

# Diffusion as Strategy? Export Controls, Open Weights, and the Global Reach of Chinese AI Models, 2018–2026

Alfredo Merlet

[alfredo.merlet@gmail.com](mailto:alfredo.merlet@gmail.com)

Universidad del Alba <https://orcid.org/0009-0007-8558-164X>

---

## Research Article

**Keywords:** export controls, artificial intelligence, open-weight models, China, technology diffusion, semiconductors, foundation models, open source, difference-in-differences, Epoch AI

**Posted Date:** September 25th, 2026

**DOI:** <https://doi.org/10.21203/rs.3.rs-11153363/v1>

**License:** © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

**Additional Declarations:** The authors declare no competing interests.

---

# Abstract

Policy analysts increasingly argue that United States export controls on advanced AI chips pushed Chinese developers towards a distinctive strategy: releasing model weights openly and competing through diffusion rather than frontier scale. This claim is widely repeated but has not been tested at the level of individual models. I assemble 3,034 AI models released between January 2018 and September 2026 from the Epoch AI database, code each model's origin, weight accessibility, training compute, training hardware and—for derivative models—the country of the base model on which it was built, and estimate difference-in-differences models around the three main rounds of US controls (October 2022, October 2023 and December 2024). Relative to US models, Chinese models became 41 percentage points more likely to be released with open weights by the end of the period ( $p = 0.002$ ), but the divergence did not follow the first two control rounds: it emerged only after December 2024 and grew through 2025–2026, and a large part of it reflects US developers closing their models rather than Chinese developers opening theirs. Within-organisation estimates are smaller (15–20 points), and a placebo test in the pre-control period is marginal, so the timing evidence is consistent with a strategic shift that consolidated in 2025 rather than with an immediate response to controls. The diffusion evidence is stronger: the share of derivative models built on a Chinese base model rose from 2.4% before 2023 to 51% in 2025, including 18% of US and 33% of other non-Chinese derivatives. Chinese models increasingly report export-compliant chips, and the frontier compute gap, after narrowing for five years, widened relative to trend after December 2024. The results support the diffusion thesis on reach but not the simple claim that controls caused openness.

**JEL codes:** O33, F52, L86, O38

## 1. Introduction

Since October 2022 the United States has restricted the export of advanced computing chips and chip-making equipment to China in three major rounds, with the stated aim of limiting China's ability to train and deploy frontier artificial intelligence systems. Over the same period Chinese developers released a sequence of widely adopted open-weight models, and by 2026 Chinese model families had become a common foundation for developers elsewhere. A growing body of policy commentary connects the two developments. Analysts at the Center for Strategic and International Studies argue that China's ability to innovate in open-weight models can be attributed in large part to the export-control regime itself (CSIS 2026). Others note that compute-constrained developers may release weights precisely because they cannot serve their models at scale (East Asia Forum 2026). The underlying intuition draws on the argument that technological leadership depends less on invention at the frontier than on diffusion across the economy (Ding 2024).

The claim that export controls pushed China towards openness and diffusion is consequential. If it holds, a policy designed to slow a rival's AI development may have accelerated the global spread of that rival's models, with implications for standard-setting, security and the design of future controls. Yet the

claim has rarely been tested. Existing assessments are either qualitative accounts of the control regime (Allen 2022; CSET 2023) or aggregate comparisons of capabilities between leading models (Maslej et al. 2025; Brookings 2026). As one recent assessment observes, the strength of Chinese models does not by itself reveal what the controls did or how Chinese development would have evolved without them (East Asia Forum 2026). What is missing is a model-level analysis with an explicit comparison group and a timing structure that can distinguish a response to controls from a secular trend.

This article provides that analysis. I use the Epoch AI database of machine learning models, which records for each model its developer, country of origin, publication date, accessibility of weights, training compute, training hardware and—for derivative models—the base model it was built on (Epoch AI 2026). The sample comprises 3,034 models released between January 2018 and September 2026, of which 807 were developed by Chinese organisations and 1,503 by US organisations. I estimate difference-in-differences models that compare the release decisions of Chinese and US developers before and after each control round, with quarter, domain and organisation fixed effects, an event-study specification, a placebo test and a battery of robustness checks. I then use the base-model information to measure diffusion directly, tracing the share of derivative models worldwide that are built on Chinese foundations.

The findings qualify the prevailing narrative in three ways. First, Chinese and US openness did diverge sharply: by the end of the period Chinese models were 41 percentage points more likely than US models to be released with open weights. But the divergence did not follow the October 2022 or October 2023 controls. It emerged after December 2024 and widened through 2025 and 2026, and a substantial part of it reflects US developers closing their models rather than Chinese developers opening theirs. Within-organisation estimates are smaller, and a pre-period placebo is marginal. The evidence is therefore consistent with a strategic shift that consolidated in 2025—a period that also includes the release of DeepSeek-V3 and DeepSeek-R1 (DeepSeek-AI 2024, 2025)—rather than with an immediate reaction to the controls. Second, the diffusion evidence is strong and unambiguous. The share of derivative models built on a Chinese base model rose from 2.4% before 2023 to 51% in 2025, and the rise extends to non-Chinese developers: in 2025, 18% of US derivatives and 33% of derivatives from other countries were built on Chinese bases. Third, Chinese developers increasingly report training on chips designed to comply with the controls, and the gap between the largest US and Chinese training runs, which narrowed steadily until 2024, widened relative to that trend after December 2024.

The article contributes to three literatures. It adds a quantitative, model-level test to the policy literature on export controls and technology denial, which has so far relied on qualitative assessment. It extends the economics of open source—where opening is understood as a strategic choice by firms seeking adoption, complementary benefits or catch-up (Lerner and Tirole 2002; von Hippel and von Krogh 2003; Henkel 2006)—to the geopolitics of AI model release. And it offers a direct measure of technological diffusion through base-model lineage, complementing platform-level studies of open AI ecosystems (Osborne et al. 2024) and the diffusion theory of great-power competition (Ding 2024).

Section 2 reviews the literature and derives hypotheses. Section 3 describes the data and the policy timeline. Section 4 sets out the empirical strategy, and Section 5 reports the results on openness. Section 6 presents the evidence on diffusion, compute and hardware, and Section 7 discusses interpretation and limitations. Section 8 concludes.

## **2. Literature and hypotheses**

### **2.1 Export controls and technology denial**

The controls exploit the concentration of the semiconductor supply chain in a few firms and countries (Miller 2022), and they target compute because, under empirical scaling laws, training compute is a principal determinant of model capability (Kaplan et al. 2020; Hoffmann et al. 2022). The October 2022 rule restricted exports to China of advanced computing chips above specified performance thresholds and of equipment for manufacturing advanced semiconductors (BIS 2022). Commentators at the time described it as an attempt to cut China off from the future of AI rather than merely to slow it (Allen 2022). The October 2023 update lowered the thresholds to capture chips that designers had adapted to fall just below the original limits and expanded the geographic scope of the controls (BIS 2023; CSET 2023). The December 2024 rule added high-bandwidth memory and further manufacturing equipment and placed 140 additional entities on the Entity List (BIS 2024). The literature on these measures is dominated by legal analysis, industry assessment and policy argument. A common theme is that controls constrained Chinese access to compute—Chinese firms were estimated to hold a small share of global AI compute by 2025—without preventing Chinese developers from producing competitive models (East Asia Forum 2026; Brookings 2026). Measurement of compute as an object of governance has advanced considerably (Sevilla et al. 2022; Sastry et al. 2024), but there is little quantitative evidence on how controls changed developers' strategic choices.

### **2.2 Openness as strategy**

The economics of open source treats the release of code as a strategic choice rather than altruism. Firms may reveal technology selectively to attract external contributions, to establish standards or to commoditise complements (Lerner and Tirole 2002; Henkel 2006), and open development can combine private investment with collective benefits (von Hippel and von Krogh 2003). In AI, the release of model weights sits on a gradient running from fully closed systems to fully open ones (Solaiman 2023), and its social consequences have been debated extensively (Bommasani et al. 2021; Kapoor et al. 2024). Critics warn that openness in AI is often partial, because weights released openly still depend on concentrated compute, data and labour (Widder et al. 2024). For a developer constrained in compute, however, open release has specific attractions. It shifts the cost of inference to users, it builds adoption where the developer cannot compete on serving capacity, and it creates an ecosystem of derivative work that reinforces the developer's models as a standard.

### **2.3 Diffusion and the measurement gap**

Ding (2024) argues that great-power advantage in general-purpose technologies has historically come from diffusion capacity rather than from leadership in frontier innovation. Applied to AI, this implies that the relevant measure of competitive position is not only who trains the largest model but whose models are most widely built upon. Platform-level evidence shows that adoption on open model hubs is highly concentrated in the models of a few companies (Osborne et al. 2024), and policy reports in 2026 describe Chinese model families as a common base for new derivative models (CSIS 2026). What the literature lacks is a systematic measure of diffusion that can be tracked over time and compared across countries. The base-model lineage recorded in the Epoch database provides one.

## 2.4 Hypotheses

Four hypotheses follow. H1 (openness): after the export controls, Chinese developers became more likely than US developers to release model weights openly. H2 (compute gap): the gap between the largest US and Chinese training runs widened after the controls. H3 (hardware substitution): Chinese developers shifted towards chips designed to comply with the controls and towards domestically produced accelerators. H4 (diffusion): the share of derivative models built on Chinese base models increased, including among non-Chinese developers. H1 and H4 are the primary hypotheses. H2 is secondary, because compute is reported for fewer than half of the models, and H3 is examined descriptively, because hardware is reported for only a quarter of Chinese models.

## 3. Data

### 3.1 The Epoch AI database

The data come from the Epoch AI “Data on AI Models” database, which tracks machine learning models from the academic literature, company releases and technical reports, is updated continuously and is distributed under a CC BY 4.0 licence (Epoch AI 2026). I use the “All AI Models” file, downloaded on 24 September 2026, which contains 3,620 models, and the “Notable AI Models” subset of 1,072 models that meet Epoch's notability criteria. I retain the 3,034 models released between 1 January 2018 and 21 September 2026. Origin is coded from the country of the developing organisation: 807 models are Chinese only, 1,503 are US only, 93 are joint Chinese–US and 631 are from other countries or of unknown origin. The analysis of openness compares Chinese-only and US-only models.

Weight accessibility is recorded in seven categories. I code a model as open if its weights are released under any open-weights category (unrestricted, restricted use or non-commercial), and use unrestricted release as an alternative outcome. Accessibility is known for 626 Chinese and 1,373 US models. Training compute is reported for 35% of Chinese and 43% of US models. Training hardware is reported for 25% of Chinese models. For derivative models, the base model is recorded by name; I match each base model to its own entry in the database and assign it the country of its developer, which succeeds for 99.6% of the 679 derivative models in the sample. Table 2 summarises the sample.

### 3.2 Policy timeline

Table 1 lists the three rounds of US export controls used as treatment dates. Two later changes fall inside the sample but are not coded as treatment rounds: the January 2025 framework for AI diffusion, which introduced controls on model weights but was rescinded before taking effect, and a January 2026 rule that moved exports of some less advanced chips to China to case-by-case review. Chinese domestic policy also changed during the period, most notably with the interim measures for generative AI services in force from August 2023 (CAC 2023). These overlapping events are discussed as threats to identification in Section 4.

Table 1  
US export-control rounds used as treatment dates

| Date                                                                                                                                                                                                                                                                                                    | Instrument                                                                                       | Main content relevant to AI training                                                                                                                                                                                                  |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 7 Oct 2022                                                                                                                                                                                                                                                                                              | BIS interim final rule on advanced computing and semiconductor manufacturing items (87 FR 62186) | Licence requirements for exports to China of advanced computing chips above performance thresholds and of equipment for advanced semiconductor manufacturing; restrictions on US persons supporting advanced chip production in China |
| 17 Oct 2023                                                                                                                                                                                                                                                                                             | BIS updates to the advanced computing and semiconductor manufacturing rules                      | Revised performance parameters to capture chips adapted to fall below the 2022 thresholds; expanded country scope and anti-circumvention measures                                                                                     |
| 2 Dec 2024                                                                                                                                                                                                                                                                                              | BIS rules on advanced computing and semiconductor manufacturing items                            | Controls on high-bandwidth memory and additional manufacturing equipment; 140 entities added to the Entity List                                                                                                                       |
| <i>Notes: Later changes not coded as treatment rounds: the January 2025 AI diffusion framework (rescinded before its compliance date in May 2025) and a January 2026 rule moving exports of some less advanced chips to China to case-by-case review. Sources: BIS (2022, 2023, 2024); CSET (2023).</i> |                                                                                                  |                                                                                                                                                                                                                                       |

Table 2  
Sample composition, January 2018 – September 2026

| Origin                                                                                                                                                                                                                                                                           | Models | Accessibility known | Open share, pre-Oct 2022 | Open share, post-Oct 2022 | Compute reported | Hardware reported |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------|---------------------|--------------------------|---------------------------|------------------|-------------------|
| China                                                                                                                                                                                                                                                                            | 807    | 626                 | 37.8%                    | 63.8%                     | 35%              | 25%               |
| United States                                                                                                                                                                                                                                                                    | 1,503  | 1,373               | 45.8%                    | 48.3%                     | 43%              | 43%               |
| <i>Notes: Open share = share of models with accessibility known that were released with open weights (any open-weights category). Joint Chinese–US models (93) and models from other countries are excluded from this table. Source: Epoch AI (2026); author's calculations.</i> |        |                     |                          |                           |                  |                   |

## 4. Empirical strategy

For model  $i$  released by developer  $j$  in quarter  $t$ , the baseline specification is a linear probability model

$$Open_{ijt} = \beta_0 China_i + \beta_1 China_i \cdot Post_{1,t} + \beta_2 China_i \cdot Post_{2,t} + \beta_3 China_i \cdot Post_{3,t} + \delta_t + \gamma_{domain} + \lambda_{orgtype} + \varepsilon_{ijt},$$

where  $Post_1$ ,  $Post_2$  and  $Post_3$  equal one for models released on or after 7 October 2022, 17 October 2023 and 2 December 2024 respectively. Because the indicators are cumulative,  $\beta_1$  measures the change in the China–US openness gap after the first round,  $\beta_2$  the additional change after the second, and  $\beta_3$  the additional change after the third; the sum  $\beta_1 + \beta_2 + \beta_3$  is the cumulative change by the end of the period. Quarter fixed effects  $\delta_t$  absorb global shifts in release practice, and domain and organisation-type fixed effects absorb compositional differences between, for example, language and vision models or industry and academic developers. A second specification replaces organisation-type effects with organisation fixed effects, so that identification comes only from changes in release behaviour within the same developer. Standard errors are clustered by developing organisation (Cameron and Miller 2015). The linear probability model is preferred for interpretability (Angrist and Pischke 2009); a logit specification is reported as a robustness check.

The dynamic specification replaces the three post indicators with China  $\times$  half-year interactions from the second half of 2022 onwards, with all earlier periods pooled as the reference. Because every Chinese model is treated at the same calendar dates, the design does not suffer from the negative-weighting problems of staggered adoption (Goodman-Bacon 2021; Callaway and Sant'Anna 2021; Sun and Abraham 2021), but the event study remains necessary to assess pre-trends and to locate the timing of any divergence (Roth et al. 2023). As a placebo, I assign a fictitious control round to October 2020 and estimate the model on releases before October 2022.

Identification rests on the assumption that, absent the controls, the openness of Chinese and US models would have evolved in parallel after conditioning on the fixed effects. Three threats are important. First, other China-specific events coincide with the treatment dates, including domestic generative AI regulation and the release of DeepSeek-V3 in December 2024, which was followed by rapid imitation among Chinese developers. Second, US release practice may have changed for reasons unrelated to China, so that a widening gap could reflect US closing rather than Chinese opening. Third, the pre-control sample of Chinese models with known accessibility is small (68 models before July 2022), limiting the precision of pre-trend tests. The results are interpreted in light of these threats, and the article does not claim that any coefficient identifies the causal effect of a single control round.

Diffusion (H4) is analysed separately. For each derivative model I record whether its base model is Chinese, and estimate the share of derivatives on Chinese bases by period and by the origin of the derivative developer, with linear probability models that include period indicators and a language-model control, clustered by organisation. The compute gap (H2) is measured as the difference between the largest US and the largest Chinese training run (in  $\log_{10}$  FLOP) in each quarter from 2019 to 2026,

regressed on post-round indicators with and without a linear trend and with heteroskedasticity- and autocorrelation-consistent standard errors. Hardware (H3) is summarised descriptively.

## 5. Results: openness

### 5.1 Descriptive divergence

Figure 1 plots the share of models released with open weights by half-year. Before October 2022, 37.8% of Chinese and 45.8% of US models with known accessibility had open weights. After October 2022 the Chinese share rose to 63.8% while the US share was essentially unchanged at 48.3%. The half-yearly series shows that the divergence is concentrated at the end of the period: the Chinese share fluctuated between 54% and 77% from 2023 onwards, while the US share fell from roughly one half in 2024 and early 2025 to between 12% and 34% from the second half of 2025. Part of the widening gap therefore reflects US developers moving towards closed release.

### 5.2 Difference-in-differences estimates

Table 3 reports the baseline estimates. With quarter, domain and organisation-type fixed effects, the China × post-October 2022 coefficient is 0.139 and the China × post-October 2023 coefficient is 0.042, neither statistically significant. The China × post-December 2024 coefficient is 0.228 ( $p = 0.011$ ). The cumulative change in the China–US openness gap by the end of the period is 0.408 ( $p = 0.002$ ): Chinese models became 41 percentage points more likely than US models to be released with open weights, relative to the pre-control difference. With organisation fixed effects the pattern is the same but the magnitudes are smaller: the post-December 2024 coefficient is 0.153 ( $p = 0.044$ ) and the cumulative change is 0.199 ( $p = 0.106$ ). About half of the cross-organisation divergence therefore reflects changes within developers, and the rest reflects entry of new, more open Chinese developers and exit or closing of more open US ones.

Table 3  
Export controls and open-weight release: difference-in-differences estimates

| China                                                                                                                                                                                                                                                                                                                                                              | (1) Baseline                       | (2) Organisation fixed effects |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|--------------------------------|
|                                                                                                                                                                                                                                                                                                                                                                    | <b>-0.129 (0.085) [0.128]</b>      | <b>absorbed</b>                |
| China × post-Oct 2022                                                                                                                                                                                                                                                                                                                                              | 0.139 (0.129) [0.283]              | 0.051 (0.133) [0.701]          |
| China × post-Oct 2023                                                                                                                                                                                                                                                                                                                                              | 0.042 (0.102) [0.683]              | -0.006 (0.093) [0.953]         |
| China × post-Dec 2024                                                                                                                                                                                                                                                                                                                                              | 0.228 (0.089) [0.011]              | 0.153 (0.076) [0.044]          |
| Cumulative effect after Dec 2024                                                                                                                                                                                                                                                                                                                                   | 0.408 (0.133) [0.002]              | 0.199 (0.123) [0.106]          |
| Observations / organization clusters                                                                                                                                                                                                                                                                                                                               | 1,996 / 313                        | 1,869 / 313                    |
| Fixed effects                                                                                                                                                                                                                                                                                                                                                      | Quarter, domain, organization type | Quarter, domain, organization  |
| <i>Notes: Linear probability models; dependent variable equals one if the model was released with open weights. Coefficients with standard errors clustered by organisation in parentheses and p-values in brackets. The cumulative effect is the sum of the three post-round interactions. Sample: Chinese and US models with known accessibility, 2018–2026.</i> |                                    |                                |

The event study in Fig. 2 locates the timing more precisely. Coefficients for the second half of 2022 through the first half of 2025 lie between  $-0.08$  and  $0.26$ , and none is statistically distinguishable from zero at the 5% level. The coefficient rises to  $0.49$  in the second half of 2025 and to  $0.57$ – $0.58$  in 2026, each with a 95% confidence interval excluding zero. There is no evidence of a discrete change following October 2022 or October 2023; the divergence is gradual and becomes significant roughly a year after the December 2024 round.

## 5.3 Robustness and placebo

Table 4 reports alternative specifications. The post-December 2024 coefficient is larger when the outcome is restricted to unrestricted open weights ( $0.332$ ) and when the sample is restricted to notable models ( $0.464$ ) or to language models ( $0.283$ ), and it is  $0.193$  ( $p = 0.064$ ) among industry developers. Excluding releases from 2026, which may be incompletely coded, reduces it to  $0.178$  ( $p = 0.050$ ). A logit specification with year fixed effects gives a post-December 2024 log-odds coefficient of  $1.011$  ( $p = 0.007$ ). The cumulative effect is positive and significant at the 5% level in every variant.

The placebo is less reassuring. Assigning a fictitious control round to October 2020 and estimating on the pre-control sample yields a placebo coefficient of  $0.254$  ( $p = 0.073$ ). The pre-control Chinese sample is small and its openness share volatile (Fig. 1), so this estimate is imprecise, but it indicates that shifts in the China–US openness gap of similar magnitude occurred before any control round. The placebo reinforces the conclusion from the event study: the data support a large divergence in openness by 2025–2026, but they do not isolate the controls as its cause.

Table 4  
Robustness and placebo

| Specification                                                                                                                                                                                                  | China × post-Dec 2024                | Cumulative effect           | N     |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|-----------------------------|-------|
| Unrestricted open weights only                                                                                                                                                                                 | 0.332 (0.092) [ $< 0.001$ ]          | 0.490 (0.098) [ $< 0.001$ ] | 1,996 |
| Notable models only                                                                                                                                                                                            | 0.464 (0.143) [0.002]                | 0.712 (0.142) [ $< 0.001$ ] | 606   |
| Language models only                                                                                                                                                                                           | 0.283 (0.097) [0.004]                | 0.441 (0.145) [0.002]       | 1,331 |
| Industry developers only                                                                                                                                                                                       | 0.193 (0.103) [0.064]                | 0.425 (0.181) [0.019]       | 1,336 |
| Excluding 2026 releases                                                                                                                                                                                        | 0.178 (0.091) [0.050]                | 0.357 (0.140) [0.011]       | 1,858 |
| Logit, year FE (log-odds)                                                                                                                                                                                      | 1.011 (0.372) [0.007]                | —                           | 1,999 |
| Placebo: fake round Oct 2020, sample before Oct 2022                                                                                                                                                           | 0.254 (0.140) [0.073] (placebo term) | —                           | 513   |
| Notes: All specifications except the logit include quarter fixed effects and the controls of Table 3, column 1, as applicable. Standard errors clustered by organisation in parentheses; p-values in brackets. |                                      |                             |       |

## 6. Results: diffusion, compute and hardware

### 6.1 Diffusion through derivative models

The evidence on diffusion is considerably stronger. Of the 676 derivative models with a matched base model, the share built on a Chinese base rose from 2.4% in 2018–2022 to 10.2% in 2023, 28.7% in 2024 and 51.0% in 2025 (Table 5 and Fig. 3). The increase occurs among developers of every origin. Among Chinese developers, the share of derivatives on Chinese bases rose from 16.7% before 2023 to 94.7% in 2025, indicating near-complete substitution away from non-Chinese foundations such as Llama. Among US developers, it rose from zero before 2023 to 16.0% in 2024 and 18.2% in 2025. Among developers from other countries or of mixed origin, it rose from zero to 32.5% in 2025. For non-Chinese derivatives as a group, the linear probability model gives increases of 5.8 percentage points in 2023, 15.9 points in 2024 and 22.7 points in 2025 relative to the pre-2023 baseline, all significant at the 5% level (the 2026 estimate rests on eight models).

Table 5  
Share of derivative models built on a Chinese base model, by period and origin of the derivative developer

| Period                                                                                                                                                                                                                                                                                                                                                                                                          | Chinese developers | US developers  | Other or mixed | All   | Non-Chinese: LPM coefficient |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|----------------|----------------|-------|------------------------------|
| 2018-22                                                                                                                                                                                                                                                                                                                                                                                                         | 16.7% (n = 12)     | 0.0% (n = 54)  | 0.0% (n = 19)  | 2.4%  | (reference)                  |
| 2023                                                                                                                                                                                                                                                                                                                                                                                                            | 23.1% (n = 39)     | 4.9% (n = 81)  | 8.1% (n = 37)  | 10.2% | 0.058 (0.023) [0.014]        |
| 2024                                                                                                                                                                                                                                                                                                                                                                                                            | 60.3% (n = 68)     | 16.0% (n = 94) | 14.7% (n = 68) | 28.7% | 0.159 (0.035) [< 0.001]      |
| 2025                                                                                                                                                                                                                                                                                                                                                                                                            | 94.7% (n = 75)     | 18.2% (n = 77) | 32.5% (n = 40) | 51.0% | 0.227 (0.066) [< 0.001]      |
| 2026                                                                                                                                                                                                                                                                                                                                                                                                            | 100.0% (n = 4)     | 57.1% (n = 7)  | 0.0% (n = 1)   | 66.7% | 0.501 (0.234) [0.033]        |
| <i>Notes: Derivative models are those for which Epoch records a base model; base models are matched to their own database entries to determine origin (99.6% matched). The last column reports period coefficients from a linear probability model for non-Chinese derivatives (N = 478), with a language-model control and standard errors clustered by organisation in parentheses; p-values in brackets.</i> |                    |                |                |       |                              |

These shares are measured within the Epoch database, which records notable and large models rather than the full population of derivatives on open model hubs, and they should not be read as shares of global downloads. They nevertheless show that, among the models that researchers and developers considered worth documenting, Chinese foundations moved within three years from almost no role to supplying half of all derivative work, and that this role extends well beyond China.

## 6.2 The frontier compute gap

Figure 4 compares the largest reported US and Chinese training runs in each quarter. Before October 2022 the US maximum exceeded the Chinese maximum by 2.04 orders of magnitude on average; after October 2022 the average gap was 0.95. Table 6 (Panel A) shows that in levels the gap fell significantly after October 2023 ( $-0.99$ ,  $p = 0.009$ ), consistent with continued Chinese catch-up. Once a linear trend is included, the gap narrows by 0.17 orders of magnitude per quarter over the whole period ( $p = 0.051$ ), and the post-December 2024 indicator is positive (1.24,  $p = 0.036$ ): after the third control round the gap stopped closing and widened relative to its trend. This pattern is consistent with the controls beginning to bind at the frontier, as several assessments have anticipated, but it rests on 28 quarterly observations of a noisy maximum and on incomplete compute reporting, particularly for Chinese frontier models, and is therefore treated as suggestive.

## 6.3 Hardware

Among Chinese models with reported training hardware, the share trained on chips designed to comply with the controls—NVIDIA's H800, A800 and H20 series—rose from zero before October 2022 to 10.8% in

the following year, 18.6% between October 2023 and December 2024 and 55.8% after December 2024 (Table 6, Panel B). The share reporting domestic accelerators such as Huawei's Ascend series fluctuated between 2% and 12%, and the share reporting the most advanced restricted NVIDIA chips remained at 7% or below. Hardware is reported for only a quarter of Chinese models, and reporting may itself respond to the controls, so these figures describe the reporting population rather than the Chinese AI sector as a whole. They are consistent with developers adapting to the controls through compliant hardware rather than through large-scale domestic substitution during the period.

Table 6

Panel A – Frontier compute gap (US maximum minus Chinese maximum, log<sub>10</sub> FLOP), quarterly regressions

| Constant (pre-Oct 2022 mean gap)                                                                                          | Levels                      | With linear trend           |
|---------------------------------------------------------------------------------------------------------------------------|-----------------------------|-----------------------------|
|                                                                                                                           | 2.041 (0.560) [ $< 0.001$ ] | 2.995 (0.572) [ $< 0.001$ ] |
| Linear quarterly trend                                                                                                    | —                           | -0.174 (0.089) [0.051]      |
| Post-Oct 2022                                                                                                             | -0.428 (0.603) [0.477]      | 0.960 (0.990) [0.332]       |
| Post-Oct 2023                                                                                                             | -0.993 (0.380) [0.009]      | -0.212 (0.470) [0.651]      |
| Post-Dec 2024                                                                                                             | 0.196 (0.431) [0.650]       | 1.237 (0.589) [0.036]       |
| Quarters                                                                                                                  | 28                          | 28                          |
| Notes: OLS on quarterly observations, 2019Q2–2026Q3; HAC standard errors (two lags) in parentheses; p-values in brackets. |                             |                             |

Table 6

Panel B – Training hardware reported for Chinese models

| Period                                                                                                                                                            | Models with hardware reported | Export-compliant NVIDIA (H800, A800, H20) | Domestic accelerators (e.g. Ascend) | Advanced restricted NVIDIA (H100 and later) |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------|-------------------------------------------|-------------------------------------|---------------------------------------------|
| pre-Oct22                                                                                                                                                         | 35                            | 0.0%                                      | 8.6%                                | 0.0%                                        |
| Oct22-Oct23                                                                                                                                                       | 37                            | 10.8%                                     | 5.4%                                | 0.0%                                        |
| Oct23-Dec24                                                                                                                                                       | 86                            | 18.6%                                     | 2.3%                                | 7.0%                                        |
| post-Dec24                                                                                                                                                        | 43                            | 55.8%                                     | 11.6%                               | 7.0%                                        |
| Notes: Shares of Chinese models with reported training hardware. Periods are delimited by the three control rounds. A model can appear in more than one category. |                               |                                           |                                     |                                             |

## 7. Discussion

## 7.1 Interpretation

The results support one part of the diffusion-as-strategy thesis and qualify another. The part that holds is diffusion itself: by 2025 Chinese models had become the foundation for half of documented derivative work, including a substantial and growing share of derivatives developed outside China. Whatever its causes, this is the outcome that the diffusion theory of technological competition identifies as decisive (Ding 2024), and it arrived quickly.

The part that requires qualification is causality through openness. If export controls had pushed Chinese developers towards open release, the openness gap should have widened after October 2022 or October 2023, when the controls first constrained access to advanced chips. It did not. The divergence emerged after December 2024 and became significant in the second half of 2025, a period in which the release of DeepSeek-V3 and DeepSeek-R1 demonstrated that competitive performance could be achieved with an open release and lower reported cost (DeepSeek-AI 2024, 2025), and in which US developers moved towards closed release. The evidence is consistent with a strategic equilibrium that consolidated in 2025, in which open release became the dominant Chinese strategy and closed release the dominant US one, with the controls as one background condition among several. Distinguishing the contribution of the December 2024 controls from the demonstration effect of DeepSeek and from shifts in US commercial strategy would require variation that the present data do not contain.

The contrast between the openness and diffusion results is itself informative. Chinese openness rose only moderately in absolute terms, from around 40% before the controls to 54–77% afterwards, while Chinese foundations' share of derivative work rose from almost nothing to half. Diffusion therefore depends less on how many Chinese models are open than on which ones are: a small number of widely adopted open model families appear to account for most derivative activity, consistent with the concentration of adoption documented on open model platforms (Osborne et al. 2024).

## 7.2 Implications

For the design of export controls, the findings suggest that policies targeting compute may affect the frontier—the compute gap stopped narrowing after December 2024—while leaving the diffusion of capable models largely untouched, because diffusion operates through open weights that are not subject to the same controls. For the debate on open AI, the results show that openness has become a dimension of geopolitical positioning as well as of research practice and commercial strategy (Widder et al. 2024; Kapoor et al. 2024). For measurement, base-model lineage provides a simple and replicable indicator of diffusion that can be tracked alongside compute and capability indices.

## 7.3 Limitations

Several limitations apply. The Epoch database is curated and records notable or large models, so its coverage of Chinese models may differ from its coverage of US models, and coverage may change over time. Accessibility is coded as of the database's latest revision, so a model reclassified after release is recorded in its final status. Compute and hardware are reported for fewer than half of the models and

less often for Chinese ones. The derivative analysis rests on 676 models and measures documented derivatives rather than downloads or usage. The treatment rounds coincide with other events, and the pre-control Chinese sample is small, which limits pre-trend tests and makes the placebo imprecise. Finally, the analysis compares China with the United States only; developers in other countries faced different constraints and are excluded from the openness comparison.

## 8. Conclusion

This article has tested, at the level of individual models, the widely repeated claim that US export controls pushed Chinese AI developers towards open release and diffusion. The evidence confirms the diffusion outcome—Chinese foundations underpinned half of documented derivative models by 2025, including a growing share outside China—and confirms that Chinese and US release strategies diverged sharply. It does not confirm that the controls caused the divergence: openness did not change after the first two control rounds, the gap widened only from 2025, and a substantial part of it reflects US closing. The controls may have shaped the conditions under which Chinese developers chose openness, and they may have begun to constrain the frontier, but the diffusion of Chinese AI appears to have been driven by the success of open model families in a period when closed release became the US norm. Evaluating future controls will require attention not only to who trains the largest models but to whose models the rest of the world builds upon.

## Declarations

## Competing interests.

The author declares no competing interests.

## Ethics declaration.

Not applicable. The study uses only publicly available data on AI models and contains no information on identifiable individuals.

## Use of artificial intelligence.

The author used Claude (Anthropic), a large language model, as an assistant for code drafting, data processing, verification of policy dates and bibliographic records against primary sources, and language editing. All research design decisions, the interpretation of results and the final text were reviewed and approved by the author, who takes full responsibility for the content. All quantitative results were produced by executable scripts included in the replication package.

## Funding.

This research received no specific funding.

## Author contributions.

A.A.M.E. conceived the study, designed the analysis, constructed the data, performed the estimation and wrote the manuscript.

## Data availability.

The Epoch AI “Data on AI Models” database is publicly available under a CC BY 4.0 licence at <https://epoch.ai/data/ai-models>. The analysis file, the base-model matching, all estimation outputs and the downloaded source files are deposited in Zenodo under CC BY 4.0: <https://doi.org/10.5281/zenodo.22938430>.

## Code availability.

Python scripts (pandas, pyfixest 0.60.0, statsmodels, matplotlib) that build the analysis file from the Epoch downloads and reproduce every table and figure are included in the Zenodo deposit cited above.

## References

1. Allen GC (2022) Choking off China's AI access to the future of AI. Center for Strategic and International Studies, Washington, DC, 11 October 2022
2. Angrist JD, Pischke J-S (2009) Mostly Harmless Econometrics: An Empiricist's Companion. Princeton University Press, Princeton
3. BIS (2022) Implementation of additional export controls: Certain advanced computing and semiconductor manufacturing items; supercomputer and semiconductor end use; entity list modification. Federal Register 87:62186. Bureau of Industry and Security, US Department of Commerce
4. BIS (2023) Commerce strengthens restrictions on advanced computing semiconductors, semiconductor manufacturing equipment, and supercomputing items to countries of concern. Press release and interim final rules, 17 October 2023. Bureau of Industry and Security, US Department of Commerce
5. BIS (2024) Commerce strengthens export controls to restrict China's capability to produce advanced semiconductors for military applications. Press release and rules, 2 December 2024. Bureau of Industry and Security, US Department of Commerce
6. Bommasani R, Hudson DA, Adeli E et al (2021) On the opportunities and risks of foundation models. arXiv:2108.07258.

7. Brookings Institution (2026) Competing AI strategies for the US and China. Brookings, Washington, DC, 10 June 2026
8. CAC (2023) Interim Measures for the Management of Generative Artificial Intelligence Services. Cyberspace Administration of China and six other agencies, in force 15 August 2023.
9. Callaway B, Sant'Anna PHC (2021) Difference-in-differences with multiple time periods. *J Econ* 225(2):200–230
10. Cameron AC, Miller DL (2015) A practitioner's guide to cluster-robust inference. *J Hum Resour* 50(2):317–372
11. CSET (2023) Explainer: The Commerce Department's October 2023 export control update. Center for Security and Emerging Technology, Georgetown University
12. CSIS (2026) China's open-weight challenge to U.S. AI leadership. Center for Strategic and International Studies, Washington, DC, 2026
13. DeepSeek-AI (2024) DeepSeek-V3 technical report. arXiv:2412.19437
14. DeepSeek-AI (2025) DeepSeek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning. arXiv:2501.12948
15. Ding J (2024) Technology and the Rise of Great Powers: How Diffusion Shapes Economic Competition. Princeton University Press, Princeton
16. East Asia Forum (2026) China's AI models aren't proof that US chip controls have failed. East Asia Forum, 9 September 2026.
17. Epoch AI (2026) Data on AI Models. Epoch AI, <https://epoch.ai/data/ai-models> (accessed 24 September 2026; CC BY 4.0).
18. Goodman-Bacon A (2021) Difference-in-differences with variation in treatment timing. *J Econ* 225(2):254–277
19. Henkel J (2006) Selective revealing in open innovation processes: The case of embedded Linux. *Res Policy* 35(7):953–969
20. Hoffmann J, Borgeaud S, Mensch A et al (2022) Training compute-optimal large language models. arXiv:2203.15556.
21. Kaplan J, McCandlish S, Henighan T et al (2020) Scaling laws for neural language models. arXiv:2001.08361
22. Kapoor S, Bommasani R, Klyman K et al (2024) On the societal impact of open foundation models. arXiv:2403.07918.
23. Lerner J, Tirole J (2002) Some simple economics of open source. *J Industrial Econ* 50(2):197–234
24. Maslej N, Fattorini L, Perrault R et al (2025) The AI Index 2025 Annual Report. Stanford Institute for Human-Centered Artificial Intelligence, Stanford University
25. Miller C (2022) Chip War: The Fight for the World's Most Critical Technology. Scribner, New York
26. Osborne C, Ding J, Kirk HR (2024) The AI community building the future? A quantitative analysis of development activity on Hugging Face Hub. *J Comput Social Sci* 7(2):2067–2105.

27. Roth J, Sant'Anna PHC, Bilinski A, Poe J (2023) What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *J Econ* 235(2):2218–2244
28. Sastry G, Heim L, Belfield H et al (2024) Computing power and the governance of artificial intelligence. *arXiv:2402.08797*
29. Sevilla J, Heim L, Ho A, Besiroglu T, Hobbhahn M, Villalobos P (2022) Compute trends across three eras of machine learning. In: 2022 International Joint Conference on Neural Networks (IJCNN). *arXiv:2202.05924*
30. Solaiman I (2023) The gradient of generative AI release: Methods and considerations. In: Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23), 111–122
31. Sun L, Abraham S (2021) Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *J Econ* 225(2):175–199
32. von Hippel E, von Krogh G (2003) Open source software and the private-collective innovation model: Issues for organization science. *Organ Sci* 14(2):209–223
33. Widder DG, Whittaker M, West SM (2024) Why 'open' AI systems are actually closed, and why this matters. *Nature* 635:827–833. <https://doi.org/10.1038/s41586-024-08141-1>

## Figures




**Figure 1**

Share of Chinese and US models released with open weights, by half-year (models with known accessibility; half-years with at least five models). Dashed lines mark the October 2022, October 2023

and December 2024 export-control rounds. Source: Epoch AI (2026).




Figure 2

Event-study estimates of the China–US difference in the probability of open-weight release, by half-year, relative to all periods before July 2022, with 95% confidence intervals. Quarter, domain and organisation-type fixed effects; standard errors clustered by organisation.




Figure 3

Share of derivative models built on a Chinese base model, by period and origin of the derivative developer. Labels give the number of derivative models in each cell. Source: Epoch AI (2026); author's matching of base models.




**Figure 4**

Largest reported training compute (log<sub>10</sub> FLOP) among US and Chinese models in each quarter, 2019–2026. Dashed lines mark the three export-control rounds. Quarterly maxima depend on which models report compute and are therefore noisy.