There is a version of American China policy in which the White House punishes Beijing's AI champions, protects Silicon Valley's dominance, and still lets American AI companies build whatever they want, however fast they want, with as few rules as possible. It is a comforting story. It is also not a strategy — it's a truce among factions that don't agree on what winning even means, and the United States is running out of time to notice the difference.

Chinese models forced this argument onto the West Wing's agenda, not the other way around. Systems like Moonshot AI's Kimi K3 have gone from curiosities to credible competitors within months, rattling parts of Silicon Valley and unsettling U.S. security officials in equal measure. The threat isn't just that Chinese labs are closing the benchmark gap. It's that they are exporting open-weight models that undercut American companies on price, giving developers around the world a Chinese technological stack to build careers and products on.

Publicly, the administration insists America is "winning the AI race." Privately, officials concede that Chinese labs are moving faster than expected, in ways the old toolkit — chip export controls, CFIUS reviews, entity lists — was never built to handle. The deepest worry isn't a single military breakthrough. It's the quieter possibility that a generation of developers across the Global South standardizes on Chinese models and infrastructure, and America finds itself locked out of markets it never even realized it was competing for.

Three camps, one incoherent policy

Inside the administration, at least three factions are pulling in different directions, and none of them has the authority — or the inclination — to overrule the other two.

The **sanctions-and-squeeze camp** wants to treat Chinese AI firms the way Washington once treated Huawei: as strategic threats to be starved of capital and market access. They point to what they call industrial-scale distillation of American models and want to deploy financial weapons — sanctions, export restrictions, limits on U.S. capital markets — against firms they accuse of laundering American research into cheaper knockoffs.

The **market-share hawks** think the entire sanctions framework is looking at the wrong scoreboard. For them, victory means the world runs on an American AI stack rather than a Chinese one, and the only way to secure that is to move faster: flood allies with U.S. chips and cloud capacity, strip away domestic guardrails, and trust that American innovation will simply outrun Beijing's.

The **security-first technocrats**, clustered mostly in cyber and science policy offices, want something narrower and less glamorous: real visibility into training data, limits on the riskiest open-weight releases, and mechanisms that can actually penalize distillation when it happens — for models built in the U.S. and in China alike.

These aren't three variations on a theme. They cancel each other out. One week, officials float banning Chinese open-weight models outright; days later, others are quietly telling Commerce and industry that no such ban is actually coming. The administration talks about "de facto bans" through sanctions and liability exposure, then pivots to insisting that innovation, not regulation, will decide this race. Both cannot be the plan.

The distillation fight nobody has resolved

Nothing exposes the confusion better than the argument over distillation and open weights. When U.S. firms accused Chinese labs of mounting what they called the largest known distillation attack on their models, hawks treated it as proof of systematic theft — compress an American model, undercut its price, and sell it back to the world as something new. Treasury officials floated sanctions against Moonshot and its peers.

But the obvious remedy — restricting Chinese open-weight releases — collides with a different article of faith inside American AI policy: that openness is itself a competitive advantage, a hedge against a handful of U.S. giants controlling the entire stack. The idea of banning Chinese open-weight models surfaced and then quietly died, killed by pushback from open-source advocates and the American startups that rely on those very models as free benchmarks and leverage against domestic incumbents.

So Washington ends up telling U.S. labs to be cautious about releasing their most capable systems as open weights, while worrying in the same breath that any American restraint just hands the openness narrative — and the global developer base that comes with it — to China. It flirts with liability rules that would make using Chinese models legally risky for American firms, all while avoiding the harder question underneath: what happens when a developer in Jakarta, Lagos, or São Paulo picks a Chinese model simply because the license is looser and the price is lower?

“Winning” means whatever the person in the room wants it to mean

Strip away the tactics and there’s a conceptual hole at the center of this debate: nobody has agreed on what winning actually means. For some, it’s narrow and defensive — keep frontier models, data centers, and chips away from the People’s Liberation Army. For the market-share hawks, it’s expansive — control the chips, the models, the clouds, and the applications that most of the world’s population touches every day.

That second vision is driving a sweeping export push, looser environmental rules for data centers, and a rollback of the prior administration’s safety-focused measures, all on the theory that speed is America’s only real edge and that any domestic friction just buys China more runway. Yet the same White House also leans on chipmakers like Nvidia not to abandon the Chinese market entirely, on the theory that a total cutoff could backfire — even as other officials argue that any sale of advanced chips to China accelerates the very progress they’re trying to blunt. It talks about “guardrails” in summits with Xi Jinping and offers little beyond the word itself.

The net effect is a loud, high-confidence confrontation with China layered over a patchwork of tools that constrain and enable Chinese AI in roughly equal measure, depending on which agency, and which week, you happen to ask.

What coherence would actually require

The United States cannot sanction its way to an AI lead, and it cannot deregulate its way there either. A serious strategy starts from an uncomfortable admission: this contest is about systems and norms as much as it is about hardware and models.

That means, first, a clear and enforceable red line on what counts as unacceptable exploitation of American IP — backed by real investment in watermarking, provenance tracking, and model-behavior forensics, not just sanctions announced after the fact. It means, second, an honest reckoning with how much openness the U.S. is willing to tolerate in its own ecosystem, since every open-weight American model is also raw material for someone else's system, foreign or domestic.

It means, third, treating allies as partners rather than props. If market share of the global stack is truly the metric that matters, U.S. policy has to make American infrastructure easy to choose — legally, financially, politically — without asking other democracies to accept permanent dependency and zero say over standards or safety. That looks like shared safety regimes and joint export frameworks, not one-way exports dressed up as loyalty tests.

And it means the White House stops treating its own AI safety debates as a distraction from the China competition. The tools that make models harder to misuse and easier to audit are the same tools that make it possible to prove, credibly, when a foreign actor has crossed a line. Instead of an endless turf war between growth advocates who want to "let it cook" and technocrats worried about externalities, the administration needs to say plainly where it believes AI creates genuine national security risk — and where it's content to let the market run.

Right now the debate in Washington is framed as a binary: break decisively with Chinese AI, or race against it unencumbered. The real choice is whether the United States builds a stable, predictable framework for how its own AI ecosystem will engage a rising Chinese one — or stays trapped in a cycle where every new Chinese model triggers another few weeks of improvisation behind closed doors.