CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

SEPTEMBER 2026

AI Breakout Capacity: A Middle-Power Path Between Racing and Dependence

Anton Leicht
Sam Winter-Levy

Technology and International Affairs Program

AI Breakout Capacity: A Middle-Power Path Between Racing and Dependence

Anton Leicht
Sam Winter-Levy

September 2026

About the Carnegie Endowment for International Peace

In a complex, changing, and increasingly contested world, the Carnegie Endowment generates strategic ideas, supports diplomacy, and trains the next generation of international scholar-practitioners to help countries and institutions take on the most difficult global problems and advance peace. With a global network of more than 200 scholars across twenty countries, Carnegie is renowned for its independent analysis of major global problems and understanding of regional contexts.

Carnegie Endowment for International Peace
Publications Department
1779 Massachusetts Avenue NW
Washington, DC 20036
P: + 1 202 483 7600
F: + 1 202 483 1840
CarnegieEndowment.org

This publication can be downloaded at no cost at CarnegieEndowment.org.

TABLE OF CONTENTS

About the Authors

Anton Leicht is a fellow with the Technology and International Affairs Program at the Carnegie Endowment for International Peace, where he researches the political economy of artificial intelligence.

Sam Winter-Levy is a senior fellow in the Technology and International Affairs Program at the Carnegie Endowment for International Peace, where his research covers emerging technology and national security, with a focus on the geopolitics of AI.

Technology and International Affairs Program

The Technology and International Affairs Program develops insights to address the governance challenges and large-scale risks of new technologies. Our experts identify actionable best practices and incentives for industry and government leaders on artificial intelligence, cyber threats, cloud security, countering influence operations, reducing the risk of biotechnologies, and ensuring global digital inclusion.

Introduction

Access to frontier artificial intelligence is becoming scarce. Three forces are converging. First, security concerns are pushing developers to only give handpicked companies access to the newest models, as with Anthropic's restricted rollout of its leading cybersecurity model, Mythos. Second, compute shortages are forcing developers to limit customer usage. And third, the U.S. government has begun to grasp for control over how frontier systems reach the market, slapping export controls on Anthropic. Access to frontier AI may soon be rationed, conditioned, and switched on or off at the discretion of a handful of American firms and the government that oversees them.

For the world's middle powers—countries such as Germany, India, and the United Kingdom, which have substantial economic and military capacity but no frontier models of their own—these trends are turning what was once a theoretical vulnerability into a more immediate problem. Their access to advanced AI systems depends on the whims of policymakers in Washington and Beijing, they remain exposed to AI's disruptive effects whether or not they share in its benefits, and they have little leverage over the terms that will shape AI's rollout.

The options available to them are unsatisfying. Sovereign moonshots are unaffordable for middle powers individually and face daunting political obstacles collectively. Adoption strategies rely on the gap between frontier models and fast followers remaining narrow, and are often cover for not doing much of anything while hoping that AI progress stalls. Bargaining with great power–based AI companies for access cannot fully insure against a decision to revoke guarantees in the event of a subsequent breakdown in relations. As we have written at length elsewhere, middle powers' best strategy is still to negotiate hard for durable access to imported frontier systems, drawing on whatever leverage they hold in the AI supply chain to help enforce those deals. But that strategy needs a fallback.

This paper examines one option: constructing and maintaining the capacity to build frontier AI quickly should the need arise—all while remaining closely aligned with a great power patron. Call it an **AI breakout posture**. This paper explains the strategy, its logic, and the policies required to implement it. A breakout posture will not guarantee frontier capability, but it provides a hedge against the worst technological futures. If well executed, this plan can strengthen middle powers' hand at the bargaining table. U.S. policymakers, for their part, should understand the options that even their closest allies may be tempted to pursue if they conclude that they cannot trust the United States as a reliable supplier of frontier AI.

Incomplete Answers

Middle powers are increasingly waking up to their vulnerability to U.S. and Chinese control over frontier AI. In response, the natural temptation is a sovereign moonshot: an attempt to close the gap by building frontier systems domestically, alone or in a bloc with other middle powers. France has signaled its intent to do this with Mistral, the strongest model developer outside the United States and China. The European Union is funding "gigafactory" computing hubs, touted as Europe's response to the U.S. data center build-out. A Franco-German initiative aims to channel investment toward frontier research.

But strategically relevant capabilities are currently out of reach for even an alliance of middle powers. They lag behind on every relevant input. Mistral raised funds at a valuation of roughly $14 billion in September 2025; the leading American labs raise multiples of that in a single funding round and spend more than that on compute alone in a year. Just one American company, Google, now controls roughly a quarter of the world's computing power. When the EU's planned gigafactories come online, they will at best match what U.S. companies built years ago. A joint effort by an alliance of middle powers would close part of this gap, but the coordination problems are daunting: If a single government struggles to move quickly on a project like this, a bloc of them will struggle more.

Right now, the politics of a frontier moonshot will be fraught. It will require tens of billions of dollars in upfront commitments and large allocations of energy and grid capacity—with no guarantee of positive returns. Legacy industries will argue that cheap power and capital should flow to them rather than to a speculative bet on an unpopular technology. Publics divided about the value of AI will ask why governments are prioritizing data centers over hospitals. These pressures are not insurmountable. But they are already building in the United States, where the AI build-out enjoys more entrenched political and commercial backing. They will weigh even heavier in middle powers, making a sustained industrial program hard to defend in ordinary political circumstances.

> **Compute asymmetries between leading U.S. developers and the rest of the world are widening, and the gap may be locked in for at least the next several years.**

Faced with the difficulties of building a serious sovereign capability, policymakers may be tempted to give up on the frontier entirely and focus instead on fast following and adopting sub-frontier technology as rapidly as possible. The intuition has some merit: For everything from economic growth to military power, diffusion will matter at least as much as invention. But a pure adoption strategy bets on the gap between the most capable systems and more widely available ones staying small enough to live with.

Until recently, that has been a reasonable bet. Fast followers have been able to track the leading edge at a manageable lag—six to nine months in many cases—by exploiting common research techniques and "distilling" from leading models (using the outputs from a larger model to train a smaller one). But the conditions that made fast following relatively easy are eroding. Compute asymmetries between leading U.S. developers and the rest of the world are widening, and the gap may be locked in for at least the next several years considering the lead times on data center build-outs and chip deployments. U.S. companies, with support from the U.S. government, are cracking down on distillation, through both technical reforms and changes in usage policies. Leading developers' revenue advantages may compound, especially if the labs' dreams of automated AI research and development (R&D) come true.

Middle powers can do many things with sub-frontier models. But for tasks where relative capability against a competitor or adversary matters—cybersecurity defense against attackers using frontier systems, drug discovery and other scientific R&D where capability gaps translate into market position, or military applications where rivals are deploying systems against each other—frontier access will be decisive. Middle powers will not need frontier AI for every application, but they will need it for some economically and strategically vital priorities. Project Glasswing, Anthropic's managed access policy for its latest model, illustrates the point: European governments and businesses anticipated serious vulnerabilities if they did not receive access, and no European fast follower or second-best model could assuage those concerns. Any sovereignty strategy needs to provide for some amount of frontier access.[1]

If neither building sovereign frontier models nor adopting significantly lagging sub-frontier ones will suffice, the third option is to bargain hard for continued access to foreign frontier systems. Middle powers can use assets the great powers need, such as critical minerals, data sources, and advanced manufacturing capabilities, to extract durable guarantees of frontier access. We have argued elsewhere that versions of this strategy could reduce the risk of being cut off, by tying great power providers into partnerships they are incentivized to maintain. And middle powers can supplement these arrangements with scaffolding layers that allow them to shift workloads between providers if terms deteriorate. But even a well-executed import policy cannot protect against every risk. If a great power decides—for security reasons, trade leverage, or political caprice—to deny a country access to imported frontier capability, no contract written in advance can reliably prevent it. Middle powers may need a backup plan.

The Breakout Posture

The most robust stance is for middle powers to retain the *option* to build advanced AI systems without committing to that project today. Guy Ward-Jackson and Keegan McBride have drawn a parallel to nuclear weapons in a piece arguing for "AI latency." For example, Japan has chosen not to bear the political, diplomatic, and financial costs of developing nuclear weapons. But with U.S. tolerance, it retains the ability to build them quickly should the need arise, by holding sufficient reserves of the necessary equipment, material, and research expertise.

This analogy is useful, but their version of this approach—a fallback capable of producing a "good enough" sub-frontier model in a crisis—may not go far enough. A breakout capability may need to aim at a full catch-up sprint to the frontier itself. Their proposal is well-suited to a scenario in which access to frontier systems is disrupted but the gap between frontier and sub-frontier capability remains sufficiently narrow that a good enough model is a tolerable substitute. In that scenario, a sub-frontier fallback does little harm, but the underlying problem it's meant to solve is relatively mild. The harder scenario, where a fallback will really matter, is the one in which frontier capability is both scarce and strategically essential. Against that contingency, a sub-frontier fallback offers little insurance.

Under this posture, middle powers' primary plan would still be to seek favorable conditions for importing foreign AI systems and to capture value and leverage elsewhere in the supply chain. But should technological, diplomatic, or economic dynamics deteriorate, they could choose to enter the race at once. In all likelihood, such a push would require a concerted effort of middle powers' private and public sectors: They would need to consolidate all their latent resources from compute to talent to supply chain leverage to conduct one singular project to build a frontier AI system.

Retaining the option of such a sprint has four distinct advantages over attempting it immediately.

> **Retaining the option of an AI sprint has four distinct advantages over attempting it immediately.**

The first is financial: The economic future of AI remains uncertain. Some analysts see risks of a financial bubble, while others expect value to accrue to smaller or specialized models, or to the ecosystems that apply them, rather than to large frontier projects. The breakout project serves as a hedge against these futures. The investment required for frontier racing is only necessary once the market structure and its strategic implications have become clearer,

which avoids overcommitting middle powers to one specific play too early. If the market develops away from the primacy of frontier developers, middle powers can simply choose not to pursue the breakout, save substantial resources, and place different strategic bets instead.

The second is political: Electorates and policymakers in many middle powers do not agree on how transformative advanced AI technology will be. The costs of a frontier project create strong political constituencies in opposition, including legacy industries that would rather they get favorable treatment on energy prices, local technology companies that argue any sovereignty bet should be placed on them instead, and so on. It is not clear that a frontier project would survive these pressures. But if frontier capability turns out to matter greatly, this opposition will be harder and harder to sustain as the importance of assured capabilities becomes more and more visible. It will be harder to recruit a nation's economy into a speculative years-long project under ordinary conditions than into a concerted months-long sprint under high urgency.

The third is geopolitical: For most middle powers, at least one great power has an interest in keeping them deeply integrated into its AI stack. Right now, given most middle powers' lack of leverage, a project to inch toward technological independence is vulnerable to interference, for example through restrictions on the export of computing hardware. But middle powers are already working to reduce these dependencies; they may be in a better position to resist that pressure when they later commit to a breakout project. They can also take specific measures to insulate the project against future pressure now, while great powers are not yet tracking these efforts as closely, for instance by securing enforceable access to the relevant hardware.

The fourth is epistemic: The technical future of AI remains uncertain. Governments know, in the abstract, that something called a frontier system will likely be relevant to geopolitical and economic competition. But they do not yet know what the difficult-to-replace features of that system will be. Would a catch-up project have to heavily focus on reinforcement learning aimed at software engineering capability in order to bootstrap into even more advanced systems? Or would it have to take the shape of one massive pre-training run at the frontier? Would success turn on building a highly efficient model and leveraging most available compute at inference time, or on racing to the highest-performance model at all costs? The American and Chinese ecosystems can afford not to answer these questions because they have many live players pursuing different versions of these strategies in competition with each other. But a middle-power project would likely get one shot only—even an alliance would not have enough resources to support *multiple* frontier players. A sovereignty project starting today would have to commit to a specific view of the technological future while the exact nature of strategically vital systems remains uncertain; a breakout play could delay that commitment until the picture is clearer.

Figure 1. The Cost of Keeping An AI Breakout Option

A delayed sprint may cost more overall, but commits major resources when more is known.

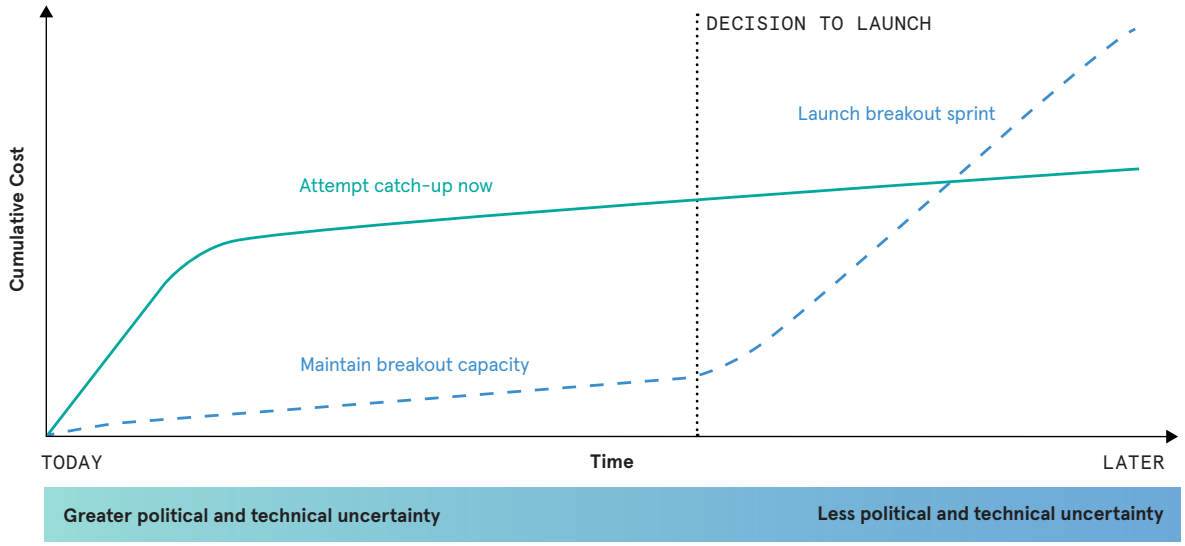

Source: Authors' illustration.
Note: Illustrative trajectories, not cost estimates. Catch-up is not guaranteed under either strategy.

Method

Any middle-power breakout posture would proceed in two phases. The first phase, addressed in this section, is to create the preconditions for a rapid sprint to the frontier: latent access to compute, private-sector capacity that can be redirected on short notice, and the policy instruments to mobilize the initiative. The second phase is the sprint itself. At the point when a government sees the risks of dependence or lost access as intolerable, it would consolidate these resources behind a champion firm or cluster of firms, supplying funding, talent, compute, and energy to pursue a frontier-level capability according to specifications set by the procuring governments and firms. The specifics would inevitably vary depending on the precise technological paradigm in play at the time. For now, the bulk of the policy work is in phase one.

Compute

The first precondition is *latent* access to a large amount of computational power. Latent access implies that a country can funnel compute into a breakout project on demand through guaranteed purchase options, government ownership clauses, or immediate

ownership. It is distinct from contingent compute that is held by foreign entities or outside its own borders and to which access could be denied, but broader than purely state-owned compute. The total volume should be large enough to start a frontier-scale AI development project and then to service the resulting model for critical applications. It should also be networked and configured for training, not only for inference.

How much compute would be necessary to secure a latent breakout capability? It remains unclear, and subject to paradigm changes. For instance, in recent years, relative compute demand in frontier model performance has shifted from singular compute concentrations required for large pre-training runs to capacity required for reinforcement learning post-training and large-scale inference provision at the point of use. Still, most trends suggest that only a small percentage of overall global deployed compute is necessary to enable frontier training and selected inference. Most AI-related compute is spent on inference, research, and post-training, with frontier training runs only making up a fairly small proportion. At the moment, global compute is distributed across AI developers, and while there is little current data available, researchers estimate that less than 10 percent of OpenAI's compute in 2024 was spent on actual frontier training. If frontier training makes up less than 10 percent of even a leading lab's compute budget, and those labs don't, for now, dominate AI compute usage, the realistic compute demand for a breakout training run is likely in the low single digits as a proportion of global capacity—an achievable goal for some middle powers or an alliance of them.

After all, much less compute is required to retain capacity for a dash to the frontier once its technical shape becomes clearer than to sustain an enduring stake in the AI race, which includes competing on inference pricing, funding speculative research, and so on. The result would never be a wholly competitive, fully integrated frontier developer, but a viable source of a frontier-level AI capability—perhaps not with sufficient inference capacity to service an entire economy, perhaps not profitably, but enough to close a critical capability gap if it emerged.

That said, in the absence of concerted action, even the single-digit percentage of global compute that would realistically be required is still far beyond the projected future capacity of even a middle-power alliance. To secure it, a middle power pursuing this strategy would likely need a portfolio of policies aimed at different access modalities.

First and most obviously, countries can choose to purchase and deploy some of that compute themselves. This is the most reliable, yet most expensive option. Government-funded infrastructure build-outs require countries to operate in a space currently occupied by the most well-capitalized firms in history, and to make investment decisions that have to compete in a private marketplace. This has historically not been a strength of allied democracies, and so direct government procurement mostly seems attractive where other options to secure required compute fail, such as onshoring compute that market incentives alone will not support (for example, high-security capacity for national security purposes).

Second, countries can incentivize their own private sectors to build out compute today. This build-out could be pursued by domestic AI developers looking for training capacity, by domestic firms with high computing needs seeking to run inference, or by neoclouds hoping to rent out compute. Domestic private build-out can be a valuable part of an allied compute stack, but it is likely to fall short of the required scale. Private sector capital in most middle powers pales in comparison to U.S. capital and is often more hesitant to take bets on AI infrastructure.

Third, countries can attract foreign private companies—neoclouds and hyperscalers—to build out infrastructure on their shores, especially in the context of compute-for-access deals in which U.S. AI developers promise to offer a country the same models as America in exchange for data centers. In many cases, that interest will be organic and only require some favorable licensing and land use, because the private sector is short of attractive locations for the compute build-outs required to keep pace with rapidly rising demand. In other cases, to accelerate build-out timelines and make hosting in middle powers more attractive, concessions may be required, such as energy price subsidies or guarantees, financial backstops provided by the state or private sector, or deal structures with governments or private sector firms that guarantee local demand or grant the operator access to valuable data.

There are other reasons to attract these build-outs, but for a breakout capability they are particularly attractive. If this strategy is executed well, middle powers can get capital-rich foreign firms to finance the infrastructure that ultimately fuels their breakout capacity. Middle powers pursuing this route would have strong incentives not simply to accept or attract that investment passively, but to negotiate terms that give them latent access to this compute. Possible measures along these lines might include laws in the shape of the U.S. Defense Production Act, which gives the executive the power to reorient industrial capacity toward specified purposes under certain extreme conditions; more general legal powers of expropriation against above-market compensation; and narrow expropriation conditions baked into specific data center deals that only trigger if the data center operators revoke frontier access to their host countries. Such measures risk making a middle-power country a less attractive location to host data centers, so they would have to be paired with greater build-out incentives.

Fourth, countries can attempt to secure access to further chip purchasing ability. A breakout project would require that countries purchase additional compute once the sprint is underway, hoping to deploy it in time to run the model they train on consolidated compute. But for that to work, they would have to retain access to compute purchasing options; it would be fatal if export controls or supply chain bottlenecks blocked the onshoring of the remaining necessary compute when it was most needed. For instance, private firms might enter long-term supply agreements with Nvidia and TSMC—if they need to sprint, they execute them; if not, they buy and resell at a moderate loss. Governments could underwrite these supply agreements to make them attractive to the suppliers wherever possible. To further

Figure 2. Two Ways to Close the Compute Gap

A breakout posture can shift much of the upfront buildout cost to hyperscalers.

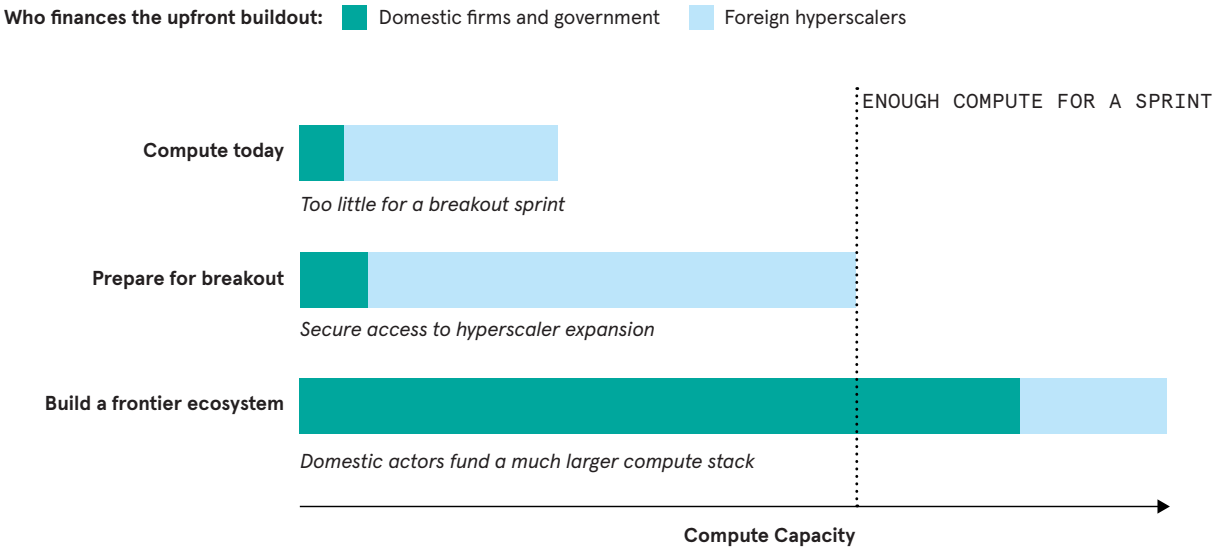

Source: Authors' illustration.
Note: Conceptual illustration of the strategy; bar lengths and funding shares are not estimates.

encourage these chip suppliers to enter deals, middle powers with positions in the semiconductor supply chain could leverage them in exchange. They could guarantee exclusive access and production expansion to spec to TSMC or Nvidia in exchange for such long-term options and purchase agreements.

Private Sector Capacity

The second condition for a successful breakout is private sector capacity—technical and organizational capability—that a state can redirect into a frontier project on short notice. That requires several core inputs.

The first is talent. A breakout project will need to draw on a critical mass of researchers and engineers who are capable of working on frontier-relevant problems. This entails retaining existing talent—through visa regimes, favorable tax treatment, an attractive R&D environment, and prestige employers willing to host sensitive workloads such as DeepMind in London or Anthropic and OpenAI's expanding European offices—and creating pipelines for new talent. A strategy oriented purely toward talent retention will struggle to hold off

the world's richest companies as they expand their hiring sprees; a strategy oriented purely toward education and development will produce talent with nowhere domestic to land. Governments invested in creating breakout capability need to simultaneously create a talent pipeline and a place for that talent to work, even if it is with foreign firms that need incentivizing to open offices within middle powers' borders.

The second is the operational capacity to run a frontier project. The required asset here is a small set of firms that have pursued frontier-adjacent work: at least some pre-training, comprehensive post-training, and integration of larger compute clusters. These firms also need to have working relationships with both domestic governments and domestic private sectors, so that an eventual project won't start from zero. A breakout project cannot be staffed by spinning up a new company in week one.

As they build that capacity, middle powers would be wise not to adopt a maximally adversarial posture toward the U.S. tech sector. There is a version of the fast-follower playbook that relies on commercial distillation of American frontier models. Through querying frontier models at scale, distillers can reproduce more efficient, less powerful versions of these models. Industry reports indicate that this technique is at least partly responsible for the success of Chinese open-source providers, and some analysts have suggested that middle powers should also resort to distillation to retain a fast-following capacity. But that suggestion mistakes current American disinterest in cracking down on distillation for a general inability to do so. In fact, American developers and the U.S. government could introduce much more restrictive rules for API access if they felt seriously threatened by distillation. If they do, this will have untenable consequences for middle-power developers. Middle powers will be able to execute neither a durable fast-following strategy nor, ultimately, the kind of chip import access required to pull off a breakout play if they have antagonized the United States while their frontier ambitions still remain vulnerable.

Today, the French company Mistral and the Emirati state-backed champion G42 are probably the only middle-power firms with the organizational capacity to even attempt a breakout; Canadian company Cohere and Germany's Aleph Alpha could be developed into candidate options. In the short term, governments can create and sustain that capacity by having these firms operate as fast followers by design. For now, the governments that currently keep them on life support would task them to stay behind the frontier, producing fallback models at a time lag of about nine months. Governments can provide a safe market for these models in, for example, public administration use cases. Counterintuitively, middle-power governments should *not* prematurely employ these companies for frontier development itself; that would overextend political resources at a time when they are unlikely to yield a frontier system, in effect misappropriating resources and exposing fast followers to political backlash for absorbing public spending without commensurate results.

A latent breakout project would also require standing of a kind that could attract political support from, if not deep integration with, multiple governments. A breakout firm that is enmeshed with only one government—preferentially financed, contractually entangled,

regulated as a national champion in a single capital—would be politically embattled in a coalition of the size required to fund and push ahead a breakout project. Mistral would run into this challenge. Its close relationship with the French state means that without credible commitments and complex negotiations over its governance, few other states would be willing to pour billions of dollars into it as the vessel for their breakout ambitions.

The final necessary element, and most concerning to the prospect of a delayed push from a breakout posture, is the availability of sufficiently advanced AI systems to assist AI development. Today, much of the software engineering at frontier AI developers is already carried out by their own internal versions of coding agents like Anthropic's Claude Code or OpenAI's Codex. The relative ability of these agents is likely to become an increasingly important factor, to the extent that today's developers are also restricting each other's access—for example, by curtailing business contracts with competing developers, or through outright restrictions on the extent to which their own models can be used for frontier AI development. Should that trend continue, a breakout project might soon find itself unable to catch up simply by virtue of lacking the software engineering AI capability to do so. The availability of coding agents is therefore a key input into when to launch a breakout project. Should it become clear that lack of access to advanced AI systems would move from being a mild handicap to a prohibitive restriction on a middle-power frontier push, a government would face a stark choice between abandoning the breakout sprint or commencing it immediately.

State Capacity

Access to compute and a well-resourced private sector are necessary preconditions for a breakout project. But any such effort will be deeply entangled with the state, and so state capacity to launch and sustain a breakout project is a final precondition for success. This involves three distinct requirements: legal authorities that allow rapid action without months of new legislation; an apparatus for deciding when to launch; and the ability to sustain what may be a multi-year commitment across electoral cycles and in the face of potentially acute geopolitical pressure.

The first requirement is legal authorities that are available to governments in advance of a development sprint. Legislating breakout powers after committing to a project would cost governments months of precious time. The relevant authorities include the power to redirect significant tranches of national energy capacity toward training and inference workloads, to consolidate domestic compute and talent under emergency conditions, and to compel cooperation from domestic firms. The U.S. Defense Production Act offers one obvious model, with its expansive definition of national defense, low threshold for invocation, and Title III provisions (which enable the executive branch to incentivize the production of certain critical materials and goods). Other middle powers have analogous authorities. France, for example, can draw on its Code de la défense and Canada on its own Defence Production Act. Middle powers serious about this posture would need to review their existing authorities to confirm they could support a breakout project, and where necessary, update them.

The second requirement is the institutional capacity to decide when to launch the breakout. This decision will be first and foremost a political question, not a technical one. Different middle powers will reach different conclusions about whether reliance on imported frontier capability still makes sense, based on their risk tolerance and security environment. But governments should regularly assess certain core questions in a formal process akin to the integrated security reviews governments currently conduct on defense policy. Such a review should assess questions including: Is the gap between frontier and sub-frontier systems widening or stable? Are existing access arrangements holding or fraying? Has a great-power supplier signaled willingness to use access as leverage? And is automated AI R&D substantially accelerating frontier development abroad? To avoid sleepwalking into a sudden loss of access, or being forced into a hasty breakout under pressure, governments would need to weigh these questions at regular intervals.

The third requirement is the ability to sustain a breakout commitment across several years and potentially multiple governments. This would require shielding it from partisan politics. Framing the project narrowly as a core national security commitment rather than a broader social issue may be one way to do this. Done well, this could look like cross-party political agreements that commit successor governments to the program, statutory protection of breakout-related budget lines, and continuity of AI leadership across multiple administrations. The French nuclear and aerospace programs are middle-power examples of sustained commitment through political change; the British government's commitment to its AI Security Institute—a flagship achievement of the Conservative prime minister Rishi Sunak, but since expanded under successive Labour governments—is another example.

A breakout posture could be mounted by a single middle power or by a coalition. A multistate project comes with the benefit of scale and divided labor. One could in principle imagine an alliance running French or Canadian models, on British-designed chips, funded by Norwegian sovereign wealth, on servers hosted in Finland. But the coordination difficulties such an effort would face will be enormous. Getting a breakout project off the ground in one government department would be hard enough; doing it cross-nationally, while negotiating benefit-sharing and governance arrangements, would be harder still.

The more effective approach would involve building national capabilities first, while embedding hooks in them that make it easier to build a coalition around them down the road. Middle powers may informally network sovereign labs and compute pools, share information on supply chain chokepoints and potential red lines that might trigger a breakout, and agree on technical standards that would let national champion firms combine resources later if necessary. They would trigger formal consolidation only if conditions required it and there were broad political buy-in. For now, champion firms, legal authorities, and political commitment would all remain national, with the option to later consolidate into a joint allied project.

Defensive Capacity

A middle-power decision to launch a breakout would be visible to the great-power suppliers it was designed to hedge against, and those suppliers may have both the capability and the incentive to disrupt it through export controls on hardware, pressure on partner firms, or broader political coercion. Middle powers will have to think seriously about how they plan to respond to such pressure. To pull off the sprint, they would have to screen investments on domestic chokepoints to block foreign acquisition of key assets, and identify the positions in the supply chain where great powers depend on them, so they can use those chokepoints, alongside other sources of leverage, to raise the cost of interference.

Under ordinary political circumstances, middle powers would struggle to leverage control of chokepoints against stronger states that are able to escalate on many other fronts. In a breakout scenario, however, the picture might change. A middle power that has committed to a breakout has already accepted that remaining dependent on imported frontier systems costs more than building its own, and may be willing to pay a high political cost to increase its autonomy. At that point, it only needs enough leverage to deter the most aggressive forms of interference and to survive sustained pressure long enough to complete the sprint.

Middle powers have absorbed great-power pressure before while pursuing high-stakes national projects. India, for example, developed nuclear weapons in the face of years of U.S. sanctions and sustained diplomatic opposition, before Washington accepted India's nuclear status with the 2008 nuclear deal. India effectively absorbed this external pressure thanks to three factors: institutional arrangements that insulated the nuclear program from changes in government; political will across multiple governments to bear the costs of sanctions rather than abandon the program; and its broader geopolitical power, which was ultimately sufficient to persuade the United States to abandon its attempts to isolate India and resort to accommodation instead.

Iran offers a more cautionary tale. Its nuclear program occupied the worst position in the latency space: It was visible and threatening enough to leading powers to invite intervention, but insufficient to deter it. And Iran compounded this error by regularly publicizing its progress in technologies relevant to building nuclear weapons, hoping that it could use its advances as leverage over Washington, instead of holding its cards close to the chest. The results were, of course, sustained sanctions, cyber operations, the assassination of nuclear scientists, and ultimately direct military strikes on Iran's nuclear facilities.

Some of this has to do with posture, but much of it has to do with a country's broader relationship to great powers, and to the United States in particular. That's why it is beneficial for a breakout candidate to maintain strong alignment with the United States throughout. A middle power that takes an adversarial posture will struggle to onshore all the resources

required to ever build frontier AI: Its vulnerability to export controls and outright acts of leverage and sabotage will be severe. There are countervailing forces—greater integration means greater dependency means greater U.S. leverage to stop a breakout sprint—but given the current structure of the chip supply chain, the United States simply controls a prohibitive bottleneck.

Alignment also changes Washington's calculus. If the breakout play were clearly framed as a back-up plan that is complementary to deep integration with the U.S. ecosystem, the United States would have good reason to tolerate it, as it has long tolerated nuclear latency in Japan. Middle powers that know they are not entirely at the mercy of their supplier will find integration easier to accept, politically and commercially, than those that feel trapped; the latter are the ones most likely to turn to Chinese offers or to costly moonshots that fragment the U.S.-led AI ecosystem. A latent capability paired with close integration into the U.S. technology stack poses fewer risks to U.S. interests than the alternatives middle powers will pursue if they conclude they have no fallback at all. The easiest course for the United States would be to meet the demands of mutual integration so that foreign governments never feel the need to trigger a breakout.

Conclusion

A breakout sprint for a frontier AI capability would be among the hardest projects any middle power has sought to pull off, fraught with economic, technological, and political difficulty. If middle powers lose access to frontier AI systems when the stakes are far higher than they are today, governments will need to do very difficult things at speed, including redirecting national energy capacity, mobilizing large quantities of capital, consolidating firms and talent under pressure, and securing access to significant quantities of chips in a tightening market.

Under the best of conditions, none of that will be easy. It will be far harder for governments that have not laid the groundwork in advance. Governments should therefore start planning today. That would involve negotiating access terms into their current data center and chip deals with U.S. companies; making it easier to build domestic compute by removing the obstacles that slow projects down; auditing their emergency authorities and closing any gaps that would otherwise cost months of legislating in a crisis, while standing up regular reviews to judge when the conditions that necessitate a sprint might arrive; and keeping a national or allied champion alive as a fast follower.

Most of these moves are sound AI policy on their own terms, under a range of possible AI futures, and none of them commits a country to a breakout—a decision middle powers should not take lightly. Their primary strategy should remain bargaining hard for reliable access to imported frontier systems, drawing on their existing and potential sources of comparative advantage to make themselves indispensable to the frontier AI value chain. They should continue to attract U.S. investment and integrate their domestic industries with foreign model providers. But the latent breakout capability would form a hedge behind that strategy—a type of insurance that middle powers should hope they never have to use.

> **Middle powers need access to American AI systems, and the United States remains the only democracy with the economic and political weight to decisively shape the trajectory of the technology.**

A latent breakout capacity can help middle powers thread the needle between two errors: overconfidence about their current position and fatalism about their future options. The overconfident view holds that middle powers can catch up to the frontier today, safely make do without frontier access, or take a hostile posture toward the world's most powerful AI companies—and the government that oversees them—without it backfiring. For the foreseeable future, middle powers need access to American AI systems, and the United States remains the only democracy with the economic and political weight to decisively shape the trajectory of the technology. The fatalist view concludes that because catching up is impossible now, it always will be, and that middle powers must therefore accept whatever terms Washington offers, indefinitely. If the United States finds its way to a stable and benign role as the AI supplier to the world through a turbulent technological transition, that approach will work. But placing every chip on that bet would be geopolitical malpractice.

A middle power that can credibly pull compute, talent, and firms into a frontier sprint will be allowed to keep that capability only so long as its great-power supplier has limited reason to fear it. Middle powers that integrate with the U.S. technological order while building the preconditions for greater autonomy underneath will therefore find themselves in the strongest position, whether or not the future demands that they exercise that option. U.S. policymakers, for their part, should understand that the more reliably the United States supplies frontier AI to its allies, the less reason foreign governments will have to pull the trigger on a large-scale effort to free themselves from U.S. technological hegemony. Washington should make itself the kind of supplier that makes a breakout sprint unnecessary.

Notes

1 We expect some of the most sensitive capabilities will remain limited to U.S. national security agencies and developers and illegible to the rest of the world. When we discuss frontier access, we refer to capabilities available to analogous actors in the United States: private firms to private firms, public sector agencies to public sector agencies, consumers to consumers.

CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

1779 Massachusetts Avenue NW
Washington, DC 20036
CarnegieEndowment.org