

## The Abundance Doctrine

America is trying to win the AI race by withholding technology. China intends to win it by giving technology away - and on the evidence, the giveaway is working.



On 16 July, two ceremonies unfolded in Shanghai within hours of each other. At the World Artificial Intelligence Conference, Beijing-based Moonshot AI unveiled [Kimi K3](#), a 2.8 trillion-parameter system that is the largest open-weight model ever released, and pledged its full weights to anyone on Earth within a fortnight. Across the city, representatives of [29 nations signed the founding agreement](#) of the World AI Cooperation Organization (WAICO), a China-led intergovernmental body headquartered in Shanghai, with the United Nations Secretary-General looking on. One event gave the technology away. The other built the institution to govern its spread. Read together, they reveal a strategy.

Western statecraft in artificial intelligence rests on scarcity. Export controls ration the chips. Frontier models sit behind interfaces whose access can be revoked at will. The operating assumption in Washington, Brussels and Canberra holds that whoever possesses the most capable model possesses the advantage, and that the capability must therefore be guarded. Beijing has made the opposite bet. Its laboratories absorb the enormous cost of training frontier-class systems and then publish the weights - the learned parameters that constitute the model itself - for the world to download, modify and run on its own infrastructure, free of charge and free of permission.

The vocabulary matters. A closed model can be used only through its maker's service, on the maker's terms. An open-weight model can be downloaded, run and adapted by anyone, though its training data and code remain private. A fully open-source model discloses those as well. Nearly every 'open' Chinese model sits in the middle category - open enough to spread, closed enough to guard its recipe.

The bet is paying off. Hugging Face, the largest repository of AI models, reported in its [Spring 2026 review](#) that Chinese models captured 41 per cent of all downloads over the preceding year, overtaking American models in both monthly and cumulative terms for the first time. Alibaba's Qwen family has displaced Meta's Llama as the most downloaded model line on the platform and has passed [a billion cumulative downloads](#). On OpenRouter, a widely used model routing service, [the most popular models are now Chinese and open](#).

The economics explain the geopolitics. An open-weight model crosses a border in the time it takes to transfer a file. a startup in Jakarta, a bank in Singapore or a ministry in Nairobi can host Chinese weights on sovereign servers, tune them to local languages and never fear a foreign capital switching them off - which is why [governments across Asia are building national AI programs on Chinese foundations](#). Every act of Western denial sharpens the appeal. When Washington suspended foreign access to Anthropic's most advanced models in June, the [Center for Strategic and International Studies observed](#) that buyers drew the obvious lesson: American access is conditional, so diversify - and many will diversify into Chinese weights they can own outright. Denial, applied to a technology that diffuses this easily, becomes a recruiting sergeant for the alternative.

America should recognise the strategy, because America invented it. The Global Positioning System began as a closed military capability; Washington opened it to civilian use and, in 2000, switched off the deliberate degradation that reserved precise signals for the Pentagon. Precision positioning became free to every user on Earth - and in return the world built its aviation, shipping, agriculture and financial timing on American infrastructure, American standards and American goodwill. Rivals spent decades and tens of billions building Galileo and BeiDou to escape a dependency Washington created by giving something away. The signal cost nothing at the point of use. The influence it purchased has lasted half a century.

Nations building on Chinese open models should plan for the day the generosity ends: the flow of free frontier weights can stop, or continue on Beijing's terms, once the world depends on it.

That is what Beijing is now buying with open weights, and the purchase compounds. A government that builds its sovereign AI program on Qwen or DeepSeek inherits Chinese architectural choices, Chinese training data and the worldview embedded at training time. Its engineers develop skills specific to that



stack, its vendors build products atop it, and its standards bodies drift towards the ecosystem's norms. The US-China Economic and Security Review Commission [describes mutually reinforcing loops](#) in which open models feed China's industrial base and the industrial base feeds the models. WAICO now supplies the institutional casing: cooperation centres with ASEAN, the Arab League and the African Union, training slots for developing nations, and a Shanghai-headquartered forum in which the norms of the intelligence age can be drafted by its most generous supplier.

There is a tell, however, that abundance is an instrument rather than a philosophy. China's commerce ministry is [weighing whether to restrict the export of future frontier model weights](#). The tap that waters the world has an owner, and the owner is deciding how long it stays open. Nations building on Chinese open models should plan for the day the generosity ends: the flow of free frontier weights can stop, or continue on Beijing's terms, once the world depends on it.

The democratic response exists in embryo. [America's AI Action Plan](#) of July 2025 directed government to encourage open-source and open-weight models as carriers of American standards. OpenAI returned to openness with its [gpt-oss releases](#) under permissive licensing. The [ATOM Project](#) is campaigning to rebuild American open-model capacity, Palantir and Nvidia are [fielding open Nemotron models in sovereign environments](#) for US government agencies, and in late July [more than 230 companies](#) - Cisco, Dell and GitHub among them - signed an open letter declaring open weights essential to American AI leadership. The ingredients are assembling. What is missing is doctrine.

Call it the Abundance Doctrine: the deliberate use of freely distributed technology to set the standards, dependencies and habits of the international system. Under it, allied governments would treat capable open models as public strategic infrastructure - funded, maintained and given away the way GPS is given away - rather than as commercial by-products. AUKUS partners already count AI among their shared advanced capabilities; an allied open-model program, trained on democratic data under democratic law and released for the world to own, would convert that ambition into distribution. For Australia, the doctrine is liberating. A middle power will never outspend the frontier laboratories, but it can co-fund, fine-tune and champion allied weights across the Indo-Pacific, where the contest for defaults is being decided now.

Honesty requires two concessions. The most dangerous capabilities - those touching biological design or large-scale offensive cyber operations - should stay unpublished, because weights, once released, can never be recalled; abundance must operate below that threshold and denial above it. And export controls on chips retain their force, since training frontier systems still demands hardware that controls can reach. The doctrine refines denial rather than abolishing it: withhold the means of production, give away the product, and let the world build on yours.



The contest for the intelligence age is being decided now, in the defaults of a thousand ministries and startups, and Beijing is arming the world with its judgement while the West audits its own. Democracies must flood the world with intelligence built to their values, or watch the world think in someone else's. Great powers have always armed themselves. The ones that shaped history armed everybody else.