May 17, 2026

Reversing Course on Semiconductor Export Controls Was a Strategic Mistake

The United States has long leveraged dominance over critical economic "chokepoints," including the dollar, fiber optic infrastructure, and the SWIFT financial messaging system, to shape global political outcomes in its favor (Krugman, 2024). Semiconductors represent the next frontier of that competition, because leadership in AI will determine which power holds the capacity to define technological standards and set the terms of the coming international order. For nearly a decade, both the Trump and Biden administrations treated American dominance over advanced semiconductor design as a strategic asset worth protecting, using export controls to deny Beijing the compute it needed to train frontier AI models. In December 2025, President Trump reversed that posture, allowing Nvidia to sell its advanced H200 processors to Chinese customers in exchange for a government share of twenty-five percent of the sales revenue (Horowitz et al., 2025). Proponents of the reversal, including Nvidia CEO Jensen Huang, argue that the controls were always ineffective, that DeepSeek proved China could innovate around hardware restrictions regardless, and that continued restrictions only accelerated Chinese technological self-reliance while costing American companies billions. The evidence suggests otherwise, and the cost of walking away from the last decisive chokepoint in the US-China AI competition is considerably higher than the reversal's proponents acknowledge. The United States should reimpose export controls on advanced semiconductors because the controls were meaningfully constraining Chinese AI development at the frontier, because neither DeepSeek's benchmark performance nor the self-sufficiency objection constitute grounds for abandoning the last leverage point that remains decisive, and because the credibility cost of repeatedly reversing course has become a strategic liability that makes future leverage harder to sustain.

American export controls were producing visible constraints on Chinese AI development at the frontier. “In every other element of the AI stack—data, talent, algorithms, applications, and electricity generation—China matches or surpasses the United States,” leaving advanced semiconductor access as

the single domain where American leverage remains decisive (Horowitz et al., 2025). The CEO of one of China's leading AI firms put it plainly, stating that "the main obstacle he faces is the embargo on high-end chips" (Amodei and Pottinger, 2025). The operational picture on the ground tells the same story, with Chinese labs smuggling high-performance Nvidia chips and engineers struggling with domestic alternatives that overheated, crashed, and lacked adequate software support (Wei et al., 2025). Morgan Stanley analysts estimated that roughly 95 out of every 100 advanced chips produced by SMIC, China's largest and most prominent contract chipmaker, were unusable (Wei et al., 2025). Training frontier AI models requires coordinating hundreds of thousands of chips across sustained compute runs, and a manufacturing process this defective cannot support development at that scale. The practical consequence is a hard ceiling on the complexity and speed of Chinese AI development, which is precisely the outcome export controls were designed to produce. What this picture describes is not a country that had found its way around the controls but one that was genuinely struggling against them, forced to smuggle what it could not buy, delay what it could not build, and accept failure rates that would be commercially unsustainable in any other competitive market.

The most compelling evidence cited against the controls is DeepSeek, the Chinese AI startup that released a model in January 2025 matching the capabilities of frontier American systems at a fraction of the cost and under the hardware constraints that strict export controls imposed. To proponents of the reversal, DeepSeek demonstrated that the controls were never the binding constraint on Chinese AI development. However, the structural conditions that made DeepSeek possible suggest the controls were never as airtight as their architects assumed. A robust black market in controlled computing chips allowed Chinese labs to acquire high-performance Nvidia chips through third-party intermediaries regardless of what Washington permitted. Beyond smuggling, physical possession of a chip is not a prerequisite for using it, and Chinese companies found ways to access computing resources located in unrestricted countries, effectively routing around the restrictions without ever violating them directly (Villasenor, 2025). Taken together, these two channels suggest that the controls created the appearance of constraint without the substance of it, and that DeepSeek's success was in large part due to the fact that the

restrictions were never as binding as initially intended.

The DeepSeek counterargument overstates the significance of what was achieved under constraint, both in terms of how innovative the gains actually were and how constrained China actually was when it achieved them. Amodei, co-founder of Anthropic, argues that DeepSeek's efficiency gains were not a unique breakthrough but an expected point on an ongoing cost reduction curve, meaning they reflected relatively standard algorithmic progress rather than evidence that hardware constraints had ceased to matter (Amodei, 2025). DeepSeek was also not operating under the kind of chip scarcity the counterargument implies, given that it had approximately 50,000 Hopper-generation chips, many acquired before export bans took effect. A company with that level of hardware investment and access is not a meaningful test case for whether export controls can impose binding constraints at scale. Even so, the controls were still visibly biting, as DeepSeek had to delay the release of its latest model because of chip shortages, indicating that the hardware ceiling was actively limiting the pace and scope of what Chinese AI development could produce (Wei et al., 2025). What DeepSeek built beneath the compute ceiling and what China could build without one are not points on the same continuum. Optimization under scarcity produces models engineered around the constraints of what is available. Unrestricted access to frontier hardware produces models engineered around the limits of what is possible, and those two development environments produce categorically different strategic outcomes at the scale where military AI superiority is actually determined. Ultimately, a benchmark model delayed by compute scarcity is not evidence that the controls had no effect. It is evidence that they were working.

The deeper problem with the DeepSeek counterargument is that it evaluates export controls against the AI competition of today rather than the one that will actually determine strategic outcomes. Frontier AI development is not static, and a field that has consistently doubled in compute requirements cannot be evaluated on the basis of what one company built with 50,000 chips in 2024. Amodei argues that reaching the compute threshold required for superintelligence-scale models and the military applications that flow from them will require millions of chips, and that no amount of algorithmic optimization substitutes for raw hardware at that scale (Amodei, 2025). Which power acquires that

threshold first will determine whether the world ends up with a durable American AI advantage or a bipolar order in which China's industrial base and military focus translate into strategic dominance (Amodei, 2025). DeepSeek's benchmark performance at moderate compute scales cannot answer that question. Only chip access can, and export controls were the one policy tool capable of keeping that access out of reach. If China could produce something this consequential while constrained to 50,000 chips, the ceiling is not an argument for liberalization. It is the strongest available evidence for why the ceiling must be maintained.

Another consequential argument against the controls is that they were counterproductive, forcing a level of indigenous Chinese investment in semiconductor development that open access would never have produced while simultaneously cutting off the American revenue that funds the next generation of innovation. Nvidia CEO Jensen Huang articulates this concern directly, arguing that the restrictions cost Nvidia access to the world's second largest economy and the commercial relationships that come with it (Wei et al., 2025). He further contends that China's position as home to half of the world's AI developers means continued restrictions only strengthen the case for a fully indigenous Chinese AI stack. That prospect carries strategic implications beyond lost revenue, as a world in which China develops a self-sufficient AI ecosystem is one in which a significant portion of the global south may come to build critical infrastructure around a Chinese AI stack that is cheaper and more accessible than its American alternatives, consolidating Chinese technological influence in global markets. The empirical record offers some support for this concern, with analysts noting that export controls have forced Chinese firms to get better across the board, producing meaningful progress in domestic chip development that would not have occurred under conditions of open access (Swanson et al., 2024). On this reading, the controls were not just ineffective but counterproductive, trading short-term constraints for a long-term Chinese semiconductor industry that American firms will have neither access to nor influence over.

However, what this argument fails to contend with is that accepting that restrictions accelerate self-sufficiency as a reason not to impose them would render economic coercion permanently unavailable

as a tool of statecraft. The United States cut Iran off from SWIFT in 2012, driving its inflation to nearly 40 percent despite the near-certain knowledge that Tehran would seek alternative financial arrangements in response (Krugman, 2024). It waged a sustained campaign against Huawei and ZTE that crippled both companies' access to American components, accepting that Chinese firms would eventually attempt to build around the restrictions. In neither case did the possibility of accelerating self-sufficiency disqualify the policy, because that risk is not a flaw unique to semiconductor controls. It is an inherent feature of every instance of economic coercion, and the decision to impose restrictions has always turned on whether the constraint imposed in the interim justifies that risk. China was pursuing indigenous semiconductor development long before export controls existed, because its leadership understood that dependence on American technology was a structural vulnerability regardless of whether Washington chose to exploit it. Open access would not have prevented Chinese self-sufficiency. It would have made China more comfortable remaining dependent for longer while still building toward the same destination on its own timeline. At the level of military AI development where raw compute determines the ceiling of what is possible, as Amodei and Pottinger note, American restrictions mean it will likely take China years if not decades to catch up in chip quality and quantity, and that gap does not close faster by handing China the chips it is still trying to learn to build (Amodei and Pottinger, 2025).

Even setting aside the empirical question of whether the controls were working, the pattern of imposing and then abandoning them has itself become a strategic liability that compounds with every reversal. The credibility of economic coercion as a tool of statecraft depends entirely on the belief that restrictions, once imposed, will be maintained long enough to matter. When Washington reverses course under industry pressure, that belief erodes, and the mechanism of the December 2025 reversal makes the credibility damage particularly hard to repair. The Trump administration lifted the controls in exchange for a twenty-five percent share of Nvidia's H200 sales revenue to Chinese customers, effectively converting a national security instrument into a licensing arrangement (Horowitz et al., 2025). That structure tells Beijing that American technology restrictions are opening positions in a commercial

negotiation, and that sufficient lobbying pressure from American firms with Chinese revenue exposure is enough to move Washington off a policy it spent years constructing and defending. Each reversal does not just damage the current policy but reduces the deterrent value of every future one. Camba's analysis of the 2025 export control showdown makes the underlying logic explicit, arguing that power in modern economic statecraft depends less on who can create chokepoints than on who can sustain them (Camba, 2026). A tool that Washington keeps putting down is not a weapon China has any reason to fear, and a twenty-five percent share of the sales revenue is not a substitute for the strategic leverage that the transfer itself surrenders.

Ultimately, the case for lifting semiconductor export controls rests on a misreading of the evidence and an underestimation of what abandoning them costs. The controls were creating meaningful constraint at exactly the frontier where American leverage over Chinese AI development remains decisive, DeepSeek's gains under scarcity reveal the limits of what algorithmic efficiency can substitute for at frontier scale, and the self-sufficiency objection, taken to its logical conclusion, would render economic coercion permanently unavailable as a tool of statecraft. What unites these arguments is that none turns on whether the controls were perfect, only on whether they were worth maintaining, and on that question the evidence points clearly in one direction. The December 2025 reversal made clear that the limiting factor was never the design of the controls but the political will to hold them against commercial pressure. Reimposing controls is necessary, but sustaining them requires the kind of multilateral coordination that was taking shape before the reversal dismantled it, because a chokepoint only holds if the states that control it hold together. Semiconductors are not a trade commodity whose export terms can be adjusted at the margin in response to industry lobbying. They are the input on which the most consequential technological competition in modern history currently turns, and the decision about whether to restrict their export to China is not one the United States can afford to keep revisiting every four years.