

# From Sovereignty to Coordination: Rethinking National AI Strategy

The Economy Research Editorial<sup>1,2</sup>

<sup>1</sup>The Economy Research, 71 Lower Baggot Street, Dublin 2, Co. Dublin, D02 P593, Ireland

<sup>2</sup>Swiss Institute of Artificial Intelligence, Chaltenbodenstrasse 26, 8834 Schindellegi, Schwyz, Switzerland

## Abstract

This article argues that the increasingly prevalent interpretation of the quest for AI development as a pursuit for national "technological sovereignty" is analytically inadequate and increasingly untrue. In contrast to the competitive notion of AI development as a process of seeking to replicate vertically-integrated "sovereign stacks", it instead suggests that the creation of AI is increasingly constituted by an inherently distributed global production process across critical technological strata such as compute infrastructure, foundational models, data sources, application stacks, and governance mechanisms. The geographies of these crucial strata of the AI technology stack are becoming geographically fragmented by national borders. Under these conditions, efforts by individual nation-states to independently build full-stack AI technology would lead to significant redundant effort and systemic inefficiency, resulting in marginal returns. Building upon new empirical observations across 2023-2026, the article argues that states derive greater value through specialization in certain AI functions which correspond most effectively with existing industrial structures, human endowments, and institutional strengths and through integration into interoperable trans-national coordination mechanisms. The article therefore, frames the future of AI strategy not as competition but rather as a coordination problem of allocating the relevant cognitive functions across sovereign boundaries; sustainable long-term advantage will be predicated less on national dominance and more on systemic coordination.

# 1 Introduction - Rethinking the AI Race: From Sovereignty to Co-ordination

The rhetoric of an AI race is politically convenient but increasingly analytically misleading. It suggests that the strategic task is to build nationally sovereign AI capabilities insulated from foreign dependence. Contemporary AI, however, is neither a self-contained good to be developed and retained entirely within national borders, nor simply a productivity augmentation that can be bolted on to existing institutions ad hoc. It is a structural force that remakes the processes of learning, the training and assessment of workers by firms, the assessment of competence by institutions and over time, the production and reproduction of inequality.

Between 2023 and 2026, the diffusion of generative AI happened at a pace that renders these structural mechanisms policy-relevant today, rather than in some envisaged future. A report from the World Bank, drawing on high-frequency traffic data, estimates that 40 leading generative AI tools were receiving roughly 3 billion monthly visits in March 2024. It also demonstrates the speed with which adoption expanded away from early high-income zones, where middle-income economies incidentally accounted for a major part of the traffic, whereas low-income economies' presence remained marginal in worldwide patterns of usage.<sup>[1]</sup> This is important because diffusion is neither politically nor economically neutral. As massive segments of the world's population use generative AI to facilitate writing, programming, editing, summarizing, data tidying, etc., the technology begins permeating cognition and output at scale: humankind adapts to the machine, not vice versa.

The frontier layers of AI production are simultaneously becoming more capital-intensive and concentrated.. Publicly documented cost calculations for training frontier AI models indicate that leading projects have announced research expenditures in the tens of millions of dollars; however, recent evidence suggests that some organizations have reached comparable model training using significantly decreased computational assets than others, which could affect future cost trajectories.<sup>[2]</sup> This cost structure not only excludes a wide swath of potential builders of frontier systems; it also revises what 'sovereignty' can be meaningfully believed to mean in most states. Sovereignty, shot through with local replication of the full stack, turns out to be in practice a high-cost proxy for the actual strategic demands of most countries: credible autonomous control over operational deployment decisions; resilience to coercion or supply shocks; and leverage over global supply chains.

The trend toward concentration is highlighted by shifts in where frontier AI development occurs. The Stanford HAI AI Index reports that nearly 90 percent of notable AI models from 2024 originated in industry, a steep increase from 2023. Academia, meanwhile, remains the primary source for highly cited research.<sup>[3]</sup> This signals a structural change: the creation of frontier models has become an industrial process, reliant upon economies of scale, proprietary infrastructure, and the financial and organizational resources to enable ongoing training. The policy takeaway is not that governments should exit the AI field, but that calculated engagement will need to evolve. Governments can still shape AI deployment, oversight, assessment, and workforce integration, even if replicating every frontier domestically is neither feasible nor desirable.

Those interactions with education and labor have dimensions more complex than can be contained by an emphasis simply on industrial policy. Teachers and schools are increasingly confronted with AI as an everyday presence that somehow compromises the authority and perceived legitimacy of assessment and learning

outcomes; according to OECD analysis of TALIS 2024 data, one in three teachers used AI in some capacity on average across OECD systems, and significant proportions feel they lack the knowledge to teach using it.<sup>[4]</sup> OECD's comparative review of early governance of generative AI in education notes that, in early 2024, none of the 18 OECD countries and jurisdictions with comparable data had implemented a specific regulation for the use of generative AI in education, with most guidance merely advisory and distancing accountability from education providers and students.<sup>[5]</sup> Workplace-facing diffusion is similarly extensive; Microsoft and LinkedIn state that, in 2024, 75% of knowledge workers had used AI at work, with usage having almost doubled in just six months.<sup>[6]</sup> This suggests that AI is already functioning as a layer of cognitive infrastructure, irrespective of how prepared formal institutions are.

A policy response that sees AI only as a standalone productivity tool, without changing incentives, will overlook important trade-offs. When organizations and work practices do not adapt, drops in productivity can hide declines in lasting skills and trust in institutions. A 2025 CHI paper, including Microsoft Research, finds that generative AI makes problem-solving, like critical thinking, much easier—a concern if this leads to overuse and unthinking reliance in everyday tasks. Studies like this remind us of a key point: AI changes how much effort thinking takes, and effort is key to learning and building skills.<sup>[7]</sup>

This report calls for a new way to think about AI strategy. Rather than treating AI as a competition between countries to build separate systems, we should see it as a coordination problem: how to share AI development across nations in global trade, by specializing according to strengths, within a system where different AI tools work together. The goal of the national strategy should be to negotiate strong positions in the global AI network that provide real benefits, not to re-create the whole system within one country. This shift in thinking helps us avoid repeating efforts, manage value fairly, keep reliable oversight, and reduce gaps between countries working with AI.

## 2 Core Critique

The Brookings argument in 'What national AI plans get wrong and how to fix them' offers a necessary intervention, given that it disputes this mistaken, fetishized and oversimplified conception of 'AI leadership'.<sup>[8]</sup> It claims AI must not be thought of in itself as a 'sector'; value creation will need to take place through the embedding of AI into existing industries; and warns against equating strategy with compute procurement, urging them to instead build 'cognitive infrastructure'- comprising information, institutions, expertise, and local contextual knowledge.<sup>[9]</sup> This proposition transpires to be correct in most of its elements.

Yet the Brookings framework—even when it recognizes the danger of duplication—still appears to be primarily organizing the AI universe by national plans, treating each country as a semi-autonomous optimizer. The underlying assumption appears to be that the main challenge resides in misallocation within countries—plans that remain overly concentrated on generic stacks rather than sectoral deployment (which are over-concentrated on generic stacks) rather than sectoral deployment, with the solution to push national AI strategy closer to national endowments. Brookings does at least acknowledge the danger that constructing a fully sovereign stack will cause significant duplication and incompatible standards and it explicitly contemplates coordination among

nations on evaluation, data, deployment frames, standards, and governance. That recognition is valuable. The point developed here is that this joint effort should not be treated as a secondary concern to the national plan approach. It is the core strategic truth.<sup>[10]</sup>

The reason is simple: AI is not something we simply "adopt." AI systems cumulative and interdependent and their layers are already deeply interconnected. As a result, nations will face a strong push for redundant planning: many governments funding overlapping compute facilities, national models, testing systems, and governance structures. These repeated investments aren't just wasted effort: they change incentives in global job markets and in domestic institutions. When the language of national sovereignty and technical independence becomes popular, governments feel more pressure to fund visible computer projects—national AI models and infrastructure—since these signal progress and independence. Meanwhile, essential work including updating curricula, reforming exams, retraining workers, strengthening procurement, and any real improvements to governance is neglected, even though it's key for sustained prosperity and equality.

Quantitative evidence on adoption reveals how easy it is to get the policy problem wrong. Business adoption accelerated sharply in 2023-24 but was nowhere close to ubiquitous: OECD found in EU27, 13.5% of enterprises with at least 10 employees had adopted AI in 2024—almost 60% more than the year before.<sup>[11]</sup> This is fast growth, but it also means that most firms were not using AI even as national governments announced ambitious sovereign plans. The temptation is to jump to the conclusion that it is a domestic supply problem, insufficient national computing capacity or national models. The reality is that most barriers to adoption are institutional: poor managerial capability to embed AI into day-to-day activity, unclear assignment of responsibility for error, low trust in outputs, or uncertainty about compliance. Building domestic models will not fix these barriers and could cause us to overlook them.




Figure 1: AI spread quickly across economies, but actual usage remained highly concentrated and uneven.

A second quantitative marker helps explain why a proliferation of frontier-model copying is becoming in-

## AI adoption by firms across the OECD more than doubled in two years



Source: OECD; 2026; AI use by individuals surges across the OECD as adoption by firms continues to expand




Figure 2: AI adoption is accelerating, yet most firms still remain outside the frontier of meaningful use.

creasingly perilous: scale and centralization of frontier model design.<sup>7</sup> An extensive cost model on frontier training runs informs that reported published training costs hover in the tens of millions of dollars for leading models and projects, which, if current trends persist, in single runs in the billion-dollar range may be achieved by 2027.<sup>[12]</sup> These trends suggest that the production of frontier models is becoming "lumpy" and scale sensitive. It is highly unlikely that most nations have the capacity to do the same domestically without transferring resources away from other high-yield investments. Additionally, the AI Index accounts for the fact that almost 90% of significant AI projects in 2024 were industry-led.<sup>[13]</sup> That is, not only is the threshold for ownership of the frontier costly; it is inherently ingrained in industrial systems built on datacentres and organizational assets of a nature that most states are unable to replicate quickly.

A third issue concerns human capital. Since race framing often assumes that nations can "train their way" from technological conservatism to world leadership through rapid upskilling and talent recruitment, it must acknowledge that high-end AI expertise is globally concentrated; its distribution is defined by strong network effects, in the form of world-leading universities, cutting-edge firms, and rising stars, and global migration flows.<sup>[14]</sup> Using LinkedIn data, the OECD recently reported that, nationwide in 2024, only 7 in 1,000 LinkedIn users could be considered AI engineering talent. Globally, the highest concentrations of AI engineering talent are clustered into only a handful of countries.<sup>[15]</sup> With this in mind, a sovereign-stack strategy may lead to a one-shot talent war that is individually rational but systemically suboptimal: each country solicits the same deeply scarce scientists but underinvests in related capabilities that are more reliably improving aggregate social welfare, such as deployment infrastructure, domain-specific datasets, and reputation, or prudent regulation.

The institutional consequences of this misframing are most evident in education. AI is not a neutral classroom tool; it alters the informational landscape in which instruction takes place. The OECD's review of generative AI governance in education argues that generative AI disrupts established pedagogical and evaluative procedures-

assignment, homework, assessment-and that much of the initial scrutiny centered upon issues like cheating, the tendency to plagiarize, and "learning loss".<sup>[16]</sup> According to TALIS 2024 data, "teachers most often use" AI to "produce summaries of topics" and "support lesson planning", yet "far fewer" employ AI to "evaluate classwork", while "a considerable proportion" feel "unprepared" to teach with AI tools, and "another substantial minority" are "doubting its role in instruction".<sup>[17]</sup> The Implication for policymakers, then, is that education authorities face a technology capable of generating convincing results much faster than institutions can innovate modalities for authenticating merit. Once the evaluation framework further degrades, inequality will be exacerbated because institutions offering schooling can channel resources into tools, tutoring, and networks that assuage the existing uncertainty, while students from lower social origins will increasingly lack a one-to-one validation proxy.

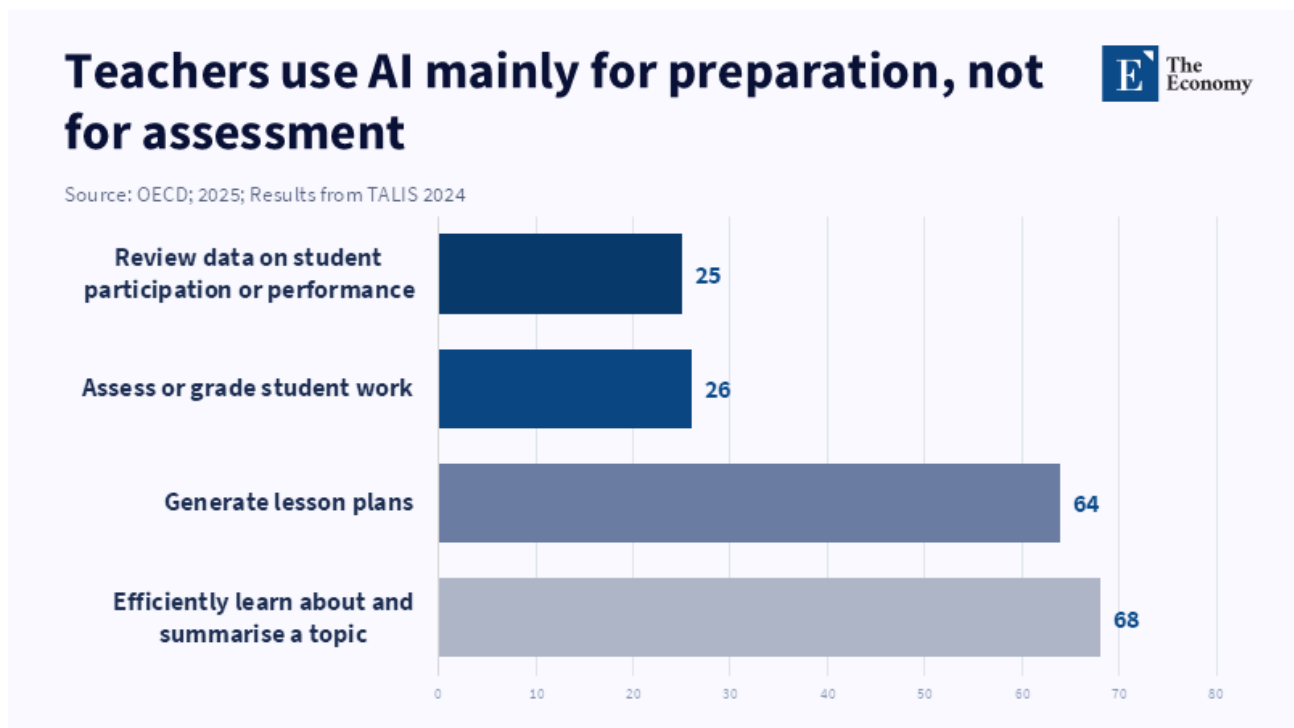


Figure 3: Education systems are integrating AI faster in preparation than in evaluation, where institutional risks are greater.

Workplace adoption deepens the same mechanism. Microsoft reports ubiquitous AI use among the world's knowledge workers.<sup>[18]</sup> However, the key issue is not simply whether AI increases productivity; it is whether it alters the knowledge work process to undermine the accumulation of skills. The CHI 2025 survey study reveals that generative AI could reduce subjective effort for complex thinking tasks, especially in routine scenarios.<sup>[19]</sup> This leads to the institutional paradox that if outputs are easier to produce, organizations can increase performance targets and shorten timeframes (thus implicitly adopting AI-generated outputs as a new norm). Under such circumstances, knowledge workers can appear to be far more productive while having less incentive and time to acquire deeper expertise. This can create a "false productivity" trap in which outputs soar and human judgment atrophies.

This is precisely why the distinction between a national optimization problem and a coordination problem does matter. Indeed, Brookings is right in the sense that nations should organize AI activities along the lines

of natural comparative advantage.<sup>[20]</sup> The crux of the matter is that the market failure is not exclusively a misallocation of resources within countries; it is also needless replication across nations in a game characterized by scale economies and network effects.<sup>[21]</sup> The creation of mismatched reporting standards, incompatible benchmarks, and parallel model development programs in different countries would result in the loss of interoperability. In a modular AI system, interoperability is not a technical afterthought; it is the mechanism through which specialization generates collective gains.

A policy-reframing is thus in order. The strategic question is not, 'How can each country develop AI sovereignty?' It is, 'How should a portfolio of mental processes be distributed across countries so as to maximize global, and not just national, capacity?' In terms of trade theory, the autonomous-multinode firm analog of the sovereign-stack motive is, under some security assumptions, inefficient and has negative welfare effects when production exhibits modularity and decreasing returns to scale. The latter is not a stance to deny the security of AI. It is acknowledged that sovereignty is more credibly achieved through governance and bargaining power, credible exit options, diversified suppliers, and interoperable standards than by replication.

### 3 The Missing Layer: AI as a Global Production System

In each national AI strategy, "the AI stack" is discussed, but only in the context of a thoroughly "nationally integrated pipeline". What is absent is a clear layers-of-production-systems concept, AI as a distributed globally configured system of production that can be integrated in different iterations of a layered stack in different jurisdictions.

Beyond its effectiveness, the OECD's framing of AI compute is revealing because it depicts the supply chain in explicitly distributed terms: it identifies chip designers, cloud compute vendors, AI developers, and users.<sup>[22]</sup> The complete chain is hardware, data centers, networking, and software ecosystems. The policy lessons to draw are straightforward. If the supply chain is globally dispersed, sovereignty cannot be secured through declaration alone; it is determined by access, standards-setting and deployment governance. A production-systems perspective explains why there is, in general, no country with a comparative advantage in all layers.

The physical compute layer is the most capital-intensive and geographically concentrated. According to OECD's analysis of the distribution of capacity in the semiconductor manufacturing industry, ten companies (> 16% of total) dominate global wafer fabrication capacity, and five economies share almost 90% of the world's capacity (China, Taiwan, Korea, Japan, and the United States).<sup>[23]</sup> Within those economies, the scope of capacity concentration is immense, or the manufacturing specialization is narrow. We can see that production is concentrated on a small handful of chip types, and is more advanced as we move into more advanced process nodes. OECD mapping of the mine-to-money semiconductor value chain provides additional evidence of the capital costs behind entry, showing that commodity memory (DRAM and NAND) is produced by IDMs like Samsung, SK Hynix, Micron, Kioxia, and YMTC that scale with continual node upgrades.<sup>[24]</sup> Our examination of the drivers of the chip landscape uncovered that Korean wafer capacity is likewise highly focused: roughly 80% of Korea's wafer capacity is concentrated between 6 nm and below 22 nm process nodes.<sup>[25]</sup> These facts reiterate a structural reality: for most countries, "sovereign compute" in the simple physical sense remains

prohibitively expensive (cost-ineffective), and the costly pursuit of sovereign compute frequently displaces other, more pragmatic courses of sovereignty: diversified purchasing, resilient alliance.

The model layer is characterized by parallel concentration dynamics. Cost modeling shows that it is so costly to train frontier models that only a small cluster of organizations can afford to do so repeatedly. According to the AI Index, organizations that led the production of all significant AI models in 2024 were also responsible for 90 percent of these models, indicating that many national strategies aiming to build a sovereign foundation model appear to focus on claiming domestic production rather than developing both the models and their supporting capabilities. Sovereignty in a modular system is often necessarily partial.<sup>[26]</sup>

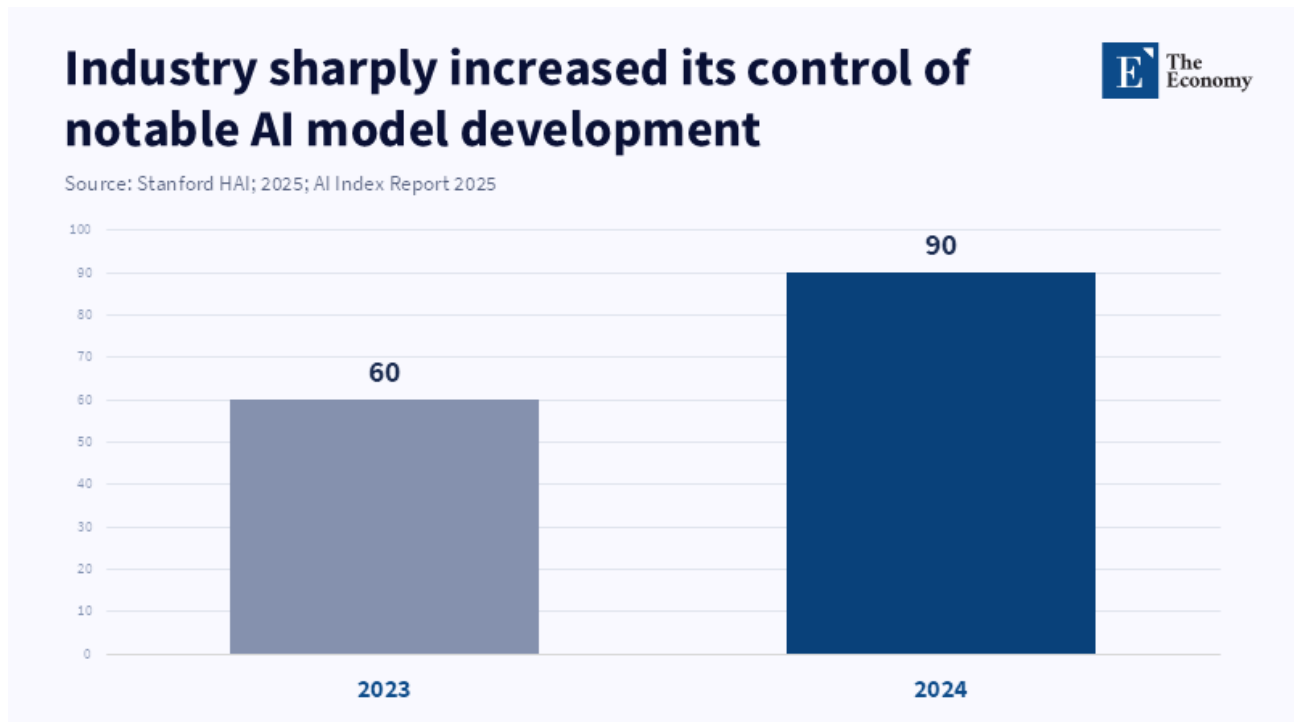


Figure 4: The frontier is no longer just advancing; it is becoming more structurally concentrated in industry.

The data layer is a largely misunderstood concept. Policymakers frequently view data as a national resource that can be collected and consolidated. In reality, however, much of the most useful data for applied AI is local, domain-specific, and embedded within institutions, such as health systems, manufacturing, taxation, settlements, and administration. Its value rests on governance, quality, ownership, consent, security, and interoperability. A country can have a wealth of data but very little institutional capacity to leverage it safely and responsibly. That is as much a human and governance challenge as it is technical.

The application layer is where comparative advantage is most elastic and most depends on human capital formation. Here, the core bottleneck of AI isn't computational but "attention": domain knowledge, workflow redesign, and organizational ability to incorporate AI in production without loss of reliability. Empirical results from World Bank analyses suggest large productivity gains for selected cognitive tasks, often greater for novices and lower-skilled workers.<sup>[27]</sup> However, they will not necessarily lead to broad-based welfare gains, and if complementary assets are unevenly distributed, they will be captured by capital owners or concentrated firms; if institutions treat AI output as a substitute for or a replacement of learning, they will undermine skill formation.

This is why AI must be treated as a structural power that determines inequality. The OECD evidence shows




Figure 5: Performance convergence weakens the case for broad national replication of frontier model development.

that even among firms in the tech revolution that adoption is uneven across firms of different sizes, with large firms far ahead of small firms that adoption is uneven by age, education, and income, with students among the highest adopters and that the global productivity split can again be explained by 'progress within digital infrastructure or digitalization levels that, although less concentrated among higher income countries, remained a bottleneck for overall productivity progress in lowermiddle- and low-income economies', along with, again, the pattern of adoption reflecting educational structures.<sup>[28]</sup> In a time when we have an AI cognitive layer, traditional digital divides are being transformed into gaps in cognitive amplification. These cases, Korea and India, demonstrate how adopting a production-systems approach toward strategic understanding of national endowments occurs.

South Korea is often described as an AI adopter trying to catch up to advanced frontiers. Production-systems literature, however, finds that the special strategic function of Korea is in the physical and industrial levels of AI: semiconductors, industrial process engineering, and high-reliability production integration. A recent IMF report on Korea's trade profile points out that Korea's exports constituted about 40% of GDP in 2023, and that semiconductors and cars are the most significant Korean export items, respectively accounting for 20% and 12% of the nation's total exports.<sup>[29]</sup> The OECD chip landscape results on Korea's concentration in process node further confirm that Korea is structurally placed in a world critical part of the physical foundation of AI.<sup>[30]</sup>

The strategic advice is not that Korea should not use any models. Korea should avoid the temptation to use frontier model competition as the dominant axis of national AI performance. Korea is not 'underperforming' in AI if it does not train frontier base models at the level of the United States or China; instead, it may be overextending itself in competitive effort. Korea's highest-return opportunities are likely to be: dominant leadership in AI-relevant memory and packaging; resilient supply-chain partnership; and embedding AI into

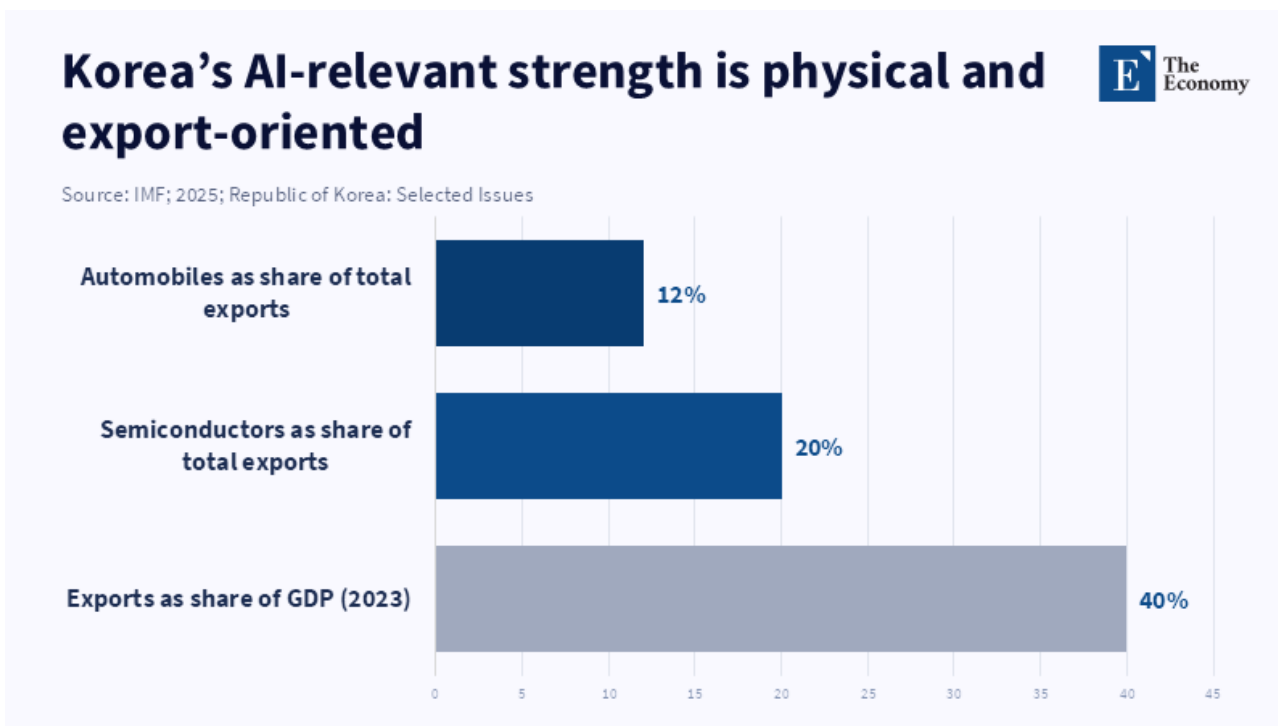


Figure 6: Korea's measurable comparative strength lies in industrial hardware and export depth, not in reproducing the full model frontier.

world-leading manufacturing, robotics, and smart-factory systems, where Korea has strong complementary capacity. Policy-wise, this involves converging education, industrial strategy, and governance. Where AI is embedded in manufacturing, it is those durable cognitive skills that matter (systems thinking, statistical process control, verification means and methods, root-cause diagnosis in conditions of uncertainty). AI can help in monitoring and simulation, but high-reliability manufacturing quality, in systems that are reliably predictable, then depends on human diagnostics. The second-order downside risk is the erosion of expertise: if AI becomes merely a shortcut to troubleshooting rather than a way to increase learning speed, then the manufacturing industrial system becomes progressively less resilient.

India is differently endowed. Its comparative advantage lies less in capital-intensive fabrication than in scaling cognitive services. According to the World Bank, evidence on the diffusion of generative AI, India ranks among the top economies in ChatGPT traffic, and middle-income economies collectively constitute a large share of the total, meaning the AI adoption adds to the service-sector competitiveness.<sup>[31]</sup> Meanwhile, the labor indicators at the index of AI, India scores high relative to AI skill penetration rates on the LinkedIn-based measures, which indicates the potentially larger base for application-layer scaling. The official Indian Government Handbooks provide evidence of robust services export dynamics, with software services contributing a considerable share of services exports, which recorded a rapid growth in FY23-FY25.<sup>[32]</sup>

From a coordination perspective, India's best opportunity may not be to lead in a capital-intensive compute race. It is to scale application development, workflow automation, service delivery systems, and agentic tools using its abundant engineers and service export infrastructure. The danger is that an application-layer<sup>[33]</sup> specialization might turn extractive if its cognitive foundations are fragile: if AI tools make certain intellectual tasks appear less effortful, as the CHI 2025 project predicts, then a service-export economy might carry a quality

risk: polished outputs that are less well grounded, cost because avoidance rather than norm-creating verification, and competition by effort-reduction rather than effort-advance. This can, over time, depress wages and increase inequality within a workforce by favoring those who can carry and check AI systems, and disfavoring those whose work becomes more replaceable.

This is where education policy takes center stage. We often expect AI to democratize knowledge by bringing high-quality information to the masses: this may be conceptually valid, but in practice, its successful democratization relies on specific supporting skills: a capacity to determine veracity, synthesize context, and apply domain judgment. OECD analysis regarding governance of AI in education focuses on the dilemma between harnessing AI for learning purposes and avoiding possible risks such as bias, privacy, inaccuracy, and skills attrition. TALIS indicates significant teacher skills gaps and diverse attitudes towards AI-guidance.<sup>[34]</sup> In an environment of this kind, openness to ad hoc institutional adoption of AI risks leading to a low-trust learning environment, with decayed incentives to learn and less credible credentials. The focus ought to shift from banning AI (it is already too late for that) to reconceptualizing curricula and assessments so that students are required to provide reasoning-based arguments rather than results generated by AI.

These are the cases that help us see where sovereignty makes the most difference. Korea should have sovereign control of some deployments (not all) related to defense or critical infrastructure, but its comparative advantage in economy-wide AI is to achieve sovereignty at the physical-layer level and embed AI in industrial systems possessing strong verification and governance. India should have sovereign control of some domains related to public service and data governance, but its comparative advantage in economy-wide AI is to gain prosperity through interoperability with models and infrastructure worldwide, maintained by domestic capacity to compete and to evaluate, audit, and change providers. Both require strong bargaining power against others; thus, the main risk is not interdependence per se, but unmanaged interdependence with weak institutional capacity.

## **4 The future of AI is not sovereign - it is modular and interdependent**

But if AI is characterized by modularity and global distribution, sovereignty cannot be assumed as the natural organizing principle of strategy. However, it would be equally unwise to respond by naively celebrating interdependence, as though markets alone could guarantee efficient and fair outcomes. The strategic challenge is to move from unorganized interdependence to systematized interdependence: a coordination regime that locates the benefits of specialization and open markets, respecting the preservation and integrity of national institutions and disciplines, and assuring impartial opportunities of human capital formation.

The economic argument opposing full-stack replication revisits a well-established critique. In contexts where production is modular and scale economies are present, pursuing complete self-sufficiency frequently leads to inefficiencies and a decline in overall welfare. Artificial intelligence demonstrates significant scale economies, particularly in frontier models and advanced fabrication processes, while simultaneously fostering network effects

via standards, data ecosystems, and platform use. Consequently, multiple countries attempting to replicate the entire stack encounter diminishing marginal returns, which contribute to a broader general-equilibrium failure: redundant investments and fragmented standards undermine the complementarities essential for specialization and interdependence.

The problem of coordination is visible in these data now. OECD collects data on AI adoption among individuals and firms, which shows that AI adoption is spreading faster in OECD countries: 20.2% of firms used AI in 2025, compared to 8.7% in 2023; the adoption is very unequal among firm sizes, with large firms being considerably more likely to have adopted it than small firms.<sup>[35]</sup> And this has meant that if the governance of AI adoption does not ensure the access of smaller firms and lagging regions to the complements and standards that enable them to participate in AI diffusion, it can magnify scale economies and market power. Therefore, the question of policy becomes "not just how to deploy AI, but how to deploy AI such that it does not kill off competition, institutional trust, and skill formation."

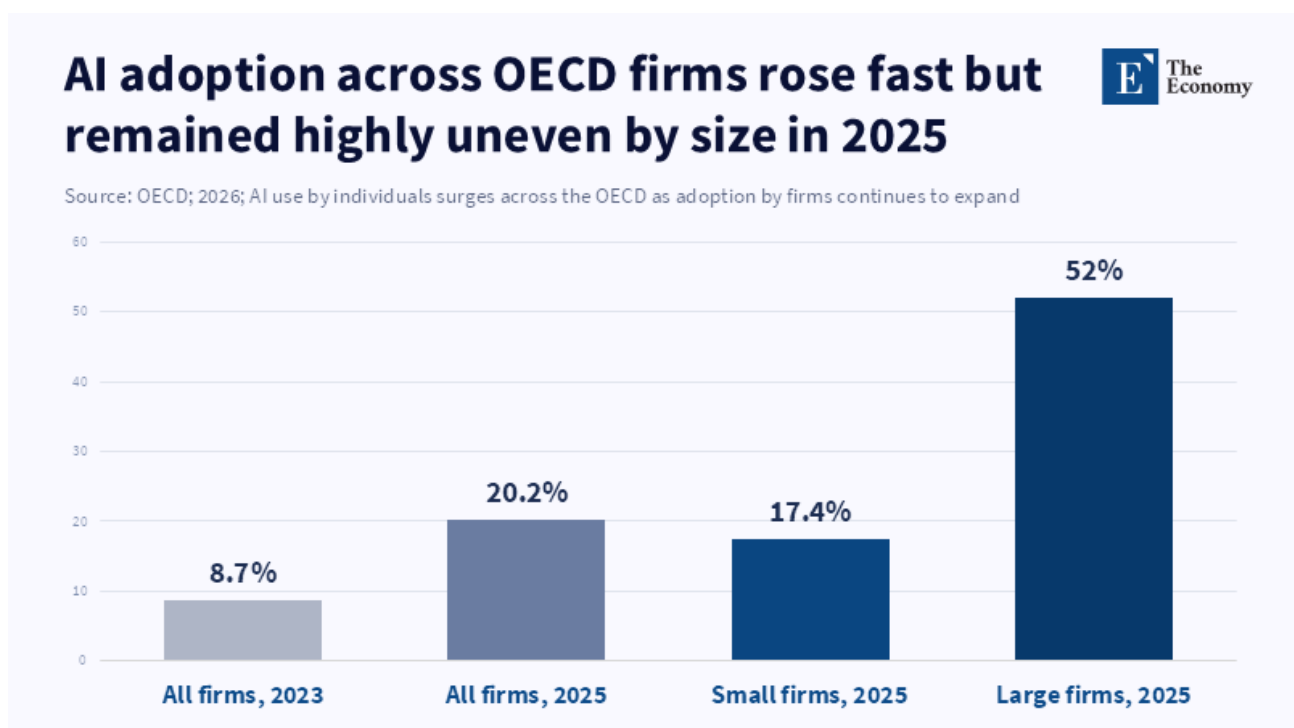


Figure 7: The central problem is not diffusion alone, but uneven diffusion across firm size and institutional capability.

The report's central thesis translates into a coordination strategy built on three pillars:

The first refers to an AI trade architecture: an explicit map of where value is created across layers, in which countries are structurally situated to supply those layers. The OECD semiconductor analysis illustrates that wafer capacity is both distilled and differentiated, while the OECD value chain analysis shows how state-of-the-art memory and scale-intensive segments are megalithic in a handful of firms. The AI index recommends that the industry control everything that constitutes models of the technological frontier. Meanwhile, India is structurally well placed to scale application-layer services. The capacity of governance is also a comparative advantage. The European Union AI Act (effective Aug 2024, implemented through 2026 and beyond) demonstrates how governance can shape global compliance regimes; it also lays down AI literacy obligations, extending

AI governance into education and human capital policy. The emerging regime is one where countries do not build every single layer but position themselves tactically in a global architecture to capture value.

The second pillar is specialization agreements. These are low-commitment, harmonized multilateral arrangements that would mitigate the incentive for unnecessary duplication while still retaining the benefits of de facto standardization. In the domain of trade, this is done not just with tariffs but with predictable rules for trade, investment, and settlements. AI requires such a bridge, but perhaps not solely with one big international framework. Consider, for example, modular generic agreements that will culminate in access to compute, regulatory, and security standards, mutual recognition of validation outcomes, joint safety trial procedures, mechanisms for secure collaboration on domain-specific datasets, etc. This would, in effect, turn the natural weakness of interdependence into a deliberate trade-off strategy.

Interoperability standards: what enables specialization. In the absence of shared technical and governance standards, the division of labor will be fragile. The OECD notes that shared terminology and definitions are the bedrock of interoperable AI policy between jurisdictions. More materially, interoperability should encompass procedures for model evaluation and benchmarking, disclosure and documentation, incident reporting, and red-teaming. The G7 Hiroshima Process International Code of Conduct for organizations creating advanced AI systems is one pragmatic effort to offer voluntary standards guidance for advanced AI development across jurisdictions.<sup>[36]</sup> Interoperability doesn't mean identical regulation; it means compatible regulation. The strategic tradeoff is that compatible standards limit redundant compliance, allowing cross-border deployment and observation. The strongest counterarguments should be confronted directly because they are as close as we get to actual policy decisions.




Figure 8: Specialization is already built into AI's material foundations; interoperability is what turns that fact into strategic advantage.

Another counterargument states that AI will truly improve productivity and expand access to knowledge,

therefore dwarfing the validity of questions of sovereignty and inequality. The first part of that statement can be backed up by evidence on productivity increases in particular tasks: experimental evidence, peer-reviewed by the World Bank, can lead to large productivity gains in specific environments, with often outsized gains being granted to beginners. However, increases in productivity in carefully controlled activities do not always equal inclusive gains to long-term welfare, as a matter of who controls the complements- access to compute, workflows, and governance- and how institutions preserve incentives to learn. If AI diminishes the so-called "effort" of engaging in high-skilled activity, and increases providing "dependence" in low-skill activities, then a society may produce output faster while accumulating skills less quickly, and democratizing output while not democratizing know-how.

A second counterargument is that the AI is simply the latest in a sequence of cognitive technologies is no different from those that occurred when new systems of cognitive tools emerged: calculators replaced manual computation, word processors replaced handwritten manuscripts, search engines replaced interlibrary loans; societies will accommodate as they have before. This analogy lacks precision, however. Generative AI not only replaces human activity in the corners of cognition (such as writing formulas or hunting for ancestors) that were replaced by earlier innovations, but outputs speech and code that can stand in as artifacts of competence in education and many workplaces. OECD analysis in particular highlights how generative AI overturns education systems' approaches to assessment in a way that previous technological gains did not. Further, generative systems' error mode is of a different character: skillful hallucinations transfer the effort of production into the domain of checking and corroborating. This transition isn't inherently unjustified-it could be the emergence of a new and efficient equilibrium, with the most important skills evolved into verification. But it could also have the consequence of stratification, with characteristics strongly associated with expertise, as entered courses of study and areas of work, identified in advance, into the secondary role of auditors and guides, leaving everyone else passively taking in the outputs of AI.

A third counterexample must be considered when approaching this in a policy-oriented report: specialization may entail becoming locked into subordinate venues and repeating patterns of unequal exchange, in which some users capture rents generated by the very infrastructure and the power to set standards, while others find themselves enmeshed in dependency. This is well taken. Comparative advantage need not be destiny, and it may turn into a well-entrenched ghetto too, if one takes it as static. Thus, the coordination approach advanced in this report calls for two safeguards.

First, specialization must be combined with explicit upgrade routes; agreements should reduce the costs of future lateral moves into neighboring capabilities, not lock in initial endowments; this is broadly consistent with Brookings' focus on adjacent diversification, but applied at the system level. Second, application layer countries must develop governance capabilities-evaluation, accounting, data stewardship-because these are points of leverage in turning dependence into a manageable state. A country that shares the global foundation but can evaluate models, set deployment restrictions, and credibly switch providers can maintain autonomous control over the local foundation.

The change from sovereignty-as-replication to coordination-as-strategy fundamentally restructures the definition of success for policy makers. Instead of defining "national base models" and "national compute" as

universal goals, countries need to be clear about which layer they can and will contribute to competently, and which layer they will consume through interoperable exchange. For Korea, this translates to identifying semiconductor and industrial AI integration as a core element of national strategy, and developing selectively exercised sovereignty through governance, assessment and procurement control, instead of competing in frontier model development. For India, it entails broadly addressing AI literacy and application layer scaling, while taking evaluation capability, privacy, and labor governance extremely seriously, and investing heavily in educational reform so that AI-enabled services leverage credible cognition instead of weak output production.

For education administrators, generative AI governance will be seen as an institution design problem rather than an isolated integrity campaign. It should be noted that OECD data indicate that binding national regulation is not commonplace, and the responsibility is increasingly shifting to individual institutions and teachers. Similarly, OECD TALIS data shows considerable gaps in training and mixed attitudes among teachers towards generative AI. Institutions need to internally define standards of acceptable use, redesign assessment methods to be sensitive to the process of reasoning, and train staff not only in using the technology but also in the verification norms and pedagogical adaptations that preserve lasting skill development.

For educators, the primary goal is to safeguard learning incentives under the conditions of pervasive cognitive assistance. Education must focus on developing skills that remain scarce when AI is ubiquitous- framing questions, judgment in a particular domain, statistical reasoning, and critical verification of AI outputs. The form of assessment must change to emphasize a process which makes reasoning legible- oral defenses, problem-solving in the classroom, staged projects with documented steps along the way, because take-home outputs have become an ambiguous indicator of capability. This is not a stance against the technology but an institutionalist stance aimed at sustained resilience, acknowledging that the increased ease with which outputs can be produced alters the incentives for learning.

The cost of failure in coordination is not only the duplication of infrastructure investments but also the corrosion of institutions such as the credibility of credentials, the skill development of individuals, and the exacerbation of disparities in terms of unequal access to the cognitive complementary technologies. A successful coordination system is capable of fostering a stable equilibrium through leveraging comparative advantage, establishing an interoperable system that facilitates exchanges between countries, and building the educational and governance infrastructure needed to transform AI from a transient efficiency tool into an engine for the development of human capital.

## 5 Conclusion - Toward Systemic Alignment

The pursuit of the "AI race" can at best distract from the actual key priority. AI is not a monolithic system that can be replicated nationally, but a globally diffused cognitive production process built from various modular layers connected across borders. Between 2023 and 2026, rapid incorporation into workplaces and education, ballooning costs at the frontier, and intense concentration in model development have made full-stack sovereignty increasingly untenable. Furthermore, second-order impacts on incentives for learning, the reliability of credentials, and institutional decision-making can degrade human capital, even if measured output grows.

An approach focused on coordination rather than replication may build a stable equilibrium. This is because it encourages countries to use AI in areas where they possess marginal comparative advantages and collaborate within a system that provides the means of interoperable exchange. The choice for Korea and India to concentrate on specific parts of the AI stack and pursue a model of conditional sovereignty via governance, evaluation and control over data and services instead of replicating all of the technology is key. Policymakers are tasked to forge mediated interdependence that looks more like trade arrangements with interoperable standards and the guarantee of upgrade pathways.

This is a pragmatic policy recommendation rooted in evidence. Governments should reassign resources away from performative duplication, towards interoperability, evaluation and designing the education system to support enduring cognitive capacity. The failure to do so could result in short-term productivity booms masked by a long-term weakening of institutional capacity and increasing inequality. The opposite could usher in an era of joint human capital development.

## References

- [1] Liu, Y., Wang, H. and Qiang, C.Z. (2024) ‘Who on earth is using generative AI?’, *World Bank Blogs*, 11 September.
- [2] Cottier, B., Rahman, R., Fattorini, L., Maslej, N. and Owen, D. (2024) ‘The rising costs of training frontier AI models’, *arXiv*. doi:10.48550/arXiv.2405.21015; Zhuang, B., Qiao, J., Liu, M., Yu, M., Hong, P., Li, R., Song, X., Xu, X., Chen, X., Ma, Y. and Gao, Y. (2025) ‘Beyond benchmarks: The economics of AI inference’, *arXiv*. doi:10.48550/arXiv.2510.26136.
- [3] Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J.C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlasci, L., Lotufo, J.B., Rome, A., Shi, A. and Oak, S. (2025) *Artificial Intelligence Index Report 2025*. Stanford, CA: Stanford Institute for Human-Centered Artificial Intelligence. doi:10.48550/arXiv.2504.07139.
- [4] OECD (2025) *Results from TALIS 2024: The State of Teaching*. Paris: OECD Publishing. doi:10.1787/90df6235-en.
- [5] OECD (2023) *OECD Digital Education Outlook 2023: Towards an Effective Digital Education Ecosystem*. Paris: OECD Publishing. doi:10.1787/c74f03de-en; Vidal, Q., Vincent-Lancrin, S. and Yun, H. (2023) ‘Emerging governance of generative AI in education’, in *OECD Digital Education Outlook 2023: Towards an Effective Digital Education Ecosystem*. Paris: OECD Publishing.
- [6] Microsoft and LinkedIn (2024) *2024 Work Trend Index Annual Report: AI at Work Is Here. Now Comes the Hard Part*. Redmond, WA: Microsoft.
- [7] Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R. and Wilson, N. (2025) ‘The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence

- effects from a survey of knowledge workers’, *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, Article 1121, pp. 1–22. doi:10.1145/3706598.3713778; Gerlich, M. (2026) ‘AI and the Rise of Societal Bifurcation: Cognitive Dependency, Inequality and Democratic Pressure’, *Societies*, 16(3), 82. doi:10.3390/soc16030082.
- [8] Kerry, C.F. and Mishra, S. (2026) ‘What national AI plans get wrong and how to fix them’, *Brookings*, 9 March.
- [9] OECD (n.d.) ‘AI compute’, *OECD Topic Pages*.
- [10] Tanner, B., Kerry, C.F., Wyckoff, A.W., Kyosovska, N., Renda, A. and Tabassi, E. (2026) ‘Is AI sovereignty possible? Balancing autonomy and interdependence’, *Brookings Institution*, 17 February; Ustun, A., Tournesac, A., Bennici, L., Schaubroeck, R., De Niese, J., Takkar, K., Krawina, M. and Drahmoune, N. (2026) ‘Sovereign AI: Building ecosystems for strategic resilience and impact’, *McKinsey*, 3 March.
- [11] OECD (2026) ‘AI use by individuals surges across the OECD as adoption by firms continues to expand’, *OECD Announcements*, 28 January; Kergroach, S. and Héritier, J. (2025) ‘Emerging divides in the transition to artificial intelligence’, *OECD Regional Development Papers*, No. 147. Paris: OECD Publishing. doi:10.1787/7376c776-en.
- [12] Cottier, B., Rahman, R., Fattorini, L., Maslej, N. and Owen, D. (2024) ‘The rising costs of training frontier AI models’, *arXiv*. doi:10.48550/arXiv.2405.21015.
- [13] Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J.C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlasci, L., Lotufo, J.B., Rome, A., Shi, A. and Oak, S. (2025) *Artificial Intelligence Index Report 2025*. Stanford, CA: Stanford Institute for Human-Centered Artificial Intelligence. doi:10.48550/arXiv.2504.07139.
- [14] Hood, R. (2025) ‘The AI workforce: What LinkedIn data reveals about “AI talent” trends in OECD.AI’s live data’, *OECD.AI*, 22 May.
- [15] OECD (2023) *OECD Digital Education Outlook 2023: Towards an Effective Digital Education Ecosystem*. Paris: OECD Publishing. doi:10.1787/c74f03de-en.
- [16] OECD (2025) *Results from TALIS 2024: The State of Teaching*. Paris: OECD Publishing. doi:10.1787/90df6235-en.
- [17] Microsoft and LinkedIn (2024) *2024 Work Trend Index Annual Report: AI at Work Is Here. Now Comes the Hard Part*. Redmond, WA: Microsoft.
- [18] Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R. and Wilson, N. (2025) ‘The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers’, *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, Article 1121, pp. 1–22. doi:10.1145/3706598.3713778.

- [19] Kerry, C.F. and Mishra, S. (2026) ‘What national AI plans get wrong and how to fix them’, *Brookings*, 9 March.
- [20] Tanner, B., Kerry, C.F., Wyckoff, A.W., Kyosovska, N., Renda, A. and Tabassi, E. (2026) ‘Is AI sovereignty possible? Balancing autonomy and interdependence’, *Brookings Institution*, 17 February.
- [21] OECD (n.d.) ‘AI compute’, *OECD Topic Pages*.
- [22] OECD (2025) ‘The chip landscape: Geographical distribution of wafer fabrication capacity’, *OECD Science, Technology and Industry Policy Papers*, No. 188. Paris: OECD Publishing. doi:10.1787/02dbd028-en.
- [23] OECD (2025) ‘Mapping the semiconductor value chain: Working towards identifying dependencies and vulnerabilities’, *OECD Science, Technology and Industry Policy Papers*, No. 182. Paris: OECD Publishing. doi:10.1787/4154cdbf-en.
- [24] OECD (2025) ‘The chip landscape: Geographical distribution of wafer fabrication capacity’, *OECD Science, Technology and Industry Policy Papers*, No. 188. Paris: OECD Publishing. doi:10.1787/02dbd028-en.
- [25] Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J.C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlaschi, L., Lotufo, J.B., Rome, A., Shi, A. and Oak, S. (2025) *Artificial Intelligence Index Report 2025*. Stanford, CA: Stanford Institute for Human-Centered Artificial Intelligence. doi:10.48550/arXiv.2504.07139.
- [26] World Bank (2025) *The Economic Impacts of Generative AI: Productivity, Jobs and Inclusion*. Washington, DC: World Bank.
- [27] Chaar, T., Filippucci, F., Jona-Lasinio, C. and Nicoletti, G. (2025) ‘AI and the global productivity divide: Fuel for the fast or a lift for the laggards?’, *OECD Artificial Intelligence Papers*, No. 51. Paris: OECD Publishing. doi:10.1787/c315ea90-en.
- [28] International Monetary Fund (2025) *Republic of Korea: Selected Issues*. Washington, DC: IMF; OECD/Korea Labor Institute (2025) *Artificial Intelligence and the Labour Market in Korea*. Paris: OECD Publishing. doi:10.1787/68ab1a5a-en.
- [29] Liu, Y., Wang, H. and Qiang, C.Z. (2024) ‘Who on earth is using generative AI?’, *World Bank Blogs*, 11 September; Liu, Y., Huang, J. and Wang, H. (2025) *Who On Earth Is Using Generative AI? Global Trends and Shifts in 2025*. World Bank Policy Research Working Paper. doi:10.60572/12rk-p733.
- [30] Hood, R. (2025) ‘The AI workforce: What LinkedIn data reveals about “AI talent” trends in OECD.AI’s live data’, *OECD.AI*, 22 May; Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J.C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlaschi, L., Lotufo, J.B., Rome, A., Shi, A. and Oak, S. (2025) *Artificial Intelligence Index Report 2025*. Stanford, CA: Stanford Institute for Human-Centered Artificial Intelligence. doi:10.48550/arXiv.2504.07139.

- [31] Press Information Bureau, Government of India (2026) ‘From stability to new frontiers, India’s services exports growth more than doubled from 7.6% in the pre-pandemic period (FY16–FY20) to 14% during FY23–FY25’, *PIB Press Release*, 29 January; D., A.A., S., C. and P., G. (2025) ‘Exit barriers as entry barriers: Explaining India’s odd development path’, *Economic and Political Weekly*, 60(1), pp. 45–53. doi:10.2139/ssrn.3701234.
- [32] OECD (2023) *OECD Digital Education Outlook 2023: Towards an Effective Digital Education Ecosystem*. Paris: OECD Publishing. doi:10.1787/c74f03de-en; OECD (2025) *Results from TALIS 2024: The State of Teaching*. Paris: OECD Publishing. doi:10.1787/90df6235-en.
- [33] OECD (2026) ‘AI use by individuals surges across the OECD as adoption by firms continues to expand’, *OECD Announcements*, 28 January; Chaar, T., Filippucci, F., Jona-Lasinio, C. and Nicoletti, G. (2025) ‘AI and the global productivity divide: Fuel for the fast or a lift for the laggards?’, *OECD Artificial Intelligence Papers*, No. 51. Paris: OECD Publishing. doi:10.1787/c315ea90-en.
- [34] European Commission (2024) ‘Regulatory framework proposal on artificial intelligence’, *Digital Strategy*; OECD (2025) *Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions*. Paris: OECD Publishing. doi:10.1787/795de142-en.
- [35] OECD (2023) *OECD Digital Education Outlook 2023: Towards an Effective Digital Education Ecosystem*. Paris: OECD Publishing. doi:10.1787/c74f03de-en.
- [36] European Commission (2023) *Hiroshima Process International Code of Conduct for Advanced AI Systems*. Brussels: European Commission.