

The Insurance Policy Nobody Wants to Buy: Why Middle Powers Need an AI Breakout Option

For nineteen days this summer, the most capable AI models on the market simply went dark. On June 12, Anthropic suspended access to its newly released Mythos-tier models to comply with U.S. Department of Commerce export controls. The controls were lifted on June 30, and access returned the next day.

The episode was brief, and it ended. But it settled an argument that capitals from Berlin to New Delhi had been having in the abstract. Frontier AI is not a utility that flows to whoever pays the bill. It is a strategic good, and the switch sits in someone else's hands.

Three pressures now point in the same direction. Developers are handpicking who gets their most sensitive systems first, as Anthropic did with Project Glasswing. Compute is scarce enough that usage gets rationed. And Washington has shown it is willing to shape how frontier models reach the market.

For the United States and China, this is a story about competition. For everyone else, it is a story about exposure. A recent paper on what its authors call an "AI breakout posture" offers one of the more serious answers yet to a question middle powers can no longer defer: what do you do if the tap is turned off when it matters most?

Its answer is counterintuitive, and I think largely right. Don't race to the frontier now. Build the machinery to race later, keep it quiet, and stay close to Washington while you do.

Three doors, all half-shut

Middle powers like Germany, India, Japan, and the United Kingdom have real economic and military weight. None has a frontier model of its own. Their instinctive responses fall into three camps, and each runs into a wall.

The first is the sovereign moonshot: build a national or European champion and close the gap by force of will. France's backing of Mistral and the EU's "gigafactory" compute hubs are the clearest examples. The trouble is arithmetic. Mistral raised money at a valuation of roughly \$14 billion in September 2025; the leading American labs raise multiples of that in single rounds and spend more on compute in a year. By the time Europe's planned clusters come online, they will roughly match what U.S. firms already built.

The moonshot also has a political problem that rarely gets enough attention. It asks voters who are ambivalent about AI to fund tens of billions in speculative infrastructure, with power and grid capacity that legacy industries want for themselves. "Why data centers instead of hospitals?" is a question no finance minister wants to answer for a decade straight.

The second door is fast following: adopt sub-frontier models quickly and let diffusion do the work. This has merit, because most economic value comes from use, not invention. But it is a bet that the gap between the best systems and the rest stays small. That gap has historically been six to nine months. It may not stay there, as compute asymmetries widen, as U.S. labs crack down on distillation, and as automated AI research begins to compound the leaders' advantage.

The third door is bargaining: trade critical minerals, data, or manufacturing capacity for guaranteed access. This remains the best primary strategy. Its flaw is simply that no contract binds a sovereign that changes its mind. June proved that.

The case for waiting


The breakout posture threads between these doors. Its logic borrows from nuclear strategy, specifically Japan's long-standing "latency": the technical capacity to build a weapon quickly, held in reserve under an American security umbrella. Tokyo pays none of the diplomatic

The Insurance Policy Nobody Wants to Buy: Why Middle Powers Need an AI Breakout Option

costs of a bomb while keeping the option alive.

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.



Applied to AI, the idea is to assemble the preconditions for a frontier sprint now, and to launch only if access collapses. Earlier proposals along these lines aimed at a “good enough” sub-frontier fallback. The breakout paper argues, persuasively, that this misreads the threat. If the gap stays narrow, a fallback is barely needed. If the gap is wide and frontier capability is decisive, a second-best model is no insurance at all.

The Insurance Policy Nobody Wants to Buy: Why Middle Powers Need an AI Breakout Option

What makes the argument compelling is that delay is not presented as timidity. It is presented as an advantage, on four fronts.

- Financial. Nobody knows yet whether value will accrue to frontier labs, to small specialized models, or to the application layer. Waiting avoids sinking tens of billions into the wrong bet.
- Political. A speculative multi-year program is hard to sell. A months-long sprint, launched when the danger is visible, is far easier to rally a country behind.
- Geopolitical. A middle power inching toward independence today invites pressure. Securing hardware access quietly, before great powers are watching closely, makes interference harder later.
- Epistemic. The technical recipe keeps changing, from giant pre-training runs to reinforcement learning to inference-heavy systems. America and China can fund many labs betting on different paths. A middle power gets one shot, so it should take that shot when the target is clearer.

The last point is the sharpest in the paper. It reframes the middle-power predicament from a resource problem into an information problem. The question is not only whether you can afford a frontier lab, but whether you know what kind to build.

The plumbing of a latent capability

An option is only worth something if it can be exercised. The breakout posture rests on three kinds of groundwork.

The first is compute, and here the paper makes its most encouraging claim. Frontier training is a small slice of total AI compute; estimates suggest less than a tenth of OpenAI's 2024 compute went to actual frontier training runs. A sprint might therefore need only low single digits of global capacity. That is still beyond any middle-power alliance today, but it is not

fantasy.

Figure 2. Two Ways to Close the Compute Gap

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

Compute today is a small bar, mostly financed by foreign hyperscalers, and falls short of the compute needed for a sprint. Preparing for breakout reaches exactly enough compute for a sprint, with domestic firms and government funding only a small slice and foreign hyperscalers funding the rest. Building a frontier ecosystem goes well past that threshold, but domestic actors fund most of the much larger compute stack.

Who finances the upfront buildout: Domestic firms and government Foreign hyperscalers

ENOUGH COMPUTE FOR A SPRINT

Compute today

Too little for a breakout sprint

Prepare for breakout

Secure access to hyperscaler expansion

Build a frontier ecosystem

Domestic actors fund a much larger compute stack

The Insurance Policy Nobody Wants to Buy: Why Middle Powers Need an AI Breakout Option

Compute Capacity

Getting there means a portfolio. Governments can build some capacity themselves, at great expense. They can nudge domestic firms to build. Most cleverly, they can let American hyperscalers and neoclouds build on their soil, then write “latent access” into the deals: Defense Production Act-style powers, compensated expropriation rights, or narrow clauses that trigger only if the operator cuts the host country off. Add long-term chip supply options with Nvidia and TSMC, executed if needed and resold at a modest loss if not.

The second is private-sector capacity. That means talent pipelines with somewhere domestic to land, and a handful of firms that have done real pre-training and post-training at scale. The paper names Mistral and the UAE’s G42 as the only plausible candidates today, with Canada’s Cohere and Germany’s Aleph Alpha as possible additions. Its advice for them is unusual: keep them deliberately about nine months behind the frontier, sustained by public-sector contracts, rather than burning political capital on a premature moonshot.

The third is state capacity. Governments need emergency legal authorities on the books before a crisis, not legislation drafted during one. They need a regular review process, modeled on defense reviews, that asks whether the frontier gap is widening and whether access deals are fraying. And they need cross-party commitments that let a multi-year effort survive elections, as France’s nuclear program and Britain’s AI Security Institute have.

Where the logic strains

The framework is strong. It is also more fragile than its tidy structure suggests, in at least four places.

First, the nuclear analogy flatters the plan. A bomb is a fixed target: the physics of 1945 still

The Insurance Policy Nobody Wants to Buy: Why Middle Powers Need an AI Breakout Option

works today. A frontier model is a moving one. Japan's latency holds because the threshold never shifts. An AI breakout sprint is chasing a line that moves further away the longer you wait, possibly faster if American labs crack automated AI research. Latency in nuclear terms means "months from a bomb." Latency in AI might mean "months from where the frontier was a year ago."

Second, and most worrying, is the dependency the paper itself flags almost in passing. Frontier labs now write much of their own code with AI coding agents. If those agents are restricted from use in competing AI development, a middle-power sprint would be trying to build a rocket without the tools the rocket makers use. This could turn the timing logic on its head. The paper concedes that a government might face a stark choice between launching immediately or abandoning the idea. If that moment arrives before the other preconditions are ready, the patient strategy collapses into the rushed moonshot it was meant to avoid.

Third, the strategy needs a level of quiet that democracies struggle to keep. Expropriation clauses in data center contracts, long-term chip options, and emergency powers reviews will not stay secret. The paper's own cautionary tale is Iran, whose program became visible enough to provoke but not strong enough to deter. Allied democracies are a world away from Iran, but the lesson about visibility still applies. Hyperscalers may simply decline to build where breakout clauses exist, and the paper acknowledges that sweeteners would be needed to compensate.

Fourth, coalition politics may be harder than the paper allows even in a crisis. Its own example is that no other state will pour billions into a firm as tied to Paris as Mistral. Urgency helps, but it does not dissolve questions of who governs the model, who gets inference first, and whose engineers sit where. India's nuclear path, often cited as proof that middle powers can absorb pressure, was a national project, not a joint venture.

None of these objections defeats the argument. Together they suggest the breakout option is

less a guaranteed escape hatch than a way to improve bad odds, and that its value lies as much in bargaining leverage as in the sprint itself.

Insurance, not independence

The paper's most important move is its insistence that breakout and alignment go together. A middle power that antagonizes the United States, by distilling American models at scale or courting Chinese offers, will find its chip imports squeezed long before any sprint begins. Washington controls the decisive bottleneck in the semiconductor supply chain, and that fact will not change soon.

The breakout posture therefore only works when framed openly as a backstop to deep integration, the way Japan's latency has long been tolerated as a complement to the alliance rather than a challenge to it. This is the version American policymakers should take seriously, because the alternatives are worse for them. Allies who feel trapped are the ones most likely to defect, fragment the ecosystem, or pour money into moonshots that weaken the U.S.-led stack. Allies who know they have a fallback can integrate with less anxiety.

That yields a quietly powerful conclusion. The more reliably the United States supplies frontier AI to its friends, the less likely any of them ever pulls the breakout trigger. Washington's best response to allied hedging is not to suppress it but to make it unnecessary.

For middle-power governments, the to-do list is mostly sensible on its own terms. Write access terms into every data center and chip deal now. Clear the permitting and grid obstacles that slow domestic compute. Audit emergency powers so they work in a crisis. Stand up a regular review of the conditions that would justify a sprint. And keep at least one national or allied champion alive as a fast follower, even when it looks like a poor return.

The Insurance Policy Nobody Wants to Buy: Why Middle Powers Need an AI Breakout Option

There is a fair critique that some of this amounts to buying an option that may expire worthless, or that latency in AI will prove much shallower than latency in nuclear weapons. Skeptics of the AI build-out would add that the whole premise assumes frontier capability is as decisive as its boosters claim. Those are real uncertainties, and a sober government should weigh them.

But June's nineteen-day blackout showed what the alternative looks like. Middle powers have spent years treating frontier access as a given. It is not. The breakout posture asks them to plan for the day it disappears, while working hard to make sure that day never comes. That is not paranoia. It is what insurance has always been: a policy you hope you never have to claim.