

The Conversion Gap: How China, Japan and South Korea Turn AI Inputs into AI Capability, 2017–2025

Alfredo Merlet

alfredo.merlet@gmail.com

Universidad de Sevilla <https://orcid.org/0009-0007-8558-164X>

Research Article

Keywords: Artificial intelligence, National AI capability, Conversion efficiency, East Asia, Innovation policy, State policy capacity

Posted Date: October 19th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11390808/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Abstract

National strategies for artificial intelligence (AI) are usually compared by the resources they mobilise: researchers, research budgets, private investment and computing infrastructure. Far less is known about how efficiently those resources are converted into AI capability, and why conversion differs among countries with comparable endowments. This article operationalises a three-layer view of national AI capability (invention, industrialisation and operationalisation) and measures conversion ratios for China, Japan and South Korea between 2017 and 2025, using the United States as the frontier reference. The analysis combines model-level data from Epoch AI, bibliometric data from OpenAlex, input indicators from the World Bank, investment, patent and diffusion data from the Stanford AI Index 2026, and pillar scores from the Government AI Readiness Index 2025. Three findings stand out. First, Japan is the lowest-converting case in almost every link: it produced 52 AI models in nine years against 87 for South Korea, despite larger research, R&D and compute endowments, and its compute-weighted conversion is an order of magnitude below that of its neighbours. Second, the Japanese gap lies in the knowledge-to-model and compute-to-model links, not in the capital link. Third, conversion aligns more closely with the state's policy capacity than with infrastructure or sector maturity: Japan and South Korea score almost identically on development and diffusion, but 70 versus 96 on policy capacity. The results are robust to fractional counting and compute weighting. The article argues that AI sovereignty is a problem of governing conversion, not only of accumulating inputs.

1 Introduction

Few areas of economic policy have produced as many national strategies in as short a time as artificial intelligence. Since China's State Council issued its New Generation Artificial Intelligence Development Plan in 2017, most advanced economies and many emerging ones have published plans that promise to build talent pipelines, expand computing capacity, mobilise private capital and govern the risks of the technology (State Council 2017; Oxford Insights 2026). Comparative assessments of these efforts typically rank countries by what they accumulate: researchers, publications, patents, venture investment, data centres. Composite indices, from government readiness scores to competitiveness rankings, aggregate such endowments into a single position (Oxford Insights 2026; Hu et al. 2026). What these exercises rarely ask is how efficiently each country turns its endowments into AI capability, and why that efficiency differs.

The question matters because endowments and outcomes can diverge sharply. Japan is the clearest example. It is among the world's largest spenders on research and development, it hosts more AI computing capacity than South Korea, and its AI research output exceeds South Korea's in absolute terms. Yet between 2017 and 2025 Japanese organisations released 52 AI models recorded in the Epoch AI database, against 87 for South Korean organisations, and only six notable models, against 23 (Epoch AI 2026a, 2026b). In 2024 and 2025 no Japanese organisation released a model that Epoch classifies as notable. China, meanwhile, moved from a distant follower to the only country that rivals the United States in frontier model development. Three East Asian economies with long traditions of state-

guided technological catch-up have therefore followed strikingly different trajectories in the same technology over the same period.

This article asks why three East Asian economies with substantial AI endowments convert those endowments into AI capability with such different efficiency. It builds on a recent conceptual contribution by Tassos (2026), who argues that national AI capability depends on converting technological advantage through three layers: invention, industrialisation and operationalisation. Tassos's framework is qualitative and is applied to the United States, China and the European Union. The present article operationalises the conversion idea quantitatively, extends it to Japan and South Korea, and connects conversion efficiency to the institutional capacity of the state.

The contribution is threefold. First, the article proposes and computes layer-specific conversion ratios: outputs of each layer divided by the inputs feeding it, lagged and normalised to the United States as the frontier reference. This distinguishes the analysis from production-frontier studies that relate AI investment to aggregate output (Sun et al. 2024), from patent-based studies of how AI enters national knowledge bases (Leusin et al. 2025), and from composite rankings that do not separate inputs from outputs (Hu et al. 2026). Second, it shows that the Japanese case is not an artefact of a single indicator: Japan's low conversion appears in publications per researcher, patents per researcher, models per publication and models per unit of compute, and survives fractional counting and compute weighting. Third, it shows that conversion aligns with the policy-capacity dimension of government AI readiness rather than with infrastructure or sector maturity, which reframes the debate on AI sovereignty as a problem of governing conversion rather than accumulating inputs.

The remainder of the article is organised as follows. Section 2 reviews the literatures on national innovation systems, the developmental state and AI capability, and derives three propositions. Section 3 describes the data, the conversion ratios and the robustness design. Section 4 presents the results layer by layer. Section 5 interprets the three national configurations and draws implications for emerging economies. Section 6 concludes.

2 Literature and analytical framework

2.1 From national innovation systems to national AI capability

The idea that countries differ systematically in their ability to turn knowledge into economic performance is central to the national innovation systems literature. Freeman (1987) introduced the concept through the Japanese case, arguing that Japan's post-war performance rested less on the volume of research than on the way firms, ministries and universities organised learning and diffusion. Lundvall (1992) and Nelson (1993) generalised the argument: innovative performance depends on interactions among institutions, not only on the stock of inputs. Furman et al. (2002) formalised national innovative capacity as the joint product of a common innovation infrastructure, cluster-specific conditions and the quality of

linkages between them, while Fagerberg and Srholec (2008) showed that the capability to absorb and use technology is a distinct dimension of national systems. At the level of the firm, the notion of absorptive capacity (Cohen and Levinthal 1990) makes the same point: what an organisation can do with external knowledge depends on prior, organisation-specific investments.

The East Asian literature added the state as a central actor. Johnson (1982) described Japan's Ministry of International Trade and Industry as the architect of industrial upgrading. Amsden (1989) and Wade (1990) showed how South Korea and Taiwan disciplined business groups through targeted credit and performance standards. Evans (1995) captured the institutional condition for such policies as embedded autonomy: a bureaucracy cohesive enough to pursue collective goals and connected enough to business to implement them. Kim (1997) and Lee (2013) documented how Korean firms moved from imitation to innovation by entering short-cycle technologies where incumbency advantages were weak. Abramovitz's (1986) notion of social capability, the institutional preconditions for exploiting a technological gap, runs through all of these accounts.

Research on AI has only partly connected with this tradition. Economists have treated AI as a general-purpose technology whose productivity effects arrive with a lag because complementary investments must first be made (Brynjolfsson et al. 2021), and as a new method of invention that changes the production of knowledge itself (Cockburn et al. 2019). Studies of AI policy have catalogued national strategies and regulatory approaches, often contrasting the state-led Chinese model with market-led and rights-based alternatives (Roberts et al. 2021). Ding (2024) argues that the decisive variable in technological competition among great powers is not the capacity to generate innovations but the capacity to diffuse them across the economy. These contributions share an emphasis on what happens after inputs are assembled, but they have not produced a comparable measure of conversion across countries.

Measurement of national AI capability has followed two routes. The first builds composite indices that aggregate many indicators of readiness, investment, talent and research into a single score, as in the Government AI Readiness Index, the AI Index's country comparisons and recent competitiveness rankings (Oxford Insights 2026; Stanford HAI 2026; Hu et al. 2026). Such indices are useful for benchmarking but, by construction, add inputs and outputs together, so a country can rank highly because it accumulates resources rather than because it uses them well. The second route relies on model-level databases that record the systems developed in each country, together with their training compute (Epoch AI 2026a). This route has shown that the training compute of leading models has doubled roughly every six months since the early 2010s (Sevilla et al. 2022), which makes compute an input in its own right and makes unweighted counts of models a potentially misleading measure of output. Neither route, on its own, separates what a country has from what it does with it.

2.2 Conversion as an analytical object

Tassos (2026) provides the most direct conceptual precedent. He defines national AI capability as the outcome of converting technological potential through three layers: invention, in which knowledge and

talent generate new techniques; industrialisation, in which those techniques become deployable models and products; and operationalisation, in which they are absorbed by organisations and institutions. The framework is valuable because it separates the possession of AI resources from their conversion. It is, however, qualitative and applied to the United States, China and the European Union, so it neither measures conversion nor examines the East Asian variety of state-guided systems.

Quantitative work has approached the problem from other angles. Sun et al. (2024) estimate a stochastic production frontier for AI investment across countries and find wide efficiency differences, with China among the least efficient, but their output is aggregate economic production rather than AI capability, and they do not decompose the process into layers. Leusin et al. (2025) use patent data for China, Japan, South Korea and the United States to show how local technological specialisations shape whether AI enters a national knowledge base by breaking into existing fields or by creating new ones; their focus is the integration pathway, not efficiency. Hu et al. (2026) construct an entropy-weighted AI competitiveness index for ten countries, and Yang and Li (2026) compare the role of higher education in the AI strategies of China, Japan and South Korea through policy documents. None of these studies measures, layer by layer, how efficiently the three East Asian economies convert AI inputs into AI outputs, or relates that efficiency to the institutional capacity of the state.

2.3 Policy capacity and three East Asian configurations

Policy capacity refers to the competences and resources a government needs to formulate and implement policy: analytical capacity to identify problems and options, operational capacity to coordinate implementation, and political capacity to secure support (Wu et al. 2015; Painter and Pierre 2005). In a fast-moving technology such as AI, policy capacity determines whether a government can align funding, compute access, data rules and talent policy quickly enough to influence where and how models are built. It is distinct from the infrastructure a country hosts and from the maturity of its private AI sector, which the Government AI Readiness Index measures as separate pillars (Oxford Insights 2026).

The three countries offer contrasting configurations of these elements. China combines very large inputs with a state that sets explicit targets, channels capital and coordinates local governments behind a national plan (State Council 2017; Roberts et al. 2021). South Korea combines a medium-sized input base with a highly coordinated relationship between the state and large business groups, a long record of entering short-cycle technologies (Lee 2013), and an assertive AI policy that culminated in the Basic AI Act, in force since January 2026 (Oxford Insights 2026). Japan has a large input base and deep industrial capabilities, but its AI governance has relied on soft law, guidelines and inter-ministerial coordination rather than binding targets or concentrated programmes (Habuka 2023); a dedicated AI Promotion Act and a central AI Strategy Headquarters were only established in 2025 (Oxford Insights 2026). These descriptions motivate three propositions.

Proposition 1

(P1). Conversion efficiency differs systematically among China, Japan and South Korea beyond differences in input levels, with Japan as the lowest-converting case.

Proposition 2

(P2). The Japanese conversion gap is concentrated in the links from knowledge to models and from compute to models, rather than in the link from capital to models.

Proposition 3

(P3). Cross-country differences in conversion align with the policy capacity of the state more closely than with infrastructure endowments or the maturity of the domestic AI sector.

Because the analysis covers three focal countries and one reference country, these propositions are examined through a structured, focused comparison with quantitative indicators, in the tradition of comparative political economy, rather than through statistical hypothesis tests. Figure 1 summarises the framework.

3 Data and method

3.1 Design and sources

The analysis covers China, Japan and South Korea, with the United States as the frontier reference, over the output period 2017–2025. The start year coincides with China's national AI plan and with the beginning of the deep-learning scaling era; the end year is the last complete year available. Data for 2026 are excluded because they cover only part of the year. All series were obtained programmatically from public sources between 14 and 15 October 2026, and every file is recorded with its source URL, download date and SHA-256 hash in a manifest that accompanies the replication package.

Six sources are combined. Model-level data come from three Epoch AI databases: all AI models, notable AI models and large-scale models (Epoch AI 2026a, 2026b). Each record identifies the country of the developing organisation, the publication date and, where available, the training compute in floating-point operations. Installed AI compute comes from the Epoch AI GPU Clusters database, which records clusters by country, first operational date and performance in H100-equivalents (Epoch AI 2025). Bibliometric data come from OpenAlex (Priem et al. 2022), queried for works tagged with the concept artificial intelligence and affiliated with institutions in each country, both in total and restricted to works with more than 50 citations. Input indicators come from the World Bank World Development Indicators (World Bank 2026): R&D expenditure as a share of GDP, researchers per million inhabitants, population and GDP at purchasing power parity. Private AI investment, granted AI patents and population-level AI diffusion come from the public data of the AI Index Report 2026 (Stanford HAI 2026), whose diffusion figures originate in the Microsoft AI Economy Institute (2026). Finally, pillar scores come from the

3.2 Conversion ratios

A conversion ratio is defined as the output of a layer divided by the input that feeds it, measured over the same period and expressed relative to the United States. Inputs are lagged by one year, so that outputs in the window 2017–2020 are related to inputs in 2016–2019, and so on. Values above one indicate that a country converts more output per unit of input than the United States; values below one indicate the reverse. Normalising to the frontier removes units and allows ratios built from different indicators to be compared on one scale. Table 1 lists the operationalisation.

Table 1
Operationalisation of the conversion layers

Layer and link	Input (lagged one year)	Output	Sources
1. Invention	Researchers in R&D; R&D expenditure (PPP)	AI publications; share with > 50 citations (to 2022); granted AI patents	World Bank; OpenAlex; Stanford HAI
2a. Knowledge → models	AI publications	AI models (all; notable)	OpenAlex; Epoch AI
2b. Capital → models	Cumulative private AI investment, 2013–2025	AI models (all)	Stanford HAI; Epoch AI
2c. Compute → models	Installed AI compute (H100-equivalents)	Large-scale AI models	Epoch AI
3. Operationalisation	–	Share of working-age population using generative AI, H2 2025	Microsoft AI Economy Institute via Stanford HAI
Institutional context	Government AI Readiness Index 2025 pillars	–	Oxford Insights
<i>Note: all ratios are expressed relative to the United States. Inputs are lagged one year relative to outputs except where only cumulative or single-year values are published.</i>			

Three design choices deserve justification. The one-year lag reflects the typical interval between the research that underpins a model and its public release; longer lags would shorten the usable panel without changing the cross-country ordering, because inputs move slowly relative to outputs. The United States is used as the reference because it defines the technological frontier in every output series, not because its configuration is normatively preferable. Finally, simple ratios are preferred to frontier-estimation techniques such as data envelopment analysis or stochastic frontier models, because with four countries and several inputs such techniques would place most units on the frontier by construction and would obscure the link-specific differences that are the object of the analysis.

Five ratios are computed for the full period. In the invention layer, AI publications per thousand researchers capture the efficiency with which the research workforce generates AI knowledge, and granted AI patents per researcher provide a complementary measure. In the industrialisation layer, three links are distinguished: AI models per ten thousand AI publications (knowledge to models), AI models per billion US dollars of cumulative private AI investment (capital to models), and large-scale models per ten thousand H100-equivalents of installed compute (compute to models). The operationalisation layer is described with population-level diffusion. Policy capacity and the other readiness pillars are treated as institutional context rather than as inputs.

3.3 Counting rules and robustness

Models developed by organisations from several countries are assigned in full to each participating country (whole counting) in the baseline. Because whole counting can inflate the output of countries that collaborate more often, the analysis is repeated with fractional counting, in which each model is divided equally among the countries listed. A second concern is that counting models treats a small fine-tuned model and a frontier system as equivalent, which mechanically penalises countries that concentrate resources in a few very large models. The analysis is therefore repeated with outputs weighted by training compute. Finally, conversion is examined across three windows, 2017–2020, 2021–2023 and 2024–2025, to check whether the pattern is stable over time. Annual counts for Japan and South Korea are too small to be interpreted year by year.

3.4 Limitations of the data

Five limitations should be stated at the outset. First, Epoch AI compiles models from publicly documented releases, which are predominantly reported in English; models documented only in Japanese, Korean or Chinese may be under-represented. This bias would affect all three focal countries, and Section 4.5 discusses why it is unlikely to explain the Japanese result. Second, the AI Index reports annual private investment by country only for the United States, China and the European Union; for Japan and South Korea only the cumulative 2013–2025 total and the 2025 value are published, so the capital link is computed for the period as a whole. Third, the GPU Clusters database covers a minority of global AI compute (Epoch AI 2025) and records only clusters located in each country, not compute rented abroad through cloud services. Fourth, World Bank R&D series end in 2022 or 2023; the last available value is carried forward for the final window and flagged. Fifth, the diffusion indicator is derived from Microsoft telemetry adjusted for platform market share (Microsoft AI Economy Institute 2026). Because China's AI market is dominated by domestic services, the Chinese value is not comparable and is excluded from the operationalisation layer. The Government AI Readiness Index is likewise measured in a single year, after the outputs it is compared with, and its authors note that China's score probably understates its capabilities (Oxford Insights 2026). These constraints make the analysis descriptive and comparative rather than causal.

4 Results

4.1 Input profiles

Table 2 summarises the inputs and outputs for 2017–2025. China leads the three focal countries on every input, with more researchers than the United States, almost as much cumulative R&D spending, and the largest stock of AI publications in the sample. Japan has the second-largest input base in East Asia: about 690,000 researchers, cumulative R&D expenditure of 1.7 trillion constant PPP dollars, 192,000 AI publications and around 5,900 H100-equivalents of installed AI compute by 2024. South Korea is smaller on every input except private AI investment, where its cumulative total of 10.8 billion dollars exceeds Japan's 7.0 billion. The output columns reverse the Japan–Korea ordering. South Korea produced 87 AI models, 23 notable models and 18 large-scale models; Japan produced 52, 6 and 4.

Table 2
Inputs and outputs, 2017–2025

Indicator	China	Japan	South Korea	United States
Researchers in R&D (mean, thousands)	2,408	690	451	1,543
R&D expenditure (cumulative, US\$ bn, constant PPP)	5,834	1,713	1,060	6,790
AI publications (thousands)	1,510	192	138	934
Installed AI compute, 2024 (H100-equivalents)	60,031	5,922	3,630	444,229
Private AI investment, 2013–2025 (US\$ bn)	131.8	7.0	10.8	757.3
AI models (all)	861	52	87	1,556
Notable AI models	147	6	23	478
Large-scale AI models	194	4	18	221
Granted AI patents, 2024 (approx.)	97,989	5,329	7,406	15,915
<i>Notes: inputs lagged one year (2016–2024). AI publications are whole-counted OpenAlex works tagged with the concept artificial intelligence. Models are whole-counted by country of the developing organisation. Patents are computed from the AI Index per-capita figures and World Bank population. Sources: Epoch AI (2025, 2026a, 2026b); OpenAlex; World Bank (2026); Stanford HAI (2026).</i>				

4.2 The invention layer

The first conversion link already separates the cases. China produces AI publications per researcher at the US rate (ratio 1.04), whereas Japan and South Korea produce roughly half as many (0.46 and 0.50). Patents tell a different story for South Korea. Using the AI Index figures for granted AI patents per 100,000 inhabitants and World Bank population data, China accounted for about 98,000 granted AI patents in 2024, South Korea for about 7,400, the United States for about 15,900 and Japan for about

5,300. Per researcher, South Korea converts at 1.59 times the US rate and China at 3.94, while Japan converts at 0.75. Japan is the only focal country below the frontier on both invention measures.

The quality of Japanese AI research is also lower on the citation measure. Within each publication year, the share of AI works that accumulated more than 50 citations was 5.3 per cent for Japan in 2017–2020, against 12.5 per cent for China, 13.3 per cent for South Korea and 13.9 per cent for the United States. In 2021–2022, the corresponding shares were 3.6, 8.2, 9.5 and 8.1 per cent. Because the comparison is made within the same year, all countries face the same citation window, so the gap is not an artefact of recency.

4.3 The industrialisation layer

Figure 2 presents the three industrialisation links together with the invention ratio. In the knowledge-to-model link, South Korea (0.39) and China (0.38) convert AI publications into models at similar rates, both well below the United States, while Japan converts at 0.16, less than half the rate of its neighbours. The gap widens for notable models: South Korea reaches 0.34 of the US rate, China 0.21 and Japan 0.06.

The capital-to-model link reverses the picture with respect to the United States: all three East Asian countries produce more models per dollar of private investment than the United States (China 3.18, Japan 3.62, South Korea 3.94). This reflects the concentration of US private capital in a small number of very expensive frontier systems, a point examined in the robustness analysis. Within East Asia, however, Japan is not the laggard in this link: it converts private capital into models slightly more efficiently than China. The Japanese gap is therefore not a shortage of capital efficiency.

The compute-to-model link is where the gap is most pronounced. Per unit of installed AI compute, South Korea produced large-scale models at 9.97 times the US rate and China at 6.50, while Japan reached 1.36. Japan has more installed AI compute than South Korea but produced four large-scale models against eighteen. Taken together, the industrialisation results support P2: Japan's conversion deficit is located in the links that turn knowledge and compute into models, not in the link that turns capital into models.

4.4 Conversion over time

Figure 3 tracks the knowledge-to-model ratio across the three windows. All countries improved between 2017–2020 and 2021–2023, as the diffusion of transformer architectures and open models lowered the cost of producing models. The trajectories diverge afterwards. China's ratio more than doubled again in 2024–2025, from 4.9 to 11.8 models per ten thousand lagged publications, and South Korea held a high level (9.1 and 9.0). Japan's ratio stalled at 3.7 and 3.8. In 2024–2025 Japan produced 17 models of any kind and no notable model, while its AI publication output remained above South Korea's. The divergence indicates that the Japanese gap is not a legacy of a late start but a persistent feature of the period in which generative AI became an industrial activity.

4.5 Robustness

Table 3 reports the robustness checks. Fractional counting leaves the knowledge-to-model ratios virtually unchanged (China 0.38, Japan 0.16, South Korea 0.40), so the results are not driven by international co-development. Weighting models by training compute changes the comparison with the United States but not the East Asian ordering. Once outputs are weighted by compute, the apparent advantage of the East Asian countries in the capital link disappears for Japan and South Korea (China 0.97, South Korea 0.60, Japan 0.10 of the US rate), confirming that the baseline capital ratio largely reflects the scale of US frontier models. In the compute link, compute-weighted output per installed H100-equivalent is at or above the US level for China (1.25) and South Korea (1.04) but only 0.07 for Japan. Japanese models are therefore not only fewer but also smaller.

Table 3
Robustness of the conversion ratios (United States = 1)

Measure	China	Japan	South Korea
Baseline: models per AI publication (whole counting)	0.38	0.16	0.39
Fractional counting: models per AI publication	0.38	0.16	0.40
Fractional counting: notable models per AI publication	0.21	0.06	0.33
Compute-weighted output per US\$ private investment	0.97	0.10	0.60
Compute-weighted output per installed H100-equivalent	1.25	0.07	1.04
Compute-weighted output per AI publication	0.12	0.01	0.06
Granted AI patents per researcher	3.94	0.75	1.59
<i>Notes: compute-weighted output is the sum of reported training compute (FLOP) of the models attributed to each country. Sources: as in Table 2; author's calculations.</i>			

Could the Japanese result reflect the language bias of model databases? Three considerations suggest it does not. First, the bias would affect South Korea at least as much as Japan, since Korean-language documentation is no more accessible than Japanese, yet South Korea converts at two to six times the Japanese rate. Second, the full Epoch database does record Japanese laboratories, including NTT, the Tokyo Institute of Technology, Preferred Networks, rinna, CyberAgent, SoftBank and Fujitsu, so Japanese developers are not invisible to it. Third, the invention-layer results, which are based on OpenAlex and patent data rather than on Epoch, show the same pattern. The AI Index count of notable models in 2025, which lists the United States, China and South Korea but not Japan, is consistent with these results but is not independent evidence, because it draws on the same Epoch data (Stanford HAI 2026).

4.6 The operationalisation layer and institutional alignment

Population-level diffusion offers a partial view of the third layer. In the second half of 2025, 30.7 per cent of the working-age population in South Korea used generative AI tools, compared with 28.3 per cent in the United States and 19.1 per cent in Japan (Microsoft AI Economy Institute 2026). South Korea's position is notable because its adoption rose sharply during 2025, a change the source attributes partly

to improvements in Korean-language models and to public policy. China is excluded for the reasons given in Section 3.4. Japan again occupies the lowest position among the comparable cases.

Figure 4 relates the notable-model conversion ratio to two pillars of the Government AI Readiness Index. On development and diffusion, which measures human capital, sector maturity and technology diffusion, Japan (66.1) and South Korea (65.9) are virtually identical, and Japan scores higher than South Korea on AI infrastructure (70.6 against 68.2). These pillars cannot account for a fivefold difference in conversion. On policy capacity, the gap is large: South Korea scores 96.0, China and the United States 92.5, and Japan 70.0, the lowest value in the sample by more than twenty points. Across the four countries, conversion is ordered as policy capacity would predict for the three East Asian cases, while the United States remains an outlier explained by its much larger private frontier sector. This pattern is consistent with P3. With four observations it cannot establish causation, but it locates the Japanese anomaly precisely in the dimension where Japan differs most from its neighbours.

The other pillars confirm that Japan is not weak across the board. Its scores for governance (80.5) and resilience (80.8) are close to South Korea's (90.7 and 82.7), and its public-sector adoption score (72.7) is far above China's (49.0). Japan's institutions for managing the risks of AI and for using digital tools in government are therefore comparable to those of its neighbours. What distinguishes Japan is the capacity to formulate and resource a coherent national AI policy, which is precisely the dimension most directly related to the links where its conversion breaks down.

5 Discussion

5.1 Three configurations of conversion

The evidence supports a reading of the three cases as distinct configurations of conversion rather than as points on a single scale of AI strength. China represents conversion through scale and direction. Its research workforce converts into publications and patents at or above the frontier rate, and its knowledge-to-model ratio accelerated sharply in 2024–2025, when Chinese developers released 547 models. The capital link shows that this was achieved with a small fraction of US private investment, consistent with the view that state coordination, local-government finance and a large domestic market substitute for venture capital in China's AI system (Roberts et al. 2021; Oxford Insights 2026). The Chinese weakness lies elsewhere: compute access is constrained by export controls, and public-sector adoption is the lowest of its readiness pillars.

The Chinese case also shows that East Asian conversion still trails the frontier in the industrialisation layer. Even after its acceleration, China produced 0.38 models per AI publication relative to the United States over the whole period, and 0.21 notable models. The US advantage in turning research into deployable systems remains large for all three countries, which suggests that the institutional and market conditions linking research to industrial model development, such as venture finance, platform

firms and researcher mobility between academia and industry, are where the frontier is most difficult to reach.

South Korea represents conversion through coordination. With the smallest input base of the three, it achieves the highest notable-model ratio in East Asia, the highest compute-to-model ratio in the sample, more granted AI patents per researcher than the United States and the highest population-level diffusion among the comparable cases. These outcomes fit a long-standing pattern in which a capable state and a small number of large business groups concentrate resources on short-cycle technologies where late entrants can compete (Amsden 1989; Lee 2013). South Korea's policy capacity score of 96, the highest of the four countries, is the institutional counterpart of this pattern.

Japan represents coordination without conversion. It possesses most of the ingredients that the national innovation systems literature associates with strong performance: a large and well-funded research system, capable industrial firms and more AI compute than South Korea. What it lacks, on the evidence assembled here, is the policy capacity to align these ingredients behind the production of AI models within the period studied. The irony is considerable, given that Freeman (1987) built the concept of a national innovation system on Japan's earlier ability to coordinate learning across firms and ministries. The institutional arrangements that served incremental improvement in manufacturing, relying on consensus, guidance and gradual adjustment, appear less suited to a technology in which capability is built through rapid, concentrated investment in model development. Japan's preference for soft-law AI governance (Habuka 2023) may be defensible on regulatory grounds, but the results suggest that the absence of an equally assertive promotion strategy until 2025 left its conversion links weak. The creation of an AI Strategy Headquarters and the AI Promotion Act in 2025 can be read as a recognition of this gap (Oxford Insights 2026); whether they close it is an empirical question for the coming years.

5.2 Implications for the debate on AI sovereignty

Discussions of AI sovereignty tend to focus on control over inputs: domestic chips, sovereign compute, national datasets and talent retention. The East Asian evidence suggests that inputs are a necessary but insufficient condition. Japan holds more compute than South Korea and more researchers and AI publications, yet it produces fewer models, fewer notable models and much less compute-weighted output. Sovereignty over inputs that are not converted does not translate into capability. Conversely, South Korea shows that a medium-sized economy can achieve frontier-level conversion in specific links when the state has the capacity to coordinate. For policy, this implies that the governance of conversion (the speed and coherence with which funding, compute access, data rules and talent policy are aligned behind development and adoption) deserves as much attention as the accumulation of inputs.

5.3 Lessons for emerging economies

The findings carry implications beyond East Asia. Emerging economies, including those in Latin America, are formulating AI strategies with input bases far smaller than Japan's or South Korea's. The South Korean case suggests that a small input base need not imply low conversion, provided the state can concentrate resources and coordinate with domestic firms. The Japanese case suggests the reverse

risk: a country can invest substantially in research and infrastructure and still convert little if policy capacity is weak. For countries that cannot hope to match the inputs of the leaders, conversion efficiency is the more realistic margin of competition. Building analytical, operational and political capacity in the agencies responsible for AI policy may therefore yield more capability per unit of public money than additional spending on inputs alone.

5.4 Limitations and further research

The analysis has the limitations of a small-N comparative design. With three focal countries and one reference, the propositions are illustrated rather than tested, and the alignment between policy capacity and conversion is a correlation that other unobserved factors, such as corporate governance, labour-market mobility or the structure of domestic demand, could also explain. The policy-capacity measure is taken from a single edition of the readiness index, recorded after most of the outputs it is compared with. The model databases are likely to under-record non-English releases, and the operationalisation layer excludes China. Future research could extend the conversion ratios to a larger panel of countries, which would allow regression-based tests; use firm-level data to examine where within national systems conversion breaks down; and track whether Japan's 2025 institutional reforms change its conversion trajectory.

6 Conclusion

This article has measured how efficiently China, Japan and South Korea convert AI inputs into AI capability. Building on the three-layer framework proposed by Tassos (2026), it computed conversion ratios for invention and industrialisation, described operationalisation, and related the results to the institutional capacity of the state. Japan emerges as the lowest-converting case in almost every link, despite larger research, R&D and compute endowments than South Korea. Its gap is concentrated in the conversion of knowledge and compute into models, and it coincides with the lowest policy-capacity score among the countries compared. South Korea achieves high conversion with a modest input base, and China combines scale with accelerating conversion.

The main implication is that national AI capability depends not only on what a country accumulates but on how its state governs the conversion of those resources. For scholarship, the article offers a transparent and replicable way of measuring conversion that can be extended to other countries and periods. For policy, it suggests that the capacity to coordinate, rather than the volume of inputs alone, is what distinguishes countries that build AI from those that mainly fund it.

Methods note on the use of generative AI

In line with the publisher's policy on artificial intelligence, the author reports that Claude (Anthropic), a large language model, was used during this research to write and debug the scripts that download and process the data, to assist with screening the literature, to prepare figures and to draft and edit parts of the manuscript under the author's direction. The author designed the study, selected the sources and

indicators, verified all data, code outputs, numerical results and references against their primary sources, and takes full responsibility for the content of the article.

Declarations

Competing interests.

The author has no competing interests to declare that are relevant to the content of this article.

Ethics approval.

Not applicable. The study uses only publicly available aggregate data.

Funding.

No funding was received for conducting this study.

Data availability.

All data, code and the download manifest are openly available in Zenodo at <https://doi.org/10.5281/zenodo.23395483>.

References

1. Abramovitz M (1986) Catching up, forging ahead, and falling behind. *Journal of Economic History* 46(2):385–406
2. Amsden AH (1989) *Asia's next giant: South Korea and late industrialization*. Oxford University Press, New York
3. Brynjolfsson E, Rock D, Syverson C (2021) The productivity J-curve: how intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics* 13(1):333–372
4. Cockburn IM, Henderson R, Stern S (2019) The impact of artificial intelligence on innovation: an exploratory analysis. In: Agrawal A, Gans J, Goldfarb A (eds) *The economics of artificial intelligence: an agenda*. University of Chicago Press, Chicago, pp 115–146
5. Cohen WM, Levinthal DA (1990) Absorptive capacity: a new perspective on learning and innovation. *Administrative Science Quarterly* 35(1):128–152
6. Ding J (2024) *Technology and the rise of great powers: how diffusion shapes economic competition*. Princeton University Press, Princeton
7. Epoch AI (2025) GPU Clusters [dataset]. <https://epoch.ai/data/gpu-clusters>. Accessed 14 October 2026

8. Epoch AI (2026a) AI Models [dataset]. https://epoch.ai/data/all_ai_models.csv. Accessed 14 October 2026
9. Epoch AI (2026b) Notable AI Models [dataset]. <https://epoch.ai/data/notable-ai-models>. Accessed 14 October 2026
10. Evans PB (1995) *Embedded autonomy: states and industrial transformation*. Princeton University Press, Princeton
11. Fagerberg J, Srholec M (2008) National innovation systems, capabilities and economic development. *Research Policy* 37(9):1417–1435
12. Freeman C (1987) *Technology policy and economic performance: lessons from Japan*. Pinter, London
13. Furman JL, Porter ME, Stern S (2002) The determinants of national innovative capacity. *Research Policy* 31(6):899–933
14. Habuka H (2023) *Japan's approach to AI regulation and its impact on the 2023 G7 presidency*. Center for Strategic and International Studies, Washington, DC
15. Hu J, Wei X, Zhou Y, Li Y, Zhou Q, Wang W (2026) A global AI competitiveness assessment and investment strategy based on entropy weighting and multi-objective optimization. In: *Proceedings of the 2026 3rd Guangdong-Hong Kong-Macao Greater Bay Area International Conference on Digital Economy and Artificial Intelligence (DEAI 2026)*. ACM, New York, pp 268–273. <https://doi.org/10.1145/3837873.3837916>
16. Johnson C (1982) *MITI and the Japanese miracle: the growth of industrial policy, 1925–1975*. Stanford University Press, Stanford
17. Kim L (1997) *Imitation to innovation: the dynamics of Korea's technological learning*. Harvard Business School Press, Boston
18. Lee K (2013) *Schumpeterian analysis of economic catch-up: knowledge, path-creation, and the middle-income trap*. Cambridge University Press, Cambridge
19. Leusin ME, Jindra B, Hain DS (2025) Breaking in or breaking through? How local specialisations shape the integration of AI technologies. *Economics of Innovation and New Technology*. <https://doi.org/10.1080/10438599.2025.2558626>
20. Lundvall BÅ (ed) (1992) *National systems of innovation: towards a theory of innovation and interactive learning*. Pinter, London
21. Microsoft AI Economy Institute (2026) *Global AI adoption in 2025: AI diffusion report H2 2025*. Microsoft, Redmond. <https://www.microsoft.com/en-us/research/wp-content/uploads/2026/01/Microsoft-AI-Diffusion-Report-2025-H2.pdf>
22. Nelson RR (ed) (1993) *National innovation systems: a comparative analysis*. Oxford University Press, New York
23. Oxford Insights (2026) *Government AI Readiness Index 2025 (January 2026 version)*. Oxford Insights, London. <https://oxfordinsights.com/wp-content/uploads/2026/01/2025-Government-AI->

24. Painter M, Pierre J (eds) (2005) Challenges to state policy capacity: global trends and comparative perspectives. Palgrave Macmillan, Basingstoke
25. Priem J, Piwowar H, Orr R (2022) OpenAlex: a fully-open index of scholarly works, authors, venues, institutions, and concepts. arXiv:2205.01833
26. Roberts H, Cowls J, Morley J, Taddeo M, Wang V, Floridi L (2021) The Chinese approach to artificial intelligence: an analysis of policy, ethics, and regulation. *AI & Society* 36(1):59–77
27. Sevilla J, Heim L, Ho A, Besiroglu T, Hobbhahn M, Villalobos P (2022) Compute trends across three eras of machine learning. In: 2022 International Joint Conference on Neural Networks (IJCNN). IEEE, Padua, pp 1–8. <https://doi.org/10.1109/IJCNN55064.2022.9891914>
28. Stanford HAI (2026) Artificial Intelligence Index Report 2026. Stanford Institute for Human-Centered Artificial Intelligence, Stanford. <https://hai.stanford.edu/ai-index/2026-ai-index-report>
29. State Council of the People's Republic of China (2017) Notice on the issuance of the New Generation Artificial Intelligence Development Plan (Guo Fa [2017] No. 35). State Council, Beijing
30. Sun Y-C, Cosgun O, Sharman R, Mulgund P, Delen D (2024) A stochastic production frontier model for evaluating the performance efficiency of artificial intelligence investment worldwide. *Decision Analytics Journal* 12:100504. <https://doi.org/10.1016/j.dajour.2024.100504>
31. Tassos A (2026) Converting artificial intelligence into national capability: a three-layer framework for comparative strategy. *Comparative Strategy*. <https://doi.org/10.1080/01495933.2026.2702077>
32. Wade R (1990) Governing the market: economic theory and the role of government in East Asian industrialization. Princeton University Press, Princeton
33. World Bank (2026) World Development Indicators. World Bank, Washington, DC. <https://databank.worldbank.org/source/world-development-indicators>. Accessed 14 October 2026
34. Wu X, Ramesh M, Howlett M (2015) Policy capacity: a conceptual framework for understanding policy competences and capabilities. *Policy and Society* 34(3–4):165–171
35. Yang Z, Li Y (2026) Comparing the sustainable role of higher education in national artificial intelligence strategies through the lens of policy documents in China, Japan, and South Korea (2017–2025). *Sustainability* 18(8):3831. <https://doi.org/10.3390/su18083831>

Figures

Conversion ratio = output of a layer / input feeding that layer, relative to the US frontier




Figure 1

Analytical framework: conversion of AI inputs through three layers, conditioned by state policy capacity.
Source: author, building on Tassos (2026)




Figure 2

Conversion ratios relative to the United States, 2017–2025 (log scale). Source: author's calculations from Epoch AI, OpenAlex, World Bank and Stanford HAI data




Figure 3

AI models per 10,000 AI publications of the preceding year, by window. Source: author's calculations from Epoch AI and OpenAlex data




Figure 4

Notable-model conversion and two pillars of the Government AI Readiness Index 2025. Source: author's calculations from Epoch AI, OpenAlex and Oxford Insights (2026)