Middle Powers and the Limits of Collective Agency in Artificial Intelligence

The Economy Research Editorial*

*The Economy Research, 71 Lower Baggot Street, Dublin 2, Co. Dublin, D02 P593, Ireland

Abstract

The article critically examines the "modular sovereignty" proposal in AI, according to which middle powers such as the European Union, Japan and South Korea could build collective capacity at select levels of the AI stack, rather than accepting complete dependence on the United States or China. Using ASEAN's recent history, particularly its inability to collectively negotiate U.S. tariffs in 2025 and the roughly three-year negotiation process in the Digital Economy Framework Agreement, the paper argues that even an institutionally mature coalition of middle powers struggles to agree on issues of governance, benefit sharing and ownership when the stakes are real. The U.N. Security Council's veto system provides a parallel example of how multilateral governance can become deadlocked when major powers' interests diverge. Instead of an ambitious, jointly governed consortium, the article concludes that middle powers are best served by a portfolio of narrow bilateral and minilateral arrangements, tailored to each country's specific strengths.

1 Introduction - AI Sovereignty Between Dependence and Cooperation

Discussions around AI and geopolitics have become entrenched in a familiar pattern: a showdown between Washington and Beijing, with every other government playing the role of a spectator who is called upon to decide which camp to join. A recent analysis by the Brookings Institution rejects this view; it argues that Europe, Japan, South Korea and other middle powers should pursue what it calls "modular sovereignty," building collective capacity on a handful of strategically chosen layers of the AI stack, rather than attempting to replicate the entire American or Chinese edifice. Its diagnosis of risk is hard to dispute: reliance on foreign models, computing power and cloud platforms can leave a country's economic and political choices hostage to decisions made elsewhere, even when its formal sovereignty remains intact.[1] What the argument assumes, more than it proves, is that a coalition of unequal governments with historically cautious attitudes can be built, governed and kept united enough to matter.

This hypothesis deserves consideration based on its own data, as the record of collective action of middle powers, even in sectors far less technically demanding than semiconductor supply chains and fundamental models, is meager. The urgency of the debate is not in doubt. The United States' private investment in artificial intelligence reached nearly $286 billion in 2025, more than twenty-three times the $12.4 billion declared by China in the same period, an amount that in any case excludes much of Beijing's state-directed funding, while only five U.S. companies are expected to spend over $700 billion in 2026 on AI-related infrastructure.[2] No group of middle powers can come close to this scale by accumulating budgets that currently run into the tens of billions.

The risk posed by this situation is not abstract. Washington briefly imposed a licensing requirement for Nvidia's H20 chips bound for China in 2025, before reversing the decision, while in June 2026 a separate export control directive forced Anthropic to exclude foreign users from two of its most powerful models.[3] Unable to verify the nationality of users quickly enough to comply, the company suspended both models worldwide, restoring access only when the relevant restrictions were lifted at the end of that month.[4] A government or business that depends on a single foreign supplier for frontier capability now has direct evidence of how quickly that dependency can turn into a liability, regardless of which capital city enforces the restriction.

Brookings responds to this risk with a concrete proposal: bringing together those layers of the AI stack that require scale, primarily computing power, large model training processes and public procurement, while leaving development, applications and industry-specific data under national control. It points to a reasonable partnership between the European Union, Japan and South Korea, each of which will contribute complementary strengths in industrial data, memory semiconductors and financial capacity, tied together by mutual public procurement commitments loosely based on the Airbus model.[5] The text is candid about the obstacles: unresolved trust issues between Tokyo and Seoul, chronic European fragmentation and the need for a governance structure solid enough to withstand disagreements over money and control.[6] Where the argument weakens is in its certainty that these obstacles can be overcome once the stakes are clear enough.

Southeast Asia offers a closer look at what happens when a comparable group of middle powers actually attempts to negotiate, develop and govern something together, under conditions of much lower technical com-

plexity than the AI stack that Brookings has in mind. The Association of Southeast Asian Nations (ASEAN) has spent six decades building exactly the kind of institutional muscle that a coalition of middle powers would need for AI: permanent advisory bodies, a culture of consensus and a shared interest in resisting the pull of larger neighbors.[7] Its recent history, examined below, shows that even this muscular system is severely strained once real money, real bargaining power and a real decision on benefit-sharing are on the table. If a coalition with ASEAN's institutional track record can't quickly align itself into a tariff response or a digital economy agreement, the assumption that Europe, Japan and South Korea, three governments with a much shorter history of joint action, could set up a jointly governed AI consortium deserves more skepticism than the modular sovereignty argument allows.

This text does not dispute that reliance on American or Chinese AI entails real costs, nor that middle powers possess genuine assets worth defending. It argues that whether these assets can be pooled is a separate question from whether they exist and that the second question has been answered too hastily. It then examines what a coalition could reasonably offer its potential members, what the experience of a comparable regional group indicates about the chances of forming and operating such a scheme and what alternatives, from bilateral alignment to loose mini-multilateral consultation, the middle powers could reasonably pursue if the ambitious version of modular sovereignty proves to be as difficult to govern as history suggests.

2 The Assets Middle Powers Could Bring to a Shared AI Strategy

The case for collective action starts from an accounting problem. Individually, no middle power has enough computing power, capital, or modeling talent to influence the evolution trajectory of cutting-edge AI. Collectively, the picture is different. Japan offers advanced robotics, precise manufacturing processes and decades of experience in developing automation on factory production lines. South Korea offers leadership in high-bandwidth memory and other semiconductors, alongside a rapidly improving base of domestic fundamental models. The European Union offers a continental market, deep capital markets and public procurement amounting to around two trillion euros per year, around 13.6 percent of its gross domestic product.[8] None of these assets individually rival the scale of Silicon Valley or Beijing's state-run programs. Combined, they represent a substantial, if still modest, counterweight, provided they can be organized.

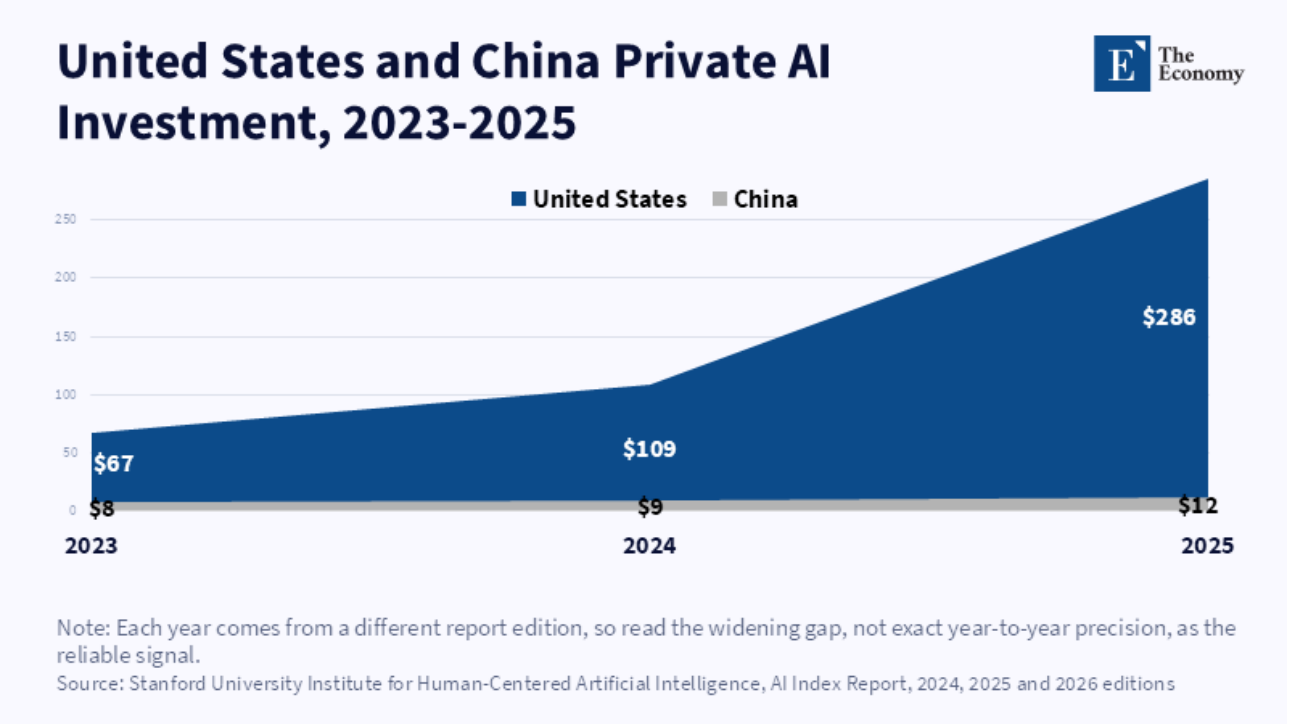

Figure 1: The US lead in private AI investment highlights the financial scale confronting middle powers. Figures are in US$ billions and exclude much state funding.

Physical and industrial AI is where Brookings identifies the clearest window of opportunity. U.S. platforms dominate cloud infrastructure and general-purpose models, but their place in robotics, manufacturing app development and other physical applications is less established, leaving room for a middle-strength effort focused on areas where jurisdiction, security and operational control matter more than gross costs.[9] China is already strong in this area: it accounted for 54 percent of new industrial robot installations globally in 2024 and has an operational inventory of about two million units, about four and a half times more than Japan.[10] This scale rules out competing on price; it does not, however, exclude competition in regulated sectors: hospital systems, electricity networks, water supply networks, ports, telecommunications and defence equipment maintenance, where an operator under foreign control may be unacceptable to regulators and insurers, regardless of price.[11]

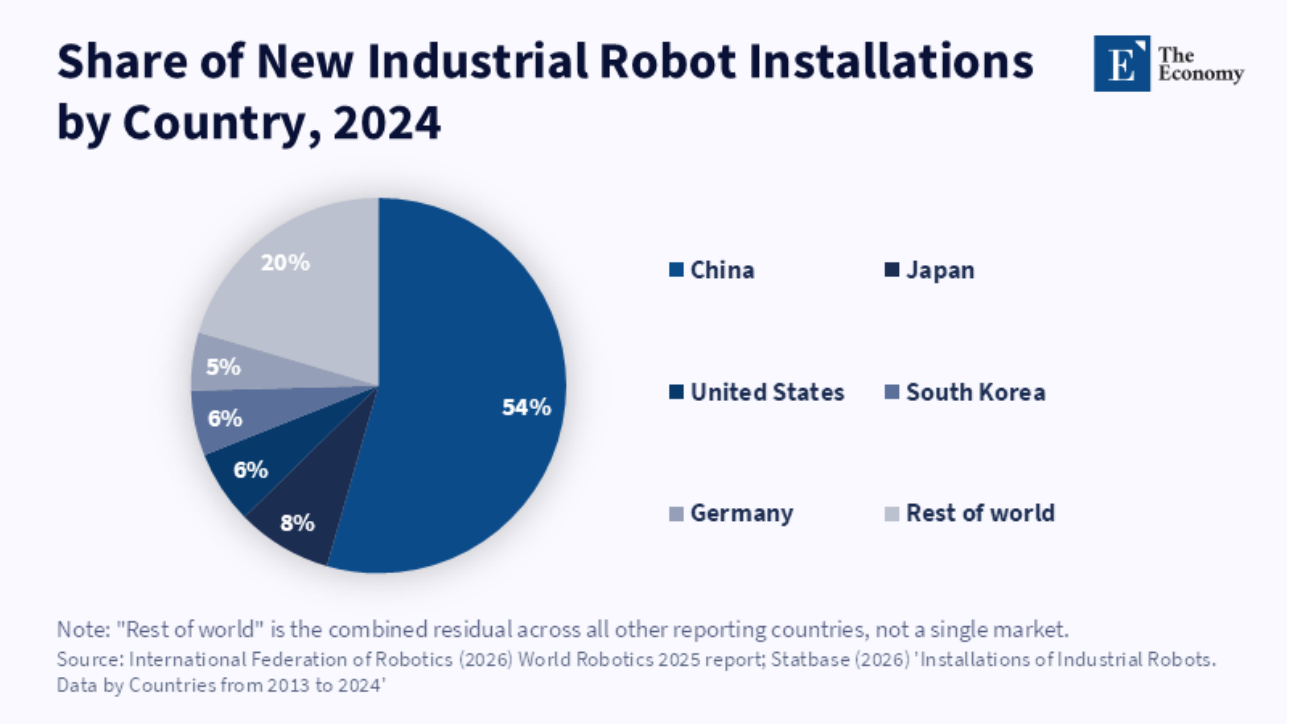

Figure 2: China installed more than half of the world's new industrial robots in 2024, setting a formidable benchmark for middle powers' deployment ambitions.

Japan's Noetra initiative shows what a national contribution to this strategy could look like. Selected by the Agency for the Development of New Energy and Industrial Technology (NEDO), the program has attracted investment from forty-four companies and institutions, including MUFG Bank, with research underpinning the National Institute of Advanced Industrial Science and Technology and Preferred Networks, alongside industry participants such as Sony, SoftBank, NEC and Honda. Its ambitions go beyond yet another national chatbot, pivoting towards multimodal models for manufacturing, supply chain, energy systems and healthcare, i.e., the physical environments where Japanese engineering and operational data have genuine value.[12] The initiative also reveals the limits of what any country can secure on its own: its planned computing cluster relies on about 27,500 Nvidia Rubin graphics processors, leaving the program exposed to U.S. jurisdiction at exactly the level of computing power, where Brookings argues that concentrated capacity would matter most.[13]

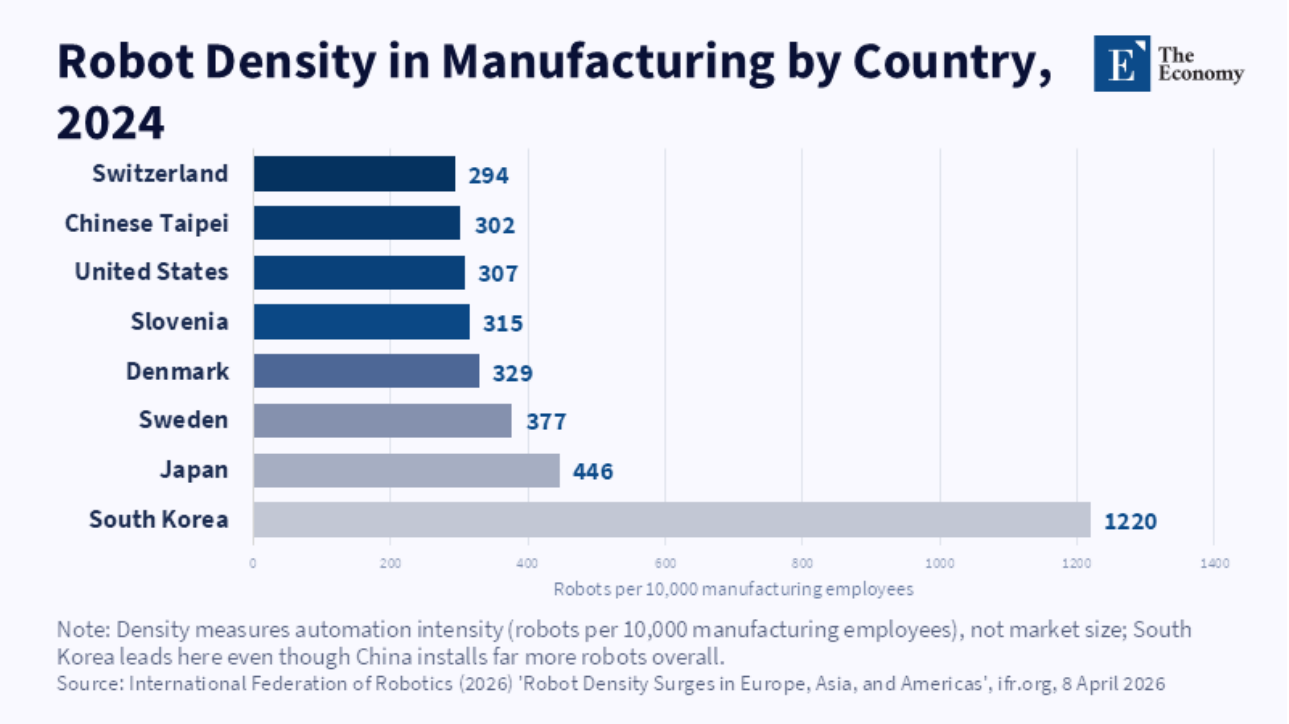

Figure 3: South Korea's robot density highlights the intensive automation that some middle powers bring to industrial AI, alongside China's advantage in total installation volumes.

South Korea occupies a complementary position. Its Sovereign AI Foundation Model program combines public support with competition among domestic consortia and by mid-2026 four Korean models had been added to Epoch AI's list of notable systems, compared to fifty-nine from the United States and thirty-five from China the previous year.[14] The gap with the edge remains large, but the path is unusual for a country the size of South Korea and it comes with leadership positions in memory semiconductors, batteries, shipbuilding and telecommunications equipment. South Korea has also demonstrated a capacity to broker compromises that neither the United States nor China can easily replicate. At the Asia-Pacific Economic Cooperation (APEC) Leaders' Meeting in Gyeongju in November 2025, Seoul helped reconcile conflicting blocs, economies that prioritised innovation and resisted regulatory language and others pushing for stronger safeguards, reformulating controversial terms such as 'trustworthy' into functional language such as 'safe and trustworthy adoption', which all sides could accept; the resulting APEC Artificial Intelligence Initiative for 2026-2030 also created a Korean-led Asia-Pacific Artificial Intelligence Center, although only after the proposal was recast as a self-funded, independently operated body, following the resistance of other members to the adoption of what resembled a single-country initiative within a multilateral text. A joint simulation of Korean and British researchers, based on global patent data, found that in a scenario of a United States-China technological decoupling, economies in place of South Korea could suffer up to sixteen percent greater loss of innovation than the United States itself, a clue that helps explain Seoul's apparent willingness to hedge against just that eventuality.[15]

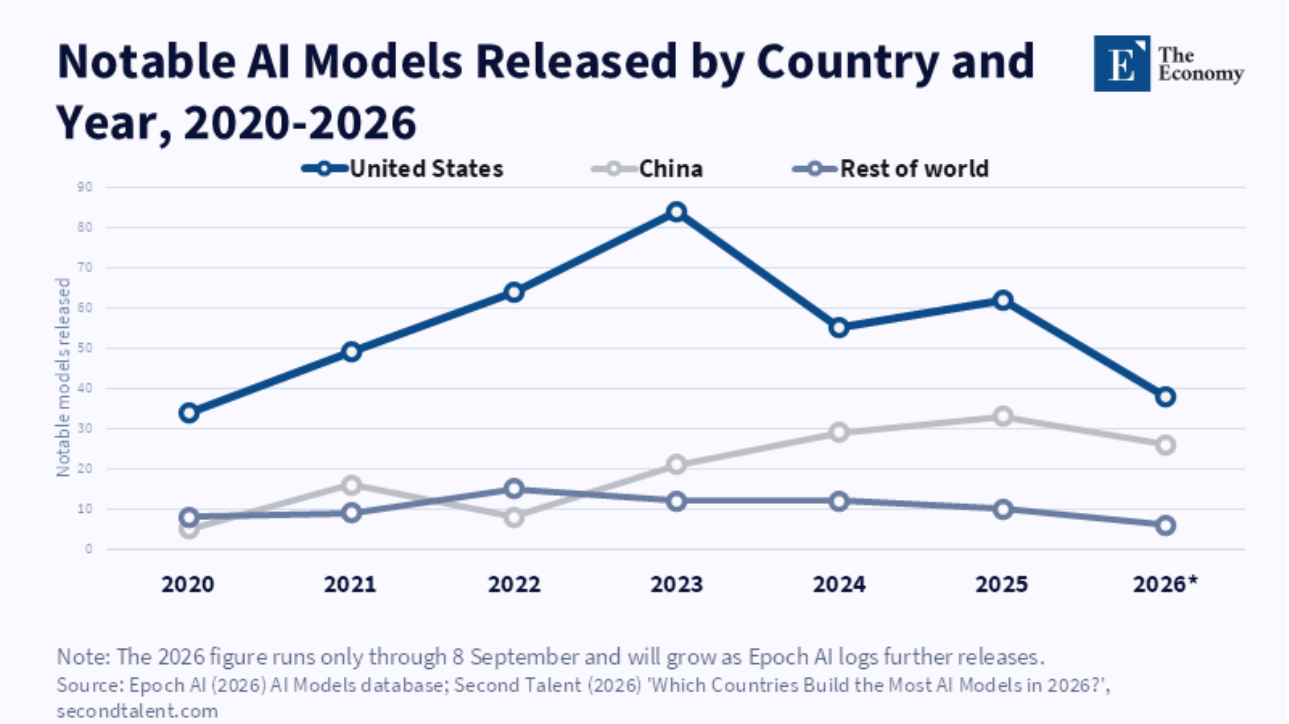

Figure 4: China's notable model output expanded during 2023–2025, while the US remained ahead. Joint US–China models are excluded; 2026 covers releases through 8 September only.

Europe's contribution is potentially the largest and so far the least implemented. The InvestAI initiative, announced in February 2025, set a target of mobilising €200 billion, of which €20 billion was earmarked for European AI "giant factories" and in July 2026 the European Commission launched a tender for up to seven such facilities, supported by up to €10 billion in combined European Union and national funding, around half of which from Brussels, to attract at least €20 billion of private capital.[16] The proposed Cloud and Artificial Intelligence Development Act, tabled in June 2026, would introduce non-pricing criteria in the procurement of cloud and AI systems by public entities, a modest step towards the kind of demand-side bargaining power that a 2023 International Monetary Fund staff paper suggests Europe could finance at scale through a common safe asset, without substantially increasing its borrowing costs. Any such procurement-based strategy would still run counter to the World Trade Organization's Agreement on Government Procurement and the European Union's own procurement directives, which limit discrimination against foreign suppliers, although there are exceptions for security reasons that should bear the brunt of reciprocal purchasing.[17] Mistral's expansion into regional inference and third-party infrastructure hosting remains France's most visible contribution to European model capability, although its growing reliance on Nvidia hardware and Microsoft funding shows how far a single national champion is from competing alone on the American or Chinese scale.[18]

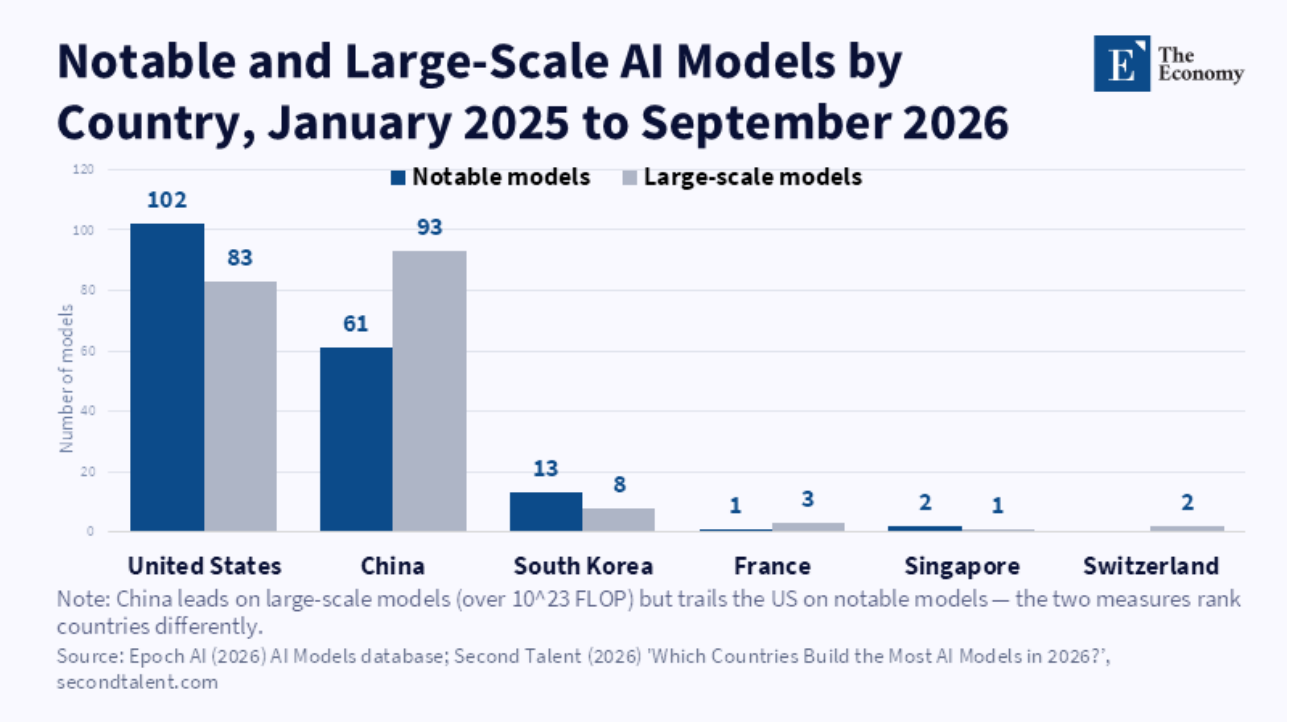

Figure 5: The U.S. leads in notable models, while China records more models above the large-scale training threshold—different measures of AI production, covering January 2025–8 September 2026.

A different kind of benefit runs through cybersecurity instead of industrial growth. Anthropic's Claude Mythos Preview and OpenAI's GPT-5.5-Cyber, both launched in 2026 with advanced offensive and defensive capabilities, were only made available to a narrow set of partners: primarily U.S. tech companies, banks and government agencies, as well as a small number of allied institutions, such as the UK's Artificial Intelligence Security Institute.[19] Everyone else, in Chatham House's perspective, faces a choice between two paths: aligning with a superpower's technology in exchange for access, or coordinating with each other to build a common assessment capacity, information exchange channels and collateral communication channels for times of crisis. The same analysis notes that there is already a loose network, connecting the European Union, the United Kingdom, Japan, Kenya and others, although it was not designed for political action and would need deliberate amplification to matter in a real crisis.[20]

New America's Shangri-La series, which brought together twenty-one policymakers, technologists and researchers from Africa, Asia, Latin America, North America and Europe in Honolulu in May 2026, has come up with a narrower but more concrete answer as to where the bargaining power of middle powers lies.[21] Participants converged on the idea that countries that do not develop frontier models can nevertheless shape how these models adapt to local languages, legislations and institutional frameworks, what the dialogue called the "level of adaptation" and that assessment and assurance standards are an undervalued lever precisely because independent testing infrastructures are almost non-existent outside of a handful of affluent institutions. The meeting minutes themselves show the ambition shrinking as soon as the real debate began: shared ownership of models and co-managed infrastructure emerged as avoidable risks, not objectives around which participants were organized, while the adaptation layer, evaluation standards and rulemaking emerged as the areas where the middle powers felt they had genuine, exercisable influence.[22]

Overall, the argument in favor of benefits is genuine but narrower than its impressive version. Japan, South Korea and the European Union each have assets that matter, in robotics and industrial data, memory semiconductors and model development, capital and market scale. Middle powers acting through mini-multilateral consultation, bilateral partnership and standard-setting have shown that they can derive real value from these assets, as indicated by both South Korea's diplomacy at APEC and the CETaS-proposed South Korea-UK partnership in standards, infrastructure and research.[23] What the evidence so far does not show is whether these same governments can move from this kind of narrow, functionally limited cooperation to jointly capitalized, qualified-majority-ruled, based on consortium procurement required by modular sovereignty in its fullest form.

Even this narrower accounting highlights an allocation problem that any coalition would have to solve before it could function. Japan's bargaining power relies on data and applications that it cannot easily extend to partners without exposing proprietary industrial know-how. South Korea's bargaining power rests on a position in semiconductors and revenue-generating memory regardless of whether a coalition is ever formed. The European Union's bargaining power rests on financial capacity and market size that for now remain mostly untapped ambition rather than developed assets. A coalition that would pool these unevenly weighted contributions would have to decide, credibly and in advance, how to share the benefits among partners who do not come to the table with equal cards. This is a question of governance, not a technical issue and it is the question that the next section deals with through the experience of a group of middle powers that has spent significantly more time trying to answer it.

3 Why Shared Interests Do Not Guarantee Collective Action

The Association of Southeast Asian Nations offers the most immediate test available of what happens when a comparable group of middle powers attempts to act as a single whole. Founded to eliminate intra-regional conflict between ten, now eleven, culturally and economically heterogeneous states, it has built its cohesion around what has been established as the "ASEAN Way": consensus instead of majority normative, non-intervention instead of public pressure and silent diplomacy instead of the tactics of public stigmatization common in Western institutions. For six decades, this approach has helped transform one of the world's most heterogeneous regions into one of the most stable, giving ASEAN a convening role in forums such as the ASEAN Regional Forum and the East Asia Summit, which neither Washington nor Beijing have been able to completely bypass.[24] If institutional longevity and common purpose were sufficient conditions for collective bargaining power, ASEAN would be the strongest available evidence that middle powers can be effectively organized. Its recent track record suggests something more complex.

In April 2025, Washington announced reciprocal tariffs that hit Southeast Asian exporters particularly hard. Malaysian Prime Minister Anwar Ibrahim, holding the rotating presidency of ASEAN, moved quickly, convening leaders by phone and pledging to coordinate efforts to present a united regional front and ensure that ASEAN's collective voice was clearly heard on the international stage. This was followed within a week by an emergency meeting of ASEAN economy ministers. Wan Shuhaimi Wan Mohd Saidi, head of economic research at Kenanga Investment Bank, summed up both the appeal and the limitation of the moment: the

combined weight of the coalition, a $3.6 trillion economy with over 650 million people, could attract real attention if exercised collectively, he pointed out, while in the same breath he warned that ASEAN, by design, is a consensus-based organization and not a negotiating bloc with "teeth".[25] That warning proved prophetic. Bilahari Kausikan, Singapore's former ambassador to the United Nations, known for his unequivocal judgments on regional diplomacy, rejected Anwar's proposal entirely, arguing that it was impossible for ASEAN, given the breadth of tariffs and the differences between its economies, to negotiate collectively. The Trump administration resolved the issue by refusing to negotiate with ASEAN as a coalition at all, insisting instead on bilateral talks with individual member states, a preference that gave Washington exactly the bargaining power that a united front intended to deprive it of.[26]

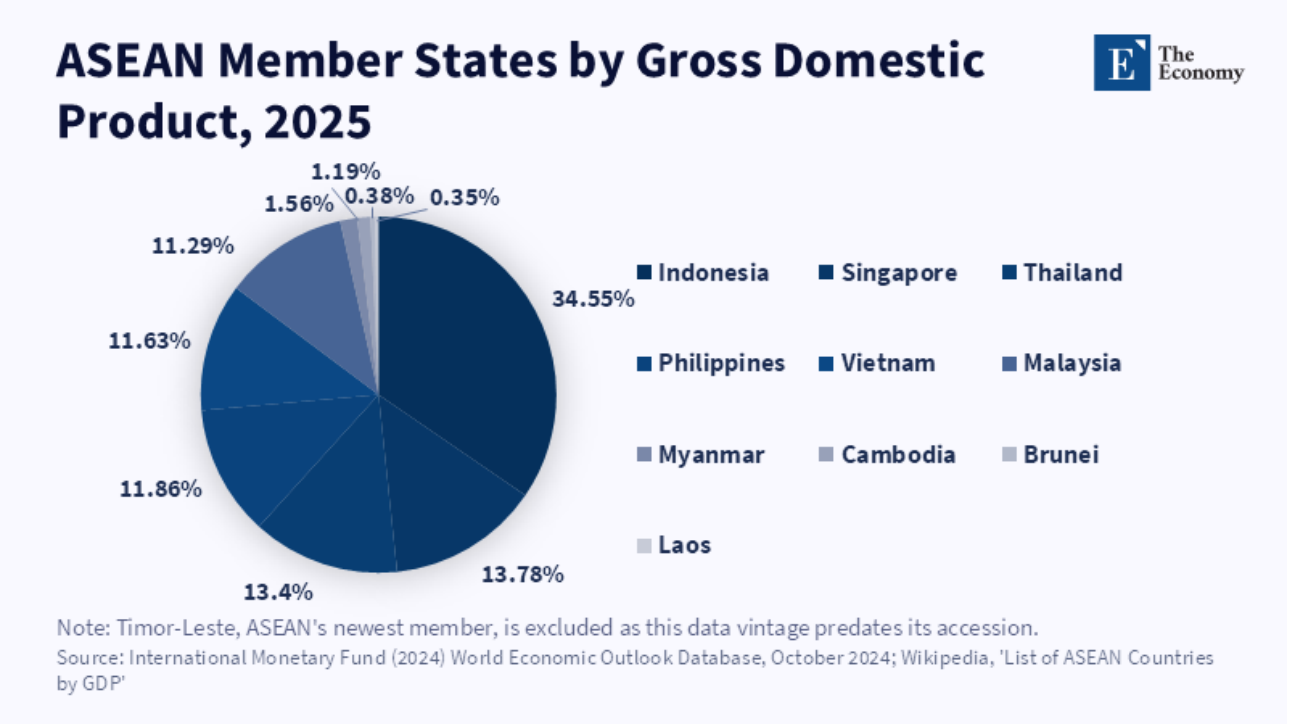

Figure 6: Indonesia's large share highlights ASEAN's unequal economic weights. Figures are 2025 forecasts from the IMF's October 2024 dataset, excluding Timor-Leste.

The same pattern was repeated later that year in a security context. When a border dispute between Cambodia and Thailand escalated into armed conflict, ASEAN was initially sidelined because both governments preferred other channels and it was Malaysia's individual mediation as the presiding country, not some permanent ASEAN mechanism, that yielded a ceasefire and the deployment of an observer group, an initiative that had already failed once before. The civil conflict in Myanmar, meanwhile, has spanned several rotating presidencies without ASEAN reaching a common position, a gap serious enough for Timor-Leste's accession as an eleventh member in 2025, the first enlargement of the coalition since the 1990s, to attract more attention as a milestone of institutional continuity than as evidence of renewed collective capacity.[27]

Even ASEAN's most tangible recent achievement shows how long consensus takes to yield results. Negotiations on the ASEAN Digital Economy Framework Agreement, an agreement aimed at harmonizing digital trade rules across the coalition and explicitly touching on artificial intelligence, financial technology and source code protection, began in 2023 under the presidency of Indonesia. They were only concluded at the fifty-seventh

round of talks of senior economic officials in Manila in May 2026, more than three years later and the agreement aims to be signed only in November 2026, with a further ratification period of 180 days in each member state thereafter. The officials involved acknowledged that the text was not yet perfect even at the time of completion and would need periodic revision to keep pace with a rapidly changing digital economy.[28] This is the timetable for a framework agreement of rules, not for a jointly capitalised industrial enterprise with concentrated computing power, joint ownership of models and binding supply commitments of the kind provided for by modular dominance for Europe, Japan and South Korea.

ASEAN's track record is not devoid of genuine successes and it would be unfair to read only its failures. The coalition's Regional Comprehensive Economic Partnership (RCEP), the world's largest trade agreement by population, has just been concluded and is proof that broad multilateral coordination between Asia-Pacific economies is possible. It is also proof of how long that achievement takes: negotiations began in November 2012 and were not concluded until 2020, eight years into the alignment of fifteen governments on tariff protocols and rules of origin.[29] This timeline hardly fits with the urgency that Brookings itself attributes to securing a place in physical and industrial AI before foreign platforms take root beyond the point of switching suppliers. The current trajectory of AI itself, to which the national strategy can shift in just one year, as shown by both the revocation of the restriction on Nvidia's H20 chip and the suspension and reinstatement of Anthropic's most powerful models, is not obviously compatible with institution-building measured in years rather than months.[30]

Europe offers a second element and it comes from within the modular sovereignty argument itself. Brookings' own description of the Pax Silica initiative notes that the European Union only joined at the end of June 2026, after France had initially opposed participation, a disagreement that the report itself calls a depiction of European fragmentation. The European Union's AI "giant factories" program tells a similar story: public funding commitments decreased, not increased, within eighteen months, even as the rhetoric around European AI dominance became more pronounced in speeches, strategies and funding announcements.[31] A continent that struggles to align its member states on where to host computing power and who to pay for it is not an obvious basis for a trilateral consortium that would also have to reconcile Japanese and South Korean positions on the same issues.

Academic studies of middle-power diplomacy explain why bilateral arrangements tend to succeed where multilaterals get stuck. Research on South Korea's New South Policy found that Seoul secured significantly deeper free trade agreements with individual ASEAN members, building a distinct middle-power identity and exercising bargaining power bilaterally, even where ASEAN's collective institutions showed initial reluctance.[32] The finding echoes a pattern identified in the fundamental literature on middle-power behavior, which traces the influence of such states to niche diplomacy, to building coalitions around specific, narrowly defined issues, rather than to attempts to build binding, jointly governed institutions from scratch.[33] A subsequent study of the Proliferation Security Initiative, a security arrangement deliberately constructed without a formal institutional structure, similarly found that the middle powers exerted more practical influence through flexible, voluntary cooperation than classical middle-power theory built around institution-building would envisage.[34] Both streams of the literature lead to the conclusion that the evidence in this section directly supports: the middle powers possess a proven ability to derive benefits from bilateral bargaining power and loosely structured

cooperation and much less demonstrated ability to build and maintain the kind of jointly capitalized, qualified majority governing consortium that would require a full coalition of modular sovereignty.

4 Alignment, Coalitions and the Case for Narrower Partnerships

Chatham House's very description of the choice facing middle powers in AI cybersecurity boils down purely to two paths: alignment with a superpower in exchange for privileged access, or coordination between the middle powers themselves to build joint capability.[35] The same dichotomy applies, with higher stakes, to the broader question addressed in this text. Multilateral coordination, what might be called a multi-play approach, promises collective bargaining power without submission to either Washington or Beijing. Its most comprehensive historical test is not encouraging.

The United Nations was built, as a recent analysis for the Lowy Institute argues, on a fundamental compromise: veto power for the five permanent members of the Security Council, adopted specifically to keep the great powers within the institution rather than risk a repeat of the collapse of the League of Nations, which fell apart when the great powers simply walked away from restrictions they didn't like. The compromise achieves this narrow goal. It also means that the Council can only act effectively where the interests of the great powers are limited or aligned and becomes powerless at the moment when they clash, a pattern visible in Liberia and South Sudan, where peacekeeping missions have worked and in Ukraine and the war involving the United States, Israel and Iran, where they have not worked. Enlarging the composition of the Council, a reform regularly proposed as a solution, would not resolve this dynamic but rather redistribute it, as new permanent members with veto power would simply multiply the points where the institution may be paralysed.[36] The evidence confirms this with unusual clarity: seven draft Security Council resolutions failed to be adopted in 2024 due to vetoes, the highest number since 1986, with eight vetoes in total on these failures, four by Russia, three by the United States, one by China and four more vetoes followed in 2025.[37]

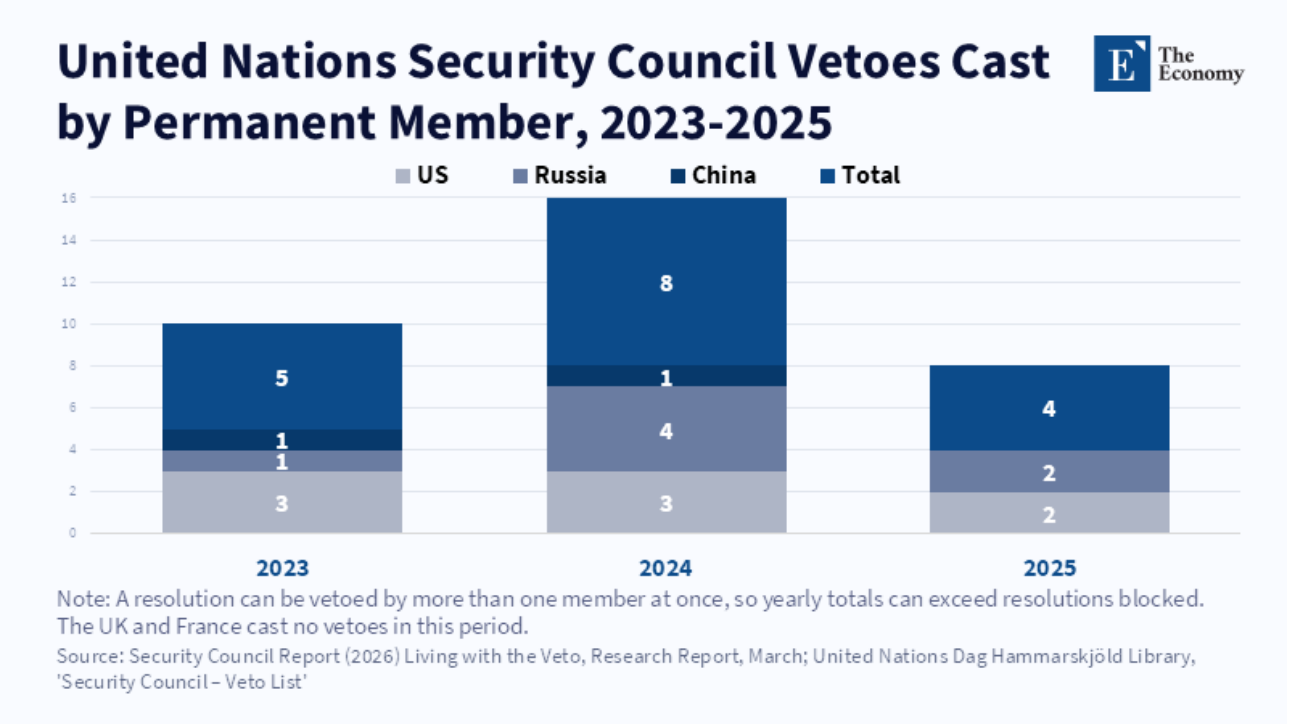

Figure 7: Vetoes illustrate how national interests can block collective action.

The consequences of this planning were harshly highlighted in cases where prevention was of greatest importance. During the Rwandan genocide in 1994, the United Nations mission was reduced from 2,500 to about 400 soldiers, even as its commander called for reinforcements, a decision determined by the United States, the United Kingdom and France, rather than any operational assessment on the ground. The following year, the peacekeeping force in Bosnia, under a commander reluctant to authorize air strikes on besieging positions, was replaced only after the Srebrenica massacre had already taken place. A way out exists in the 1950 Uniting for Peace resolution, which allows the General Assembly to recommend collective action when the stalemate in the Security Council prevents it and which has been invoked eleven times, most recently since the Russian invasion of Ukraine in 2022. Its rarity, eleven uses in seventy-five years, demonstrates the limits of the mechanism as much as its promise: a way out that first requires the failure of the Council and remains available only when several governments are willing to circumvent precisely the veto structure that created the stalemate in the first place.[38]

None of this is a perfect analogy for an AI coalition, which would not carry the life-and-death stakes of a Security Council and would not need anything analogous to a veto on military power. However, the underlying mechanism, which requires unanimity or qualified majority agreement between governments with divergent interests before any collective action can proceed, comes close to the governance structure proposed by Brookings itself: a qualified majority approval process for major allocations, combined with a ”contribution versus access” principle that determines who gets what in return for what they contribute. This structure solves the theoretical question of how the benefits would be distributed. It doesn't solve the practical question of what happens when a Japanese, a South Korean and a European government, or a broader circle that includes Canada, India and the Gulf states, disagree, as they will eventually argue, about whose data they trained which model, whose factory hosts a shared computing cluster, or which country the regulator has jurisdiction over a piece of critical

infrastructure. The very Brookings benchmark for institutional failure, the Gaia-X initiative, which devoted years to standards and pilot programs without producing a European competitor in the cloud because no single party had responsibility for implementation, is itself a small-scale preview of the same dynamic.[39]

Proponents of an effort despite obstacles could reasonably respond that waiting for perfect governance guarantees failure by definition, as Brookings itself argues that the window of opportunity for physical and industrial AI is narrowing as foreign platforms accumulate proprietary operational data and become harder to displace. Airbus, according to this perspective, did not start out as a fully settled institution either; it built credibility and implementation capacity gradually, through repeated launch orders, rather than settling every governance issue in advance. This argument carries real weight and is the strongest version of the argument in favor of moving forward despite the unresolved issues of presidency and ownership. It is, however, weaker when applied to artificial intelligence than to Airbus, because Airbus solved the ownership issue early on by entrusting implementation responsibility to a single corporate entity with its own balance sheet, a structure that the governments of Europe, Japan and South Korea have yet to agree to replicate and because the very life cycle of an AI product, in which entire generations of capabilities are renewed within a year, leaves much less room for the kind of gradual, ten-year trust-building that made Airbus' slow start viable.[40]

The alternative to a multi-play approach is not necessarily full alignment with a single superpower, although this option retains real appeal. A draft letter reportedly released by the State Department in August 2026 informed the thirty-five signatories of an earlier American AI framework that commitments under the Pax Silica Declaration were incompatible with engaging in "overlapping initiatives," a wording broad enough to cover almost any independent effort by middle powers, not just overt cooperation with China. China's own World Artificial Intelligence Cooperation Organization, launched in July with twenty-nine founding members, offers a parallel pull of alignment and Kazakhstan's decision to join both organizations at the same time shows that some governments are already hedging the risks between the two rather than waiting for a middle-power alternative to mature.[41] Japan and South Korea had joined Pax Silica long before the draft letter was reportedly released, a demonstration that, in the absence of an already built credible alternative, alignment remains the path of least resistance, even when it limits precisely the independence that a coalition is meant to protect.[42]

An intermediate path between full alignment and full coalition building can be seen in arrangements that the middle powers have actually managed to maintain. South Korea's bilateral free-trade diplomacy with individual ASEAN members under the New South Policy secured deeper agreements than collective bargaining yielded, precisely because it operated outside the demand for pan-Asian consensus.[43] The CETaS-proposed South Korea-UK partnership, built around three narrow pillars, interoperable assessment standards, joint investment in infrastructure targeted at the Global South and shared open research resources, follows the same logic: two governments, a defined scope and no requirement for a third or fourth party to agree before anything can move forward.[44] The findings of the Shangri-La meeting themselves point in the same direction; when the representatives of the middle powers sat down together at the table, their most concrete proposals were gathered around the adaptation layer, evaluation standards and common open tools specialized by task, functions that do not require shared ownership of frontier compute or models, but rather around the kind of jointly capitalized industrial consortium that modular dominance envisions.[45]

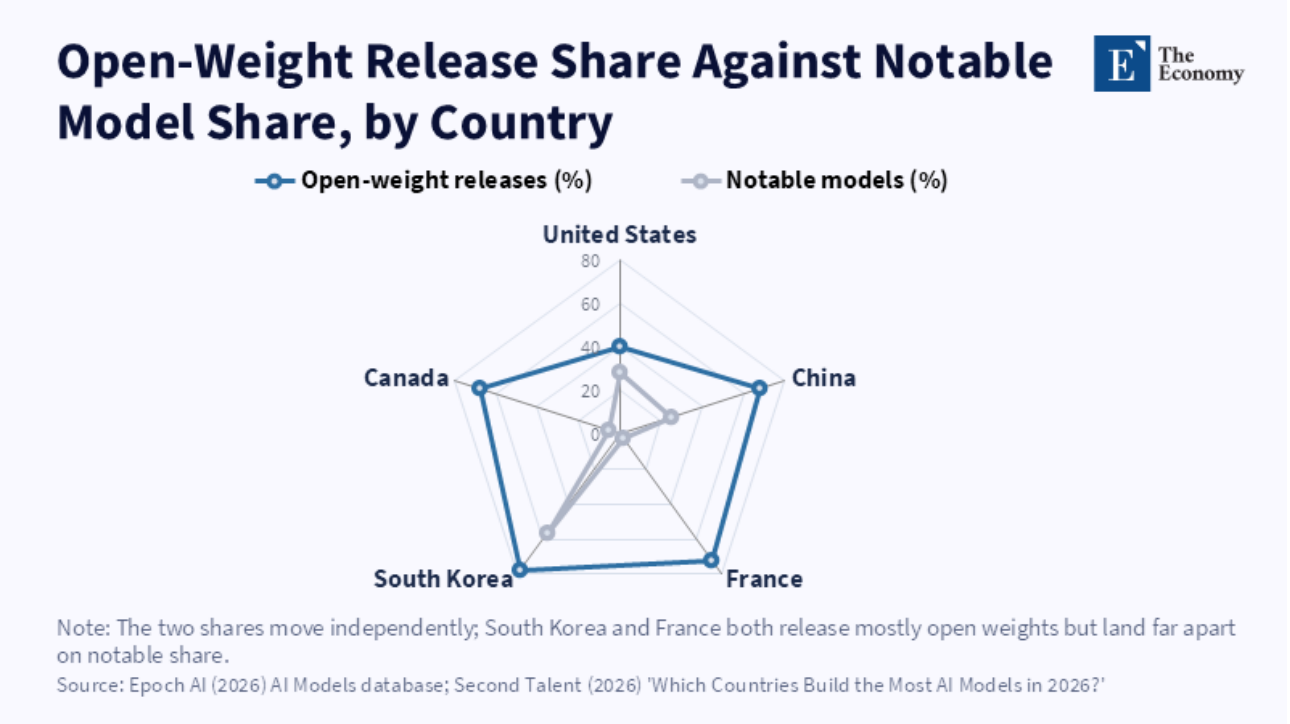

Figure 8: Openness alone does not establish technological autonomy. Within each country, the open-weight share covers models with a known release type; the notable share covers all recorded models.

Other loosely structured middle-power groups tell a similar story. MIKTA, the informal consultative forum that has brought together Mexico, Indonesia, South Korea, Turkey and Australia since 2013, has never issued binding commitments or pooled resources and precisely because it requires so few of its members, it has outlasted several more ambitious multilateral proposals that promised more and yielded less.[46] Its moderation is the source of its duration, not a failure of greater ambition. A middle-power effort for AI modeled on this scheme, consultation without centralized ownership, would miss the scale advantages that Brookings identifies as the meaning of modular dominance in the first place. That compromise is real and it is the compromise that the evidence in this text suggests that the middle powers are actually capable of making.

This is not an argument that pooling resources is always impossible, or that scale never matters. A separate analysis by the Brookings Institution defines AI dominance as a spectrum of strategies rather than a binary option and calls instead for ”managed interdependence” through alliances that diversify risk across the stack.[47] A recent standard treatment of AI dominance advances the same logic further, modeling it as a continuum across several separate pillars, data, computing power, models and rules, arguing that a government’s optimal position varies by pillar rather than applying uniformly across the stack.[48] Read this way, the debate on middle-power AI strategy has confused two different claims: that valuable cooperation on specific, narrow pillars is feasible, which the evidence broadly supports and that a single, totally governed coalition spanning computing power, models, procurement and industrial development is feasible on a timeline that matters, something that the evidence, from ASEAN’s response to tariffs to the Digital Economy Framework Agreement to the historic Security Council veto, does not support. The middle powers seem to be better served by treating cooperation as a portfolio of close, bilateral and mini-multilateral arrangements tailored to specific pillars and specific partners, rather than as a single grand coalition waiting for the right governance formula.

5 Conclusion - Building AI Agency Through Targeted Cooperation

The European Union and the Association of Southeast Asian Nations represent the world's two longest-standing examples of middle powers building lasting collective mechanisms: seventy years of integration in the former, six decades of consensus-based diplomacy in the latter. Both offer genuine lessons and both, based on the evidence examined here, point to persistent limits even within their areas of greatest success, a common market that still cannot quickly agree on where to accommodate computing power, a regional coalition that took three years to complete a framework agreement that touches on artificial intelligence only marginally and was unable to present a unified position to Washington on tariffs that hit each member state at the same time. A coalition spanning Europe, Japan and South Korea, or a wider circle encompassing Canada, India and the Gulf states, would start from a weaker institutional base than either.

Without an orderly answer to who presides over a joint effort, who owns the resulting models and who reaps the largest share of the benefit, ambition tends to converge toward the fate of the Security Council: a structure built to keep the key players within the institution that becomes unable to act when their interests actually diverge. That doesn't make coordination worthless. It makes it a different kind of project than what modular sovereignty proposes today: narrower in scope, bilateral or mini-multi-party in structure and built around specific pillars rather than a single, jointly governed consortium waiting for the right qualified majority formula. The middle powers retain real assets and real choices. The pragmatic path passes through the accumulation of narrow, functional arrangements rather than through a grand coalition that history gives little reason to expect to remain united when money and recognition are actually on the table.

References

[1, 2, 3, 5, 6, 8, 9, 10, 11, 12, 13, 14, 16, 17, 18, 30, 31, 39, 40, 41, 42] Rediker, D.A. (2026) 'Middle power AI agency: preserving choice between the United States and China', *Brookings*, 8 September.

[4] Anthropic (2026) 'An update on Claude Fable and Mythos access', *Anthropic News*, 1 July.

[7, 24] Elizabeth, R. (2025) 'Southeast Asian diplomacy in the era of great power rivalry: the ASEAN way', *Asia Scotland Institute*, 1 July.

[15, 23, 44] Lee, S. and Lee, H. (2026) 'Asia-Pacific AI governance: a call for middle-power partnerships', *CETaS Commentary*, The Alan Turing Institute, 7 May.

[19, 20, 35] Wilkinson, I. (2026) 'In the face of growing AI cyber threats, do middle powers have agency?', *Chatham House*, 3 June.

[21, 22, 45] Coral, L. (2026) 'The Shangri-La Series: AI for middle powers', *New America*, 29 June.

[25] Sarmiento, P. (2025) 'ASEAN unity vital in fight amid tariffs', *China Daily* (Hong Kong edition), 11 April.

[26] MalaysiaNow (2025) 'Prominent Singapore diplomat rubbishes Anwar's proposal for ASEAN to collectively discuss tariff with US', *MalaysiaNow*, 10 May.

[27] Real Instituto Elcano (2026) 'La ASEAN en 2025: entre el avance y el estancamiento', *Real Instituto Elcano*, Comentario.

[28] Rajah & Tann Asia (2026) 'ASEAN Digital Economy Framework Agreement: negotiations concluded, targeted for signing in November 2026', *Rajah & Tann Asia*, Legal Update.

[29] Department of Trade and Industry, Philippines (2025) *Regional Comprehensive Economic Partnership (RCEP).*

[32, 43] Lee, S.Z. (2023) 'Middle power and power asymmetry: how South Korea's free trade agreement strategy with ASEAN changed under the New Southern Policy', *Contemporary Politics*, 29(3), pp. 318–338.

[33] Cooper, A.F., Higgott, R.A. and Nossal, K.R. (1993) *Relocating Middle Powers: Australia and Canada in a Changing World Order.* Vancouver: UBC Press.

[34] Cooper, D.A. (2011) 'Challenging contemporary notions of middle power influence: implications of the proliferation security initiative for "middle power theory"', *Foreign Policy Analysis*, 7(3), pp. 317–336.

[36] Kumar, A. (2026) 'The UN is not broken — it was built this way', *Lowy Institute, The Interpreter*, 5 May.

[37] Security Council Report (2026) *Living with the Veto.* Research Report, March.

[38] Stanton, G.H. (2024) 'Looking backward into the future: why the United Nations has failed to prevent genocide', *Cairo Review of Global Affairs*, Summer.

[46] Patrick, S.M. and Feng, A. (2018) 'MIKTA in the middle: a little-known multilateral group turns five', *Council on Foreign Relations*, 18 September.

[47] Tanner, B., Kerry, C.F., Wyckoff, A.W., Kyosovska, N., Renda, A. and Tabassi, E. (2026) *Is AI Sovereignty Possible? Balancing Autonomy and Interdependence.* Washington, DC: Brookings Institution.

[48] Singh, S.K. and Sengupta, S. (2025) 'Sovereign AI: rethinking autonomy in the age of global interdependence', *arXiv preprint*, arXiv:2511.15734.