

OCTOBER 2026

How Middle Powers Can Survive the AI Race: Four Scenarios for an Uncertain Future

Alasdair Phillips-Robins, Teddy Tawil, and Sam Winter-Levy

U.S. allies may have breathed a sigh of relief earlier this year when AI developer Anthropic [announced](#) that it was extending access to Mythos, its most advanced AI model, to governments and institutions in fifteen partner countries. Any relief, however, was short-lived. Less than two weeks later, the U.S. government directed Anthropic to cut off all foreign access to Mythos and Fable, a more limited public version of the model. The administration also appeared to strong-arm OpenAI into [limiting](#) access to its newest model, before lifting restrictions on both companies.

The episode provided a stark reminder that much of the world depends on the whims of the U.S. government for access to frontier AI. American companies produce the most powerful models, and the United States and China command a large share of the world's top AI researchers, computing power, and investment. Foreign governments, fearful of missing out on the benefits of advanced AI while still bearing the brunt of its harms, are scrambling to respond.


The right strategy will depend on each country's current position and on the trajectory of AI progress. Some middle powers are deeply embedded in frontier AI

supply chains, while others are primarily exposed to its effects. On the technological front, two variables should drive middle power choices: the speed of frontier progress and what we call the “frontier premium,” the gap between the best systems and open-source or other cheap alternatives. Together, these create four possible worlds for middle powers, each with its own best strategy (see figure 1).

Of course, the variables are not fully independent—compute-intensive scaling may produce both rapid progress and advantages for the leading country—and the world may move between quadrants over time. But the matrix provides a guide to the range of possibilities. Outlining each scenario illustrates how it might arise, who would win and lose, and what governments can do to improve their position before knowing which outcome will appear.

So far, governments and analysts have advocated many potential strategies, but few are considering the full range of possible futures (see figure 2). Some governments want to build sovereign models. Some argue they should pool their resources to build frontier models and hardware together. Some believe countries can gain advantage by

Figure 1. Four Possible Scenarios for AI Middle Powers



finding a chokepoint in the AI supply chain. For others, faster adoption will bring the biggest returns. Still others are betting on new technological directions entirely. Many have yet to make serious moves of any kind.

Two variables should drive middle power choices: the speed of frontier progress and what we call the “frontier premium.”

Even in the face of uncertainty, governments can find strategies that remain robust across multiple scenarios. In particular, finding sources of leverage at supply chain chokepoints is valuable in almost every technological future—and most valuable in scenarios where AI poses the largest risks to middle powers’ economic and political standing. Where the right choice isn’t clear, policymakers can preserve optionality, gather information, and hedge against the most dangerous futures. Getting big bets right

can pay off handsomely: If Taiwan had not invested in semiconductors in the early 1980s, or China had not built out its infrastructure for electric vehicles in the 2010s, each would be worse off today. Countries that wait for total clarity may find it too late to take advantage of the world that emerges.

**Scenario 1: Chokepoint World
(fast progress, large gap)**

This is the world of the artificial general intelligence bulls: One company or a small group develops advanced AI systems that are not easily copied or recreated. Access to large quantities of compute drives dramatic performance gains, and labs figure out how to bar techniques like [distillation](#), which involves training a less capable model on the outputs of a stronger one, and secure their algorithmic and data insights from competitors. The

Figure 2. Middle Powers Have Invested Across the AI Stack



Source: Carnegie Endowment for International Peace analysis.

Note: Flagship initiatives are highlighted; most countries listed have pursued policy efforts across the AI stack. Figures are converted into USD based on May 2026 rates.

United States maintains a lead in chip production, in part by continuing to enforce its export control regime. As a result, the gap between U.S. and Chinese systems grows rapidly.

In this scenario, the United States has both the incentives and the capacity to restrict access to the best systems. The beginnings of such a regime are already visible: In June 2026, the U.S. government revoked foreign access to Claude Mythos and Fable and barred OpenAI from publicly releasing GPT-5.6. (It later removed both restrictions.) If AI progress accelerates, restrictions will become more aggressive. Companies will develop models with cyber and military capabilities far exceeding those of [Mythos](#), and national security agencies will push to limit distribution. U.S. allies and companies may lobby for a tiered access model, in particular for government agencies, but public access to the most powerful systems will be sharply restricted. The world will [divide](#) into AI haves and have-nots.

For some middle powers, the best bet may be to find a position in the frontier AI supply chain that can be leveraged to ensure continued access, or some other point of interdependence that will make the United States reluctant to cut it off. Theoretical leverage will not be enough; governments will also need the political will to use it.

Finding a true chokepoint—like TSMC and ASML, key companies in the semiconductor supply chain—is a heavy lift. But there are other ways for countries to gain influential positions. If the AI hardware landscape becomes more [varied](#), for example, specialized chip firms may win by betting on a particular type of model or workload. The United Kingdom has embraced this approach, attempting to retain promising [companies](#) by [offering](#) public funding and expedited visa processing for employees. Or countries may be able to specialize elsewhere, such as in advanced testing and packaging (where chip firms compete heavily on cost) or semiconductor manufacturing equipment (an area of success for Malaysia, where tool maker Lam Research [opened](#) its largest global plant in 2021). Countries may

also develop expertise in AI safety and evaluation, as the United Kingdom has done, or band together to offer market access, a potential strategy for the EU.

Theoretical leverage will not be enough; governments will also need the political will to use it.

Compute offers another possibility. Countries with readily available energy and fast-track permitting processes may be able to attract large data center investments; countries that serve as regional or global [compute](#) hubs will be in a stronger negotiating position relative to both the U.S. AI industry and the U.S. government. American firms are eager to build new compute capacity, and countries that can provide it will be able to negotiate for concessions, from tax revenue to frontier model access.

In this world, countries that control critical chokepoints in the AI supply chain hold the strongest cards. [Germany](#), [Japan](#), [the Netherlands](#), [South Korea](#), and [Taiwan](#) are the obvious examples today. Those that can successfully develop new footholds may be able to bargain for access, though likely with less leverage than current players. Compute hosts, possibly including Australia, France, India, Malaysia, the Nordics, and the UAE, would hold leverage roughly proportional to the share of global compute they control. Important U.S. security partners, like the other Five Eyes countries, would also have a claim based on the broader value of their contributions to U.S. national security policy.

Scenario 2: Transform or Die (fast progress, small gap)

In this scenario, highly advanced AI exists, but can be copied or developed relatively easily by countries with access to compute and talent. The United States and China are locked in a close race, and the task for middle powers is to help preserve stability, build resilience,

and promote productive adoption. A large edge in AI capabilities could be hard to defend, giving rise to this situation, if techniques like [distillation](#) continue to work on powerful models, security [lapses](#) allow the theft of model weights, or new techniques reduce the importance of additional computing power. Many observers contemplated such a scenario in early 2025, after the Chinese company DeepSeek [released](#) its R1 model, which some claimed reflected greater training and processing efficiencies than contemporary U.S. models. Such concerns have resurfaced more recently in the wake of Kimi K3, an efficient and highly capable open-weight model [published](#) in July 2026.

In this scenario, the bottleneck is not access to AI but the speed at which institutions can adopt it. AI developers may struggle to capture returns, as profits are competed away, and strategic and economic value will flow to other bottlenecks: militaries with the strongest hardware and AI orchestration, compute providers capable of turning energy into tokens, and businesses that best convert tokens into economic value. That would make adoption capacity—the ability to reorganize institutions and workflows around abundant AI—critical.

To that end, countries may strive to attract highly skilled immigrants and boost top educational programs. Risk-taking entrepreneurs, power users with domain expertise, and skilled engineers will be critical to turning cheap, capable models into productivity gains. And even if powerful AI models are broadly available, compute may not be. Governments may want to provide public computing infrastructure, perhaps with [subsidized](#) access for researchers or startups, since early and widespread use will build the knowledge to adopt AI effectively. Just as the internet made Amazon possible and the smartphone created Instagram, powerful and broadly available AI will generate new business models and powerful companies; the countries that enable them will have a significant advantage.

Flexibility will come with difficult tradeoffs, however. The same lax regulations that could facilitate growth

may also make it more difficult to block AI-enabled cyberattacks and manipulation campaigns. Disruption will drive political backlash. Permissive rules for the use of AI in government and critical infrastructure could embed unsafe technologies in essential systems, and inadequate social safety nets could lead to runaway inequality.

The same lax regulations that could facilitate growth may also make it more difficult to block AI-enabled cyberattacks.

Countries may also lean into the relational or the physical economy. The Google economist Alex Imas has [argued](#) that in a world where powerful AI systems can produce many goods cheaply, the gains will flow toward high-touch service sectors like arts, care, education, hospitality, therapy, and craftsmanship. These areas are harder to automate because human involvement is a fundamental part of the product. Countries that have existing advantages in these sectors may choose to lean in further, particularly where the products are exportable—handcrafted goods, for example, rather than live theater.

Countries will also likely want to build industrial capacity along the AI supply chain, producing actuators for humanoid [robots](#), for example, or semiconductor manufacturing equipment. Other physical industries may provide advantages, particularly if jobs require dexterity, benefit from tacit process knowledge, and need costly hardware to automate. And even in largely automated sectors, countries can reap economic benefits if they rely on scarce natural resources or can develop an early lead.

China's vast physical production capacity and ability to turn economic forces at central direction may position it well. Fast followers rich with talent, like Canada, Singapore, and the United Kingdom, could prove adaptable, although service-heavy economies might see their current profit centers eroded. Industrial powers like

Germany, Japan, and South Korea would benefit from their manufacturing bases. Countries with economies heavily reliant on low value-added services—India and the Philippines, for example—could be badly exposed as AI replicates the skills on which their export sectors depend.

Scenario 3: Adopt for Advantage (slow progress, small gap)

In this scenario, AI is useful but its capabilities remain [jagged](#). The gains seen in coding and mathematics—where solutions are verifiable and there is less uncoded process knowledge—fail to replicate in other sectors, and coding tools do not themselves reshape the economy. World [models](#) prove difficult, AI systems struggle to learn from real-world experience, and capabilities remain jagged and domain-specific. As the researcher Tom Reed [puts it](#), “there is no GitHub for closing a Series B, winning a knife fight, or negotiating with Houthis.”

If progress slows, governments and companies will have less reason to lock down the most advanced systems—which, in this scenario, won’t be much better than the second-best. AI company valuations might collapse if investors decide that transformative results will take longer than expected, and frontier AI model developers would likely operate on thin margins. Gains would instead flow to sectors where AI tools are most directly productive, such as finance and consumer tech, and to the companies doing the hands-on work of integrating AI into industries like law and healthcare, building bespoke tools, and reshaping internal systems around new AI products.

This is the world in which sovereign AI efforts like open models fine-tuned for expertise in local languages are most likely to succeed. Second-best and open-source models would follow closely behind the frontier; customization and integration will be the hard part. Efforts to create national databases and develop models

trained on local language data might add significant value, and national AI champions could compete effectively in particular markets by building on open models through fine-tuning, model harnesses, and other custom infrastructure. Sovereign programs could offer greater security and control, help spark AI startup ecosystems, and improve leverage in negotiations with American providers if they can offer a realistic alternative. Yet precisely because progress would have slowed in this scenario, pursuing sovereign AI would not be an overriding national priority.

If progress slows, governments and companies will have less reason to lock down the most advanced systems.

Governments would still want to promote adoption, but doing so would be less of an imperative than in worlds where progress is rapid. The technology would resemble other general-purpose technologies, where adoption boosts growth and could gradually reshape the world order, but without necessarily creating a narrow window for decisive government action. India’s huge IT services sector and domestic market, as well as its large pool of engineers, would position it well in this world. China’s [ability](#) to mobilize to promote mass AI diffusion could also pay off, though markets might [beat](#) mandates in promoting adoption. Moving slowly—Europe’s great sin—would be far less costly here than in a rapid-progress world.

Scenario 4: Tiered Access (slow progress, large gap)

In this scenario, AI progress plateaus, but a few clear leaders maintain a durable edge. This could happen if the strongest talent concentrates in one or two companies, or the frontier labs collect so much data from implementation that they can parlay [early leads](#) into a lasting advantage. [Partnerships](#) with other industries

may also help build moats. Or producing and running frontier models may require so much computing power that only a handful of companies in rich countries can afford to build and access them. Frontier developers and the U.S. government would likely restrict access to the best systems, but not as aggressively as they would in a world with faster progress, because the range of dangerous capabilities (in areas like cyber and military operations) would be narrower. Both frontier developers and skilled AI users would capture economic value.

In these circumstances, adoption and leverage to bargain for access may both help countries secure their economic and national security positions, though less urgently than in other scenarios. One significant difference between this scenario and those in which the frontier premium shrinks, though, is that the most ambitious sovereign AI model efforts that involve attempts to build independent frontier-model capabilities from the ground up would be particularly inadvisable. If well-funded American and Chinese companies struggle to catch up to the leaders, international competitors will find doing so even harder, perhaps wasting public funds and technical talent in the process.

Middle-income countries with large technology and service sectors, like India, might be better positioned here than in faster-moving scenarios. They are likely not sufficiently close to the United States to gain full access in a fast-moving world, but if progress were slower and restrictions less widespread, their scale would let them play an important role in providing AI services like model customization, systems integration, and AI-driven process transformation. Countries without supply chain leverage, large domestic markets, or strong adoption capacity might find themselves in a slow squeeze, falling steadily behind in relative terms.

If well-funded American and Chinese companies struggle to catch up to the leaders, international competitors will find doing so even harder.

Implications for Middle Powers

If the right strategy depends on which scenario emerges, what should countries do now, while the future remains unclear? Countries enter the AI transition in a wide range of economic and strategic positions (see figure 3). Taiwan's problem is how to convert hardware leverage into software access. India's is how to capture implementation upside while managing exposure in tradable services. The UAE's is how to convert its capital and energy into durable strategic relevance. But some principles apply broadly, and governments should not wait for one of these worlds to fully materialize before they act.

Figure 3. The Six Types of AI Middle Powers by Current Economic Position

TYPE	EXAMPLES
 Semiconductor chokepoint powers	Japan, Netherlands, South Korea, Taiwan
 Energy-rich compute hosts	Canada, France, Malaysia, Nordics, UAE
 Physical and industrial powers	Germany, India, Japan, Mexico, South Korea
 Talent- and software-rich fast followers	Canada, India, Israel, Singapore, Sweden, UK
 Regulated-sector incumbents	France, Germany, Japan, Switzerland
 AI-exposed services exporters	Egypt, India, Latin America, Philippines, Vietnam

Note: The categories are not mutually exclusive.


Certainty won't be possible, but there are some indicators to watch for. If distillation continues to help competitors quickly mimic the performance of frontier models, that will point toward a small-gap world. If unit economics at the top firms remain strong, that will suggest the opposite. If the United States formalizes its nascent ad hoc access regime, that could indicate it expects a large-gap, restriction-prone future. And if major labs begin reporting diminishing performance returns on new training runs, spending more compute for smaller capability gains, that might suggest progress is slowing. These indicators will not point cleanly to a single one of the four scenarios outlined here, but they can help governments figure out how urgently they need to move.

Building leverage with supply chain chokepoints or scarce resources is valuable in most scenarios—and most valuable in the worlds where middle powers have the most to lose. If AI progress continues apace and the frontier pulls away from cheap alternatives, middle

powers risk watching as AI winners race ahead. It is in this scenario that cornering the market on part of the supply chain will pay the highest dividends. In scenarios where progress is slow, or the frontier is easy to replicate, advantage in AI will be both easier to obtain and less significant. Countries will have more time to figure out their strategies, and mistakes today will be less likely to harden into permanence.

Promoting adoption capacity also pays off in most scenarios, though it is far harder for governments to do well. Industrial policy is often ill-suited to encouraging businesses to use new technologies, and centralized programs risk subsidizing unproductive or inevitable applications. A few levers, however, make general sense. High-skilled immigration, investments in education, flexible labor markets, and openness to startups create conditions in which new organizational structures can emerge that take advantage of AI without requiring governments to predict which applications will matter.

Figure 4. Leverage and Adoption Are Safer Strategies Than Sovereign AI



In scenarios where frontier access matters most, attempts to replicate the frontier through sovereign model development will be the hardest to pull off. If the leaders open up a large gap with second-best alternatives, developing powerful sovereign options will be difficult and expensive, perhaps [costing](#) as much as \$900 billion, largely because of compute and talent needs.

The track record of such moonshots is not promising. A few years ago, the UAE hoped that its national champions would rival the best U.S. offerings. In 2024, presidential adviser Faisal Al Bannai [declared](#) that the country's new Falcon model showed the UAE could "compete with the best players globally." But as the AI race has accelerated, Emirati companies have lagged. Sovereign efforts are most likely to succeed in scenarios where frontier progress is slow, and the gap between the best and second-best models is small—precisely the worlds where frontier access will be less important. More limited efforts to build domestic models—local language models, small models designed to fit specific use cases, and open-source platforms on which local developers can build—may attract talent and make adoption easier. But none of them will solve the economic and strategic problems that stem from a lack of frontier access.

The best path across all scenarios demands smart engagement, not autarky.

Establishing a stronger position at supply-chain chokepoints may help middle powers secure that access, but once a country has accumulated leverage, it needs the will to use it. This requires an ability to stomach pain. For example, if the Netherlands' frontier access is cut off, the country could threaten to restrict access to ASML's advanced lithography machines, but actually doing so would risk economic damage and harsh retaliation. Countries may do better by banding together (although the track record of the EU is mixed on this front) and

focusing on their highest-priority asks—ones with enough domestic support that their political systems can endure a sustained effort.

The best path across all scenarios demands smart engagement, not autarky. Adopting cutting-edge AI, innovating on top of it, and carving out a place in the value chain all involve various forms of industrial learning. The way to accomplish that learning is through partnerships that disseminate know-how from firms on the bleeding edge, not by trying to recreate it from scratch. This is the playbook that successful companies have followed in the past. South Korea's Hyundai didn't attempt "sovereign" auto manufacturing; it licensed access to engines from Japan's Mitsubishi, visited as many international car makers as it could, and embraced foreign consultants. The playbook for AI is likely to be similar.

Middle powers do not control the rate of progress or the size of the frontier gap. But there is plenty they do control, and the costs of waiting for certainty are high.

About the Authors

Alasdair Phillips-Robins was a fellow in the Technology and International Affairs Program at the Carnegie Endowment for International Peace.

Teddy Tawil was a visiting researcher in the Technology and International Affairs Program at the Carnegie Endowment for International Peace.

Sam Winter-Levy is a senior fellow in the Technology and International Affairs Program at the Carnegie Endowment for International Peace.

For complete source notes, please read this article at [CarnegieEndowment.org](https://www.carnegieendowment.org).