institutions (MSIT, 2025e). This is evidence of actual public-purpose use at the reported date, not an audited 2026 availability figure. The ministry later selected Naver Cloud, NHN Cloud, and Kakao as partners in a 2025 procurement intended to secure approximately 13,000 H200/B200 GPUs; the announcement specified phased deployment and expected government allocations rather than confirming that all equipment was already in service (MSIT, 2025a).

A second observable allocation channel is the Sovereign AI Foundation Model initiative. MSIT's August 2026 phase-two announcement described approximately 768 B200 units supporting participating teams in the first half of 2026 and a planned rise to approximately 1,000 in the second half (MSIT, 2026d). These are program-specific support figures, not an independently established nationwide inventory, and they must not be added to overlapping public procurement totals without asset-level reconciliation.

The National AI Computing Center is different again. Earlier plans created a procurement and public-private governance framework (MSIT, 2025c). According to MSIT, Korea AI Computing Center (KOACC) was established on June 10, 2026 by a public-private consortium led by Samsung SDS, ground was broken at Haenam on August 3, and the project aims to equip the center with approximately 15,000 AI chips by 2028. MSIT envisaged some advance service in 2027 through Samsung SDS data centers. This establishes a governance and construction pathway, **not an operational 15,000-chip pool in October 2026** (MSIT, 2026b). The relevant ANP test is not the headline GPU count but allocable GPU-hours by hardware type, admission rules, queue times, delivery continuity, and demonstrated domestic workloads.

1.3 Domestic Models: Development, Project Eligibility, and Procurement Are Distinct

Domestic model development is supported by a national policy architecture that extends beyond individual model competitions (MSIT, 2025d). It is substantive but must not be collapsed into a single 'sovereign-model' ranking. MSIT initially selected five teams led by NAVER Cloud, Upstage, SK Telecom, NC AI, and LG AI Research. The first-stage review advanced LG, SK Telecom, and Upstage; NAVER's result also involved the program's independent-model qualification standard, and MSIT called for an additional team. Following that supplementary selection, four teams entered the next stage, and the second-stage evaluation advanced three: Upstage, SK Telecom, and LG AI Research (MSIT, 2025b, 2026c, 2026d). Model and benchmark descriptions in these notices are program-reported and should not be read as independently established parity on all institutional tasks.

Policy Brief

Three claims require separate evidence: a **Korean-developed model** exists; a model **meets government criteria for a sovereign-model initiative**; and a model **outperforms feasible alternatives for a particular institution**. National research and training support strengthens the first two pathways, while repeat task-based procurement, reliability, integration costs, and workload continuity establish the third. KIC makes the distinction concrete without showing that domestic models are generally uncompetitive: there is no publicly available matched comparison of its chosen model variants against contemporary Korean alternatives on identical tasks and hardware.

1.4 Industrial and Public-Sector Conversion

In August 2026, the Ministry of Trade, Industry and Resources (MOTIR) reported support across more than 170 manufacturing sites under its M.AX industrial AI effort. Within a subset of 42 relatively early participating sites with measurable results, the ministry reported a mean productivity increase of 30.1% and a mean defect-rate reduction of 15.5% (MOTIR, 2026). This is evidence that concrete process applications have been implemented at selected sites. Without a consistent baseline, control group, full participant accounting, and follow-up, it is not a nationally representative causal effect or proof that the reported gains will persist.

Administrative conversion has a narrower observed basis. In July 2026, the Ministry of the Interior and Safety announced the pilot of an AI legal assistant developed for public officials, using institutional legal materials for assistance rather than replacing accountable human legal judgment (MOIS, 2026). An announced pilot and defined workflow support evidence of an implementation pathway; actual accuracy, active users, processing-time reductions, and maintenance obligations require separate evaluation. A similar caution applies to SK Telecom’s July 2026 announcement that existing A-One and Spider systems would support a summer network-management program targeting approximately 1,100 locations: a forward-looking deployment announcement is not itself an audited post-event coverage result (SK Telecom, 2026).

Table 1. Selected Evidence of Korean Capacity Conversion

Layer and Interface	Publicly Supported State	What Remains Unestablished
AI memory	Firm-reported HBM4 commercial shipments (2026)	Interface-specific control rights, documented restrictions against named customers, available substitutes, and attributable customer responses.
Public compute	416 H100 government-secured units reported in use in Gwangju (2025); targeted model-project GPU support (2026)	Current allocable hours, uptime, waiting times, overlap with procurement totals
National AI center	KOACC formed and Haenam construction initiated (2026); approximately 15,000 chips targeted for 2028	Routine availability of planned center resources
Domestic models	Multi-stage national evaluation and continuing development support	Repeat institutional wins on equivalent workloads and total-cost conditions
Industrial AI	Selected-site deployments and ministry-reported performance changes	Representative, sustained or independently estimated effects
Administrative AI	Public-sector legal-assistant pilot and specified support functions	Measured accuracy, throughput, adoption and continuity

Source. Author’s synthesis of MSIT (2025a, 2025b, 2025e, 2026b–d), Samsung Electronics (2026), SK hynix (2026), MOTIR (2026), and MOIS (2026). These are heterogeneous evidence states, not additive components of a single score.

Policy Brief

The reviewed evidence establishes selected pathways of national AI-capacity conversion, rather than comprehensive control over the Korean AI stack or demonstrated resilience across all critical functions. Organizational concentration may facilitate coordination, but whether it systematically shortens conversion time relative to more dispersed systems remains a hypothesis requiring comparison.

2. External Interfaces: Boundary Control and Boundary Autonomy

The direction and extent of control vary across Korea's AI technology stack. Korean HBM suppliers occupy consequential positions in advanced memory production, but their relational leverage must be assessed through specific contractual and operational interfaces rather than market importance alone. Qualification procedures, allocation decisions, delivery commitments, and potential supply restrictions involve different rights and constraints. Establishing AI Boundary Power requires identifying who can alter a specified crossing condition, whether that authority is enforceable, and how target-side substitutes condition its effectiveness. Actual interventions and resulting customer responses provide further evidence of exercised control and its consequences.

At semiconductor manufacturing and related investment interfaces, Korean firms are also subject to externally governed access conditions. On October 15, Reuters reported, citing sources familiar with the discussions, that SK hynix and its U.S. subsidiary Solidigm were seeking authorization to export additional U.S.-origin semiconductor equipment to the Dalian NAND flash memory facility in China, reportedly linking the request to a potential new NAND manufacturing investment in the United States. SK hynix stated that equipment installation at Dalian's Mod B facility was proceeding under U.S. authorization and denied seeking a capacity expansion there. The reporting does not establish approval of the additional licenses or a concluded investment arrangement (Reuters, 2026). The case illustrates how Korean corporate ownership, production in China, U.S. export-licensing authority, and investment decisions remain distinct but interconnected dimensions of technological interdependence.

Model and cloud relations present a different control structure. An external hosted API may retain per-request authentication and service termination rights. By contrast, lawfully acquired weights running under a genuinely independent hosting and update arrangement can reduce continuing dependence on the original model developer for inference. Neither condition follows solely from where a model was developed. Access to weights does not eliminate contractual, software-dependency, vulnerability-management, cloud, or GPU constraints. Korean public authorities also retain procurement and data-governance powers: the AI Basic Act's entry into force on January 22, 2026 establishes a domestic legal framework, but not an automatic restriction on every foreign-developed model or evidence of a particular procurement exclusion (MSIT, 2026a).

Policy Brief

3. KIC as a Diagnostic Case of Mixed-Origin Institutional AI

3.1 Reported Design and Chronology

Yonhap reported on October 19, 2026, citing materials submitted by the Korea Investment Corporation (KIC) to National Assembly member Kang Min-kuk, that KIC's internal AI system incorporated Alibaba's Qwen3 for language processing, the Beijing Academy of Artificial Intelligence's (BAAI) BGE-M3 for retrieval embeddings, Meta's Llama 3 for harmful-content screening, and Google's Gemma 3 for image-related functions. Reported applications included document summarization, translation, internal rules and weekly reports, groupware retrieval, and support for parliamentary-question materials (Han, 2026). This represents a mixed-origin, functionally differentiated architecture, but does not establish that the models participate in investment selection, trade execution, or portfolio management.

The procurement chronology requires separate treatment. KIC reportedly introduced BGE-M3 on December 15, 2025. On May 7, 2026, it made its internal AI service available to approximately 350 employees and signed an infrastructure expansion contract worth KRW 950.4 million. The project was expected to increase GPU capacity from four to twelve units, with completion planned for November 2026. KIC subsequently explained that the expansion was intended to support an internal AI platform enhancement project commissioned in August, citing rising GPU and memory prices and equipment delivery lead times. The relationship between the May infrastructure contract and the later platform enhancement has not been independently established through the underlying procurement and project documents (Han, 2026).

Reported internal AI requests declined from 3,755 in May to 798 in September, a decrease of 78.7%, while recorded users fell from 339 to 147. In September, employees made 36,881 ChatGPT requests, approximately 46.2 times the internal-system count. However, KIC stated that its internal AI was restricted to nine categories of supporting tasks. Differences in task coverage and potentially in counting conventions prevent these figures from establishing the systems' relative operational importance, security exposure, or substitutability. Lower request volume does not necessarily imply lower task criticality or reduced demand for secure internal processing.

The model-selection process also became a subject of parliamentary scrutiny. According to Yonhap, Kang criticized the absence of developer nationality and U.S. export-control listing status from the submitted evaluation criteria, while questioning the scale of infrastructure expansion relative to recorded usage. KIC attributed its model choices to accuracy, reasoning, business suitability, hardware constraints, cost efficiency, available GPU memory, and multilingual performance. These competing positions raise legitimate procurement-governance and resource-allocation questions, but neither the underlying model evaluations nor the projected infrastructure requirements have been independently audited. The case therefore supports analysis of institutional selection and oversight rather than a definitive finding of security failure, regulatory noncompliance, or unnecessary expenditure.

Table 2. KIC Evidence Ledger: Observed Disclosure Versus Remaining Uncertainty

Question	Evidence From Yonhap’s Report	What the Evidence Does Not Establish
Model provenance and tasks	Qwen3 language processing; BGE-M3 retrieval; Llama 3 screening; Gemma 3 image functions	Exact model versions, deployment topology, or functional substitutability
Selection and security review	KIC cited performance, cost, memory, and multilingual suitability; a legislator criticized the omission of developer nationality and U.S. export-control status	Whether independent security assessments were completed or regulatory requirements violated
Procurement chronology	BGE-M3 introduced December 2025; KRW 950.4 million expansion contract signed May 7, 2026; platform enhancement commissioned in August	Complete contractual relationships, original demand forecasts, or subsequent procurement revisions
Recorded activity	Internal requests and users declined between May and September; ChatGPT requests were substantially higher; KIC identified nine supporting-task categories	Comparable workload importance, sensitive-data exposure, or relative cost-effectiveness
Continuing technical control	Public reporting did not disclose hosting architecture or administrative control paths	External data access, update privileges, service termination rights, or security compromise

Source. Author’s synthesis of Han (2026), based on KIC-submitted materials, the institution’s reported explanations, and parliamentary criticism. The table does not constitute an independent audit of KIC’s procurement, system architecture, or operational performance.

3.2 Export Controls, Software Provenance, and Security Inference

BAAI’s inclusion on the U.S. Entity List is independently verifiable. A Bureau of Industry and Security final rule, effective March 25, 2025 and published March 28, imposed licensing requirements on specified transactions involving the institution and items subject to the Export Administration Regulations (BIS, 2025). The legislator’s criticism concerning KIC’s selection criteria raises a distinct procurement-governance question: whether developer provenance and applicable external restrictions received adequate consideration. However, the developer’s listing does not itself establish that KIC’s deployment of a separately obtainable model artifact violated U.S. export controls. A transaction-specific compliance determination would require additional evidence concerning the parties, items, licenses, and applicable jurisdiction.

Public Qwen3 and BGE-M3 model repositories illustrate how downloadable model artifacts can circulate beyond the developer’s operating environment (Qwen Team, 2025; BAAI, n.d.). The cited Qwen3-8B model card is evidence concerning one publicly released variant, not proof that KIC deployed 8B. The public report does not identify KIC’s precise variants, model revisions, hosting sites, patching privileges, service providers, or license chain. Three states must remain separate: architecture undisclosed, external control not observed in the reviewed public record, and external control independently ruled out. Only the first two are supported here. Model origin alone neither proves data transfer to China nor negates plausible software and supply-chain exposure.

The strongest defensible case classification is therefore reported adoption and use of foreign-developed model components for specified internal support functions, with continuing control, switching costs, and measured organizational benefits unresolved. This is a diagnostic observation of institutional selection, not evidence of a security breach, demonstrated Chinese ABP over KIC, or China-driven systemic influence in Korea.

Policy Brief

4. Competing Mechanisms and the Autonomy Test

Three interpretations merit testing rather than assuming. Task-constrained selection predicts that an organization may choose a foreign model because language coverage, inference memory, support-task accuracy, or cost meet a local workload better than available alternatives. KIC cited several of these considerations, but its evaluation files and matched Korean-model benchmarks are unavailable. Artifact diffusion predicts that transferable weights may broaden adoption without preserving the original developer’s ability to authorize every subsequent inference; the actual hosting and maintenance contract determines whether this mechanism applies to KIC. Institutional boundary management predicts that procurement, authentication, data placement, and replacement testing shape which dependencies persist after adoption.

A fourth consideration prevents over-reading the mixed pipeline: functional diversity within a pipeline does not establish substitutability within any one function. Llama 3 screening and Gemma 3 image processing are not demonstrated alternatives to Qwen3 text generation or BGE-M3 embeddings. A proper exit test would separately assess language-model migration; re-embedding the document corpus and rebuilding retrieval indexes where required; permission and metadata preservation; accuracy regression; acceptable latency and unit costs; security evaluation; and measured recovery after provider or component interruption. Shared GPU, cloud, or software dependencies must be included in the same test.

5. Implementation Priorities and Verification

The following actions are author-proposed timelines, not existing government mandates. Sensitive system diagrams and vendor privileges should be maintained for authorized audit while publicly released reporting can disclose methods and aggregated results. Local hosting can increase maintenance burden, model diversity increases integration costs, and additional security checks can slow procurement; these costs should be measured alongside the resulting reduction in exposure.

Table 3. Proposed Implementation Responsibilities and Deliverables

Responsible Actor	Proposed Horizon	Action and Accountable Deliverable	Verification Evidence
Public-institution CIO, security and procurement leads	90 days	Register critical AI functions and map models, data flows, licenses, operators, update privileges and hosting jurisdictions	Versioned asset register; independently reviewable control-path diagram
Program ministry and independent evaluation team	6 months	Benchmark Korean and foreign candidates on the same authorized tasks and hardware budgets	Accuracy, failure rate, latency, human rework, cost per task, security tests
Critical workflow owners	6–12 months	Run model and embedding replacements and supplier-interruption exercises	Functional retention, index rebuild effort, recovery time and migration cost
Public compute program operators	Quarterly	Report resource maturity and delivered service, retaining unique project and asset identifiers	GPU type, allocable hours, admission rules, queue time, utilization, outages
Industry ministry and participating firms	Annually	Review named HBM, equipment and model-service interfaces for functional substitutes and external restrictions	Controller, interface, instrument, change, response, competing explanations

Source. Author’s implementation proposal based on the evidence limits identified above; proposed horizons and responsibilities are normative recommendations, not empirical claims.

Policy Brief

6. Political Implications

South Korea's AI position illustrates why geopolitical alignment and technological interdependence cannot be understood through exclusive blocs. Its strategic importance in advanced semiconductor supply chains coexists with reliance on externally controlled computing, software, and model ecosystems. This asymmetric interdependence creates distinct forms of leverage and vulnerability across technological interfaces, making supply-chain importance, national capability, and relational control analytically separate.

The KIC case brings this complexity into institutional practice. Its reported adoption of Chinese-developed models alongside U.S.-developed components demonstrates that procurement choices need not follow geopolitical alignment. Such adoption, however, does not establish strategic realignment, foreign control, or durable operational dependence. Its political significance lies in the distinction between national technological priorities and the functional, economic, and security considerations shaping institutional decisions.

For **South Korea**, the central challenge is to convert specialized technological strengths into sustainable national capabilities while preserving control over critical functions. Strategic autonomy does not require complete technological self-sufficiency, but it depends on credible substitution options, institutional control over sensitive operations, and resilience against changes in external access conditions.

For **the United States**, the case highlights the limits of assuming that security partnerships and upstream technology restrictions necessarily determine downstream model adoption. Deployable model weights can reduce continuing dependence on original developers, although infrastructure, licensing, and software dependencies may persist. Effective boundary control must therefore be assessed at specific interfaces rather than inferred from geopolitical alignment.

For **China**, the KIC case illustrates one pathway through which Chinese-developed AI models can enter institutions within a U.S.-aligned technological and security environment. Such adoption demonstrates the potential for technological diffusion across geopolitical divisions but does not itself establish AI Boundary Power or systemic influence. Whether broader diffusion translates into durable relational leverage depends on identifiable control rights, institutional dependence, switching costs, and observable changes in technological relationships.

For **other technologically interdependent states**, South Korea illustrates both the opportunities and limitations of combining capabilities from competing ecosystems. Such arrangements can expand technological options, but supplier diversity alone does not guarantee autonomy when critical functions remain exposed to shared infrastructure or externally controlled interfaces.

Taken together, these implications suggest that the emerging global AI order is shaped less by exclusive technological blocs than by overlapping, function-specific relationships. Its political consequences depend on how states convert technological resources into usable capacity, exercise control at consequential interfaces, and preserve autonomy under changing conditions of interdependence.

Policy Brief

7. Five Monitoring Priorities Through 2030

Five indicators can help assess whether South Korea's specialized technological strengths translate into durable national capacity, boundary autonomy, and external influence.

- **Compute conversion:** Track contracted, delivered, operational, and actually utilized computing resources, distinguishing infrastructure commitments from allocable capacity.
- **Domestic-model adoption:** Evaluate Korean-developed models through institutional procurement, sustained use, and renewal, accounting for task performance, cost, security, and foreign alternatives.
- **Industrial AI durability:** Measure productivity and quality improvements across participating manufacturing sites, including unsuccessful deployments and discontinued projects.
- **Boundary autonomy:** Test whether critical functions can survive model withdrawal, API restrictions, or infrastructure disruption through verified substitution, recovery time, and switching costs.
- **External boundary leverage:** Document specific control rights, their exercise at consequential interfaces, target responses, and available alternatives to distinguish supply-chain importance from effective relational power.

Each assessment should establish its baseline, observation period, and evaluation criteria in advance. Evidence of systemic influence requires additional demonstration of persistent, sufficiently broad, and credibly attributable effects beyond individual transactions or institutional choices.

8. Limitations

This policy brief offers a selective, theory-guided assessment of South Korea's AI power rather than a comprehensive national evaluation. Evidence varies across technological domains, and reported investments, deployments, and project outcomes do not consistently establish sustained operational capacity or resilience under disruption.

The KIC case relies primarily on secondary reporting rather than direct access to procurement contracts, system architecture, or operational records. Available evidence confirms reported model adoption but does not establish external control, functional dependence, switching costs, or organizational benefits.

Finally, semiconductor supply-chain importance does not itself demonstrate AI Boundary Power, and a single institutional case cannot establish broader adoption patterns or systemic influence. Stronger conclusions require longitudinal, interface-specific, and independently verified evidence.

Policy Brief

Conclusion

South Korea’s position in the emerging global AI power order reflects the interaction between advanced semiconductor specialization, growing domestic AI capabilities, and continued dependence on externally controlled technological interfaces. Its importance in global memory supply chains does not automatically establish comprehensive national capacity, unilateral boundary control, or operational autonomy.

The KIC case illustrates how institutional AI adoption can cross geopolitical and technological boundaries without necessarily producing foreign control or durable dependence. Its significance lies in the conditions governing deployment, continued access, and functional substitution rather than model provenance alone.

South Korea’s future AI power will therefore depend on its ability to convert specialized resources into sustainable national functions, exercise effective control at consequential interfaces, and preserve credible alternatives under disruption. These outcomes, rather than investment announcements or isolated institutional choices, provide the basis for assessing its evolving role in the global AI order.

Policy Brief

References

Beijing Academy of Artificial Intelligence (BAAI). (n.d.). *BGE-M3 model card* [Model repository; version may change]. Hugging Face. <https://huggingface.co/BAAI/bge-m3>

Bureau of Industry and Security (BIS). (2025, March 28). *Additions to the Entity List* (Effective March 25, 2025). *Federal Register*, 90, 14046–14052. <https://www.federalregister.gov/documents/2025/03/28/2025-05427/additions-to-the-entity-list>

Han, J.-h. (2026, October 19). 339조 굴리는 국부펀드, 사내AI 핵심은 중국산... “보안의식 허술” [Sovereign wealth fund managing KRW 339 trillion uses Chinese models for core internal AI]. *Yonhap News Agency*. <https://www.yna.co.kr/amp/view/AKR20261018027100002>

Ministry of the Interior and Safety (MOIS). (2026, July 13). 공무원이 직접 만든 ‘AI 법령 비서’ 시범 개시 [Pilot launch of an AI legal assistant developed by public officials; author translation]. https://www.mois.go.kr/frt/bbs/type010/commonSelectBoardArticle.do?bbsId=BBSMSTR_000000000008&nttId=127770

Ministry of Science and ICT (MSIT). (2025a). *GPU Procurement Project (First Supplementary Budget, KRW 1.46 Trillion) Selects Participants*. <https://www.msit.go.kr/eng/bbs/view.do?bbsSeqNo=42&mId=4&mPid=2&nttSeqNo=1155&sCode=eng>

Ministry of Science and ICT (MSIT). (2025b). *MSIT Selects Five Elite Teams for the Sovereign AI Foundation Model Project*. <https://www.msit.go.kr/eng/bbs/view.do?bbsSeqNo=42&mId=4&mPid=2&nttSeqNo=1152&sCode=eng>

Ministry of Science and ICT (MSIT). (2025c). *MSIT to Establish the National AI Computing Center as a “Catalyst for Genuine Growth in AI”*. <https://www.msit.go.kr/eng/bbs/view.do?bbsSeqNo=42&mId=4&mPid=2&nttSeqNo=1166&sCode=eng>

Ministry of Science and ICT (MSIT). (2025d). *National AI Strategy Committee launched as Korea’s top-level AI policy body*. <https://www.msit.go.kr/eng/bbs/view.do?bbsSeqNo=42&mId=4&mPid=2&nttSeqNo=1165&sCode=eng>

Ministry of Science and ICT (MSIT). (2025e, February 20). *Korea to Expand AI Computing Infrastructure to Strengthen National AI Capabilities and Achieve Global Leadership*. <https://www.msit.go.kr/eng/bbs/view.do?bbsSeqNo=42&mId=4&mPid=2&nttSeqNo=1092&sCode=eng>

Ministry of Science and ICT (MSIT). (2026a). *AI Basic Act and Enforcement Decree Enter into Force*. <https://www.msit.go.kr/eng/bbs/view.do?bbsSeqNo=42&mId=4&nttSeqNo=1214&sCode=eng>

Ministry of Science and ICT (MSIT). (2026b). *National AI Computing Center, a Key Infrastructure for Korea’s AI Highway, Breaks Ground in Haenam*. <https://www.msit.go.kr/eng/bbs/view.do?bbsSeqNo=42&mId=4&mPid=2&nttSeqNo=1291&sCode=eng>

Ministry of Science and ICT (MSIT). (2026c). *Phase 1 Evaluation Results Released for Sovereign AI Foundation Model Project*. <https://www.msit.go.kr/eng/bbs/view.do?bbsSeqNo=42&mId=4&mPid=2&nttSeqNo=1212&sCode=eng>

Ministry of Science and ICT (MSIT). (2026d). *Phase 2 Evaluation Results Unveiled for the Sovereign AI Foundation Model Project*. <https://www.msit.go.kr/eng/bbs/view.do?bbsSeqNo=42&mId=4&mPid=2&nttSeqNo=1292&sCode=eng>

Ministry of Trade, Industry and Resources (MOTIR). (2026, August 5). *M.AX Delivers Tangible Results on the Factory Floor* [Web publication; briefing described as August 4]. <https://english.motir.go.kr/eng/article/EATCLdfa319ada/2703/view>

Qwen Team. (2025). *Qwen3-8B model card* [A public variant; does not establish the variant used by KIC]. Hugging Face. <https://huggingface.co/Qwen/Qwen3-8B>

Reuters. (2026, October 15). *SK Hynix conditions new US factory on more US chip equipment exports to China, sources say* [Source-based news report]. <https://www.reuters.com/world/asia-pacific/sk-hynix-conditions-new-us-factory-more-us-chip-equipment-exports-china-sources-2026-10-15/>

Policy Brief

Samsung Electronics. (2026, February 12). *Samsung Ships Industry-First Commercial HBM4 with Ultimate Performance for AI Computing*. <https://news.samsung.com/global/samsung-ships-industry-first-commercial-hbm4-with-ultimate-performance-for-ai-computing>

SK hynix. (2024). *SK hynix Partners with TSMC to Strengthen HBM Technological Leadership*. <https://news.skhynix.com/en/sk-hynix-partners-with-tsmc-to-strengthen-hbm-technological-leadership/>

SK hynix. (2026, July 29). *SK hynix Announces 2Q26 Financial Results*. <https://news.skhynix.com/en/q2-2026-business-results/>

SK Telecom. (2026, July 16). *SKT, AI로 휴가철 통신 품질 미리 살핀다* [SK Telecom plans AI-supported summer network quality monitoring; author translation]. <https://news.sktelecom.com/227991>

Wu, S. (2026a). *Beyond centrality: Conditional connectivity and boundary control in interdependent networks* (EPINOVA Working Paper No. EPINOVA-WP-T-2026-01). Global AI Governance and Policy Research Center, EPINOVA LLC. <https://doi.org/10.67037/epinova.wp.t.2026.01>

Wu, S. (2026b). *Two structures of AI power: National capacity, relational boundary control, and systemic influence in the emerging AI order* (EPINOVA Working Paper No. EPINOVA-WP-T-2026-02). Global AI Governance and Policy Research Center, EPINOVA LLC. <https://doi.org/10.67037/epinova.wp.t.2026.002>