AI Diffusion as Complementary Innovation: Evidence from GenAI Tutorials and GitHub Production in China and the United States

Binxin Gao

June 4, 2026

Abstract

Generative AI may diffuse less through access to frontier models than through the complementary innovations that make those models useful. We study online tutorials as a market for these complements. Unlike earlier tutorials, which primarily transmitted knowledge that viewers then had to implement, GenAI tutorials increasingly distribute *executable economic artifacts*: prompts, workflows, agents, templates, mini-apps, automation scripts, and marketing funnels that viewers can copy, adapt, and deploy with limited expertise. We assemble a bilingual corpus of **21,780** GenAI tutorials from Bilibili and YouTube, spanning 2018–2026; code each tutorial at the artifact level using validated large-language-model classifications; map **496,453** Bilibili recommendation edges; and link tutorial demand to **26,505** GitHub repositories. Three facts emerge. First, the object of diffusion has changed. The share of tutorials distributing an executable artifact rises sharply after major model releases, reaching roughly three-quarters of Bilibili tutorials by 2026. Second, executable artifacts diffuse nonlinearly. Conditional on views, topic, and time, artifact-bearing tutorials generate 14 percent more downstream AI-tutorial recommendations than knowledge-only tutorials, and bookmark-style engagement rises more than fourfold with executability, consistent with viewers saving tools for reuse rather than merely consuming information. Third, ecosystems differ in the speed and form of operationalization. Chinese tutorials are more executable, rely more on domestic models and no-code agents, and monetize through private communities and paid courses; U.S. tutorials are more likely to remain explainers and monetize through affiliates and product links. A difference-in-differences around DeepSeek-R1 shows that China converts a domestic model shock into deployable artifacts about 12 percentage points faster than the United States. These same artifacts both lower and concentrate entry. Turnkey tools sharply reduce barriers and roughly triple the likelihood that a tutorial points to monetization, but copyability also raises imitation and attention remains winner-take-most. On the supply side, GitHub production mirrors tutorial demand: building is dominated by individual developers, Chinese developers build disproportionately on Chinese-origin models, and a tool's tutorial demand tracks its repository supply. We interpret tutorial-borne artifacts as complementary innovations that carry AI from point solutions to system solutions. The economics of GenAI diffusion is therefore less about model access alone than about how quickly ecosystems package, replicate, and monetize the complements that make those models productive.

Keywords: generative AI, technology diffusion, complementary innovation, platforms, entrepreneurship, China.
JEL: O33, L86, D83, O53, L26.

1 Introduction

The economic study of technology diffusion has long treated the diffusing object as *information*: a hybrid seed variety, a managerial practice, or a software skill spreads as agents learn that it exists and learn how to use it (Griliches, 1957; Rogers, 2003; Foster and Rosenzweig, 1995; Conley and Udry, 2010). Tutorials, whether printed manuals or instructional videos, accelerated this process by lowering the cost of acquiring knowledge. But the viewer still had to convert knowledge into action: learning how to format a spreadsheet did not hand you the spreadsheet.

Generative AI changes the unit of diffusion. Because GenAI tools turn natural-language instructions into working outputs, a tutorial can do more than teach—it can *distribute the working output itself*. A video about resume writing can hand viewers a copy-paste prompt that drafts resumes; a video about customer service can share a configured agent that answers customers; a video about side income can package an end-to-end automation that produces and posts content while you sleep. We call these *executable economic artifacts*: reusable objects that compress a productive capability into a form deployable with minimal technical skill. The central premise of this paper is that GenAI diffusion is increasingly the spread not of ideas but of *executable capabilities*.

This reframing connects directly to the dominant economic account of how AI creates value. Agrawal et al. (2022) argue that AI's economic payoff arrives only as adopters climb a ladder from *point solutions* (drop AI into an existing task) to *application* and ultimately *system solutions* that redesign workflows and business models—and that the binding constraint is rarely the AI itself but the *complementary innovations* needed to embed it. An executable artifact is precisely such a complement: a workflow, agent, or script is the packaged, transferable form of the system redesign that Agrawal et al. (2022) place at the top of the ladder. In the language of Goldfarb and Tucker (2019), distributing artifacts collapses the *replication, search, and verification costs* of adopting an AI workflow—a viewer no longer has to discover, rebuild, and validate the complement, only to copy it. Tutorials are thus a market institution for diffusing AI's complements outside the boundary of the firm, and a knowledge-only tutorial (a point-solution explainer) and an artifact-bearing tutorial (a transferable system solution) are economically distinct objects even when they describe the same model. Our empirical task is to measure that distinction at scale and trace its consequences.

This distinction matters for economics because it changes the diffusion dynamics. When tutorials transmit only knowledge, the number of new adopters is bounded by the number of viewers who both watch and successfully act. When tutorials transmit artifacts, three things happen. (i) The conversion from viewing to capability is near-instant, because the artifact *is* the capability. (ii) Artifacts are *reproductive*: many of them (a content-automation workflow, a product-image generator) themselves produce more sellable outputs, and often more tutorials, generating a branching process rather than a one-shot transfer. (iii) Because artifacts are cheap to copy, the same capability proliferates rapidly, raising both genuine entry and low-quality replication. The net effect on entrepreneurship—does GenAI lower barriers to entry, or mainly amplify replication and winner-take-most dynamics?—is theoretically ambiguous and, we argue, depends on platform incentives, monetization channels, and local AI infrastructure.

We study this empirically by comparing the two largest tutorial ecosystems on either side of the AI frontier: Bilibili in China and YouTube in the United States. The comparison is of first-order policy interest—it is the explicit focus of this conference and a growing literature on AI in China (Beraja et al., 2023)—and it is also methodologically useful, because the two platforms differ sharply in monetization (China's knowledge-payment and private-traffic economy vs. the U.S. advertising-and-affiliate economy), in available foundation models (a rich domestic stack including DeepSeek, Doubao, Tongyi/Qwen, ERNIE, Kling, and Coze), and in regulation. These differences let us ask not just *whether* executable capabilities diffuse, but *which kinds* spread where, and *who captures the value*.

Data and measurement. We assemble a corpus of 21,780 GenAI tutorials—10,866 on Bilibili and 10,914 on YouTube—collected through public platform interfaces over 2018–2026, with complete engagement statistics, upload timestamps, creator identifiers, and full text (titles and descriptions). For China we additionally crawl the Bilibili recommendation graph by snowball, recording 496,453 "related-video" edges among more than one hundred thousand nodes; this graph is the platform's own diffusion substrate and lets us measure how far an artifact propagates. Our central measurement innovation is *artifact-level coding*: we use a large language model (Claude Haiku 4.5 via Amazon Bedrock) to read each tutorial and emit a structured record—whether it distributes an executable artifact and of which type(s), an executability score (0–4), a reusability score, whether the artifact is reproductive, the suggested monetization channel(s), the target user, the named AI tools, and the concrete economic practice it enables. We validate this instrument against a frontier model (Claude Sonnet 4.6) on 758 dual-coded tutorials, obtaining Cohen's $\kappa = 0.80$ for the core "is this a GenAI tutorial" classification, $\kappa = 0.81$ for domain, and correlation $r = 0.79$ for the ordinal executability score—reliability comparable to trained human coders on subjective constructs.

Findings. We organize the evidence around four claims.

1. The artifact turn (Section 4). GenAI tutorials are not mostly explainers. Sixty-three percent of Chinese and 59% of U.S. tutorials distribute at least one executable artifact, and the share has risen steeply—from roughly one-third in early 2023 to three-quarters by 2026 on Bilibili—with visible discontinuities around major model releases. The most common artifact is the multi-step *workflow* ($\sim$50%), followed by automation scripts, prompts, and agents.

2. Nonlinear diffusion (Section 5). Using the Bilibili recommendation graph, we show that executable tutorials reproduce more. In a negative-binomial model with views, domain, and time fixed effects, distributing an executable artifact raises the count of downstream AI-tutorial "offspring" by an incidence-rate ratio of 1.14 ($p < 10^{-9}$); crucially, executable and knowledge-only tutorials have nearly identical *total* recommendation degree but executable ones are embedded in denser *AI-capability* clusters. On the demand side, the "save-to-reuse" rate—favorites or coins per view, a revealed intention to return and use the content—rises monotonically with executability, from 1.1%

at the lowest level to 4.8% at the highest ($p < 10^{-80}$). Both margins are consistent with capability diffusion behaving more like a branching process than a one-shot transfer.

3. Cross-country divergence (Section 6). China's tutorials are more executable and more operational: relative to the U.S., they over-index on automation scripts (+8.2%), workflows, prompt packs, and structured courses, and they reference a distinctly domestic, no-code-heavy tool stack (DeepSeek, Doubao, Tongyi/Qwen, Coze). U.S. tutorials retain more knowledge-only explainers (+6.9%) and lean on the developer stack (Claude, Copilot, n8n). Monetization diverges accordingly: Chinese tutorials more often route value through private-traffic communities and paid courses, U.S. tutorials through affiliate links (+9.9%) and product sales. Event studies sharpen the picture: around the release of DeepSeek-R1 (January 2025) and Sora (December 2024), executable share *rose* in China while it *fell* in the U.S., suggesting China more rapidly operationalizes model shocks—especially domestic ones—into reusable tools.

4. Value capture and replication (Section 7). Attention is extraordinarily concentrated on both platforms, and substantially more so in the U.S.: the top 1% of tutorials capture 37% (China) and 53% (U.S.) of views, with view Ginis of 0.84 and 0.91 and creator-level Ginis of 0.87 and 0.95. Higher copyability predicts more near-duplicate replication, and the relationship is stronger in the U.S. (replica rate rising from 17% to 23% across reusability levels) than in China (5% to 17%)—evidence that lowering the cost of copying capabilities amplifies both entry and low-quality imitation, with the mix shaped by the platform.

Contributions. We make four contributions. First, *conceptual*: we reframe GenAI diffusion as the spread of executable capabilities and identify tutorial-borne artifacts as the *complementary innovations* that, in the language of Agrawal et al. (2022), carry adopters from point solutions to system solutions—making concrete and measurable a mechanism that account has so far treated theoretically. Second, *measurement*: we introduce a scalable, validated artifact-level coding instrument that converts free-text tutorials into structured economic objects, in the spirit of the large-scale usage measurement of Handa et al. (2025); and we show platform recommendation graphs can serve as *observed* diffusion substrates. Third, *a two-sided design*: we pair the *demand and learning* side (tutorials, where capabilities are taught) with the *supply and production* side (26,505 GitHub repositories, where the artifacts are built and forked), and show the two track each other—grounding the diffusion in both what creators teach and what developers build. Fourth, *empirical*: we provide the first large-scale China–U.S. comparison of *what* GenAI economic practices spread, *who* captures value, and *how* platform institutions shape the answer, including a difference-in-differences showing China operationalizes domestic model shocks into deployable artifacts faster than the U.S. Our results speak directly to debates on AI and entrepreneurship (Otis et al., 2024; Brynjolfsson et al., 2025), AI adoption (Bick et al., 2024; Humlum and Vestergaard, 2024; McElheran et al., 2024), AI in China (Beraja et al., 2023), and the macroeconomics of AI (Acemoglu, 2024), reconciling the equalize-vs-concentrate debate through the rung of the ladder an artifact occupies, and grounding diffusion in what users actually do with these tools.

2 Conceptual Framework

We present a framework that distinguishes knowledge diffusion from capability diffusion and yields the predictions we test.

Tutorials as objects with two attributes. Index tutorials by i. Each carries two attributes beyond its topic: an *executability* $x_i \in \{0, 1, 2, 3, 4\}$, capturing how directly a viewer can copy-and-deploy what the tutorial distributes ($0 =$ pure concept; $4 =$ a ready-to-run artifact is provided), and a binary *reproductive* flag $\rho_i \in \{0, 1\}$, equal to one if the distributed artifact itself generates more sellable outputs or content at scale (e.g., a content-automation pipeline) rather than producing a single result.

Knowledge diffusion. In the classical benchmark (Griliches, 1957; Bass, 1969), a tutorial with V_i viewers produces adopters $A_i = \phi\, V_i$, where ϕ is the share who both understand and successfully implement. Adoption is linear in viewership and bounded by the implementation hazard ϕ, which is decreasing in required skill.

Capability diffusion. When a tutorial distributes an executable artifact, two things change. First, the implementation hazard collapses: $\phi(x_i) \to 1$ as $x_i \to 4$, because the artifact *is* the capability and requires little skill to deploy (*barrier-lowering*). Second, and crucially, a reproductive artifact ($\rho_i = 1$) lets adopters produce further artifacts and tutorials, so each adopter seeds, in expectation, β further pieces of capability-bearing content. The number of "offspring" tutorials generated by tutorial i is then a branching process with reproduction number

$$\hat{R}_i \;=\; \underbrace{\phi(x_i)}_{\text{conversion}} \;\cdot\; \underbrace{\beta\,\rho_i}_{\text{reproduction}} \;\cdot\; \underbrace{m(\text{platform})}_{\text{amplification}}, \tag{1}$$

where $m(\cdot)$ captures platform amplification (recommendation, monetization incentives to repackage). If $\hat{R}_i > 1$ for a class of artifacts, capability diffusion is self-sustaining: a single super-spreader artifact can generate a cascade of derivative tools, content, and businesses. Knowledge-only tutorials have $\rho_i = 0$ and thus $\hat{R}_i \approx 0$ in this content-generating sense, even if they are widely viewed. In the ladder of Agrawal et al. (2022), executability x_i indexes the rung: low x_i is a point-solution explainer, and only at $x_i = 4$—a turnkey artifact—is the system solution itself transferred, so that the conversion $\phi(x_i)$ that lowers the barrier to deployment is steepest at the top of the ladder, not in the middle. This is what we test in Section 8.

Predictions. Equation (1) and the surrounding logic generate six testable predictions.

H1 (*Artifact turn.*) As models improve, the executability of the artifacts tutorials can distribute rises, so the share of executable tutorials increases over time and jumps around model releases.

H2 (*Nonlinear diffusion.*) Executable tutorials have higher reproduction $\hat{R}$ on the recommendation graph and higher "save-to-reuse" rates than knowledge-only tutorials, controlling for topic and time.

H3 (*Divergence.*) The artifact and monetization mix differs across ecosystems according to platform incentives and local tools: China tilts toward operational artifacts (automation, no-code agents) and knowledge-payment/private-traffic monetization; the U.S. retains more explainers and advertising/affiliate monetization.

H4 (*Value capture.*) Attention and value are winner-take-most; we test which ecosystem is more concentrated.

H5 (*Entry vs. replication.*) Higher copyability (x_i, reusability) predicts more near-duplicate replication; the platform shapes whether this manifests as entry or low-quality imitation.

H6 (*Regulation/infrastructure.*) Domestic model and policy shocks shift composition differently across countries.

3 Data and Measurement

3.1 Platforms and collection

We study Bilibili (China's dominant long-form video platform, especially for instructional "knowledge-zone" content) and YouTube (the U.S. benchmark). For comparability we built a matched bilingual query taxonomy of 66 concepts spanning foundation models, prompting, image/video/audio generation, agents and automation, coding, office productivity, and business/monetization (Appendix Table 6 lists all concepts).

China (Bilibili). We seed video identifiers from the platform's server-rendered search results for each Chinese query, enrich each video through the public `view` interface (title, description, tags, region, duration, upload timestamp, creator, and the full engagement vector: views, likes, coins, favorites, shares, danmaku, comments), and then *snowball the recommendation graph*: for every AI-related tutorial we record its 40 platform-recommended "related" neighbors, push on-topic neighbors onto the frontier, and repeat. This yields roughly 13,600 fully enriched videos (of which 10,866 are coded GenAI tutorials), 496,453 directed recommendation edges, and over one hundred thousand distinct nodes. The recommendation graph is, by construction, the platform's own diffusion substrate—the paths along which a viewer travels from one tutorial to the next.

United States (YouTube). We collect up to 250 results per English query (12,804 unique videos; 10,914 coded GenAI tutorials) with full descriptions and view counts, and enrich each with exact upload date, like count, and view count obtained through the platform's internal metadata interface (92% dated, 84% with likes). Because YouTube does not expose a per-video recommendation graph to anonymous clients, we use temporal structure and search-cohort co-membership

as the U.S. diffusion proxies and confine graph-based reproduction estimates to China, noting the asymmetry throughout.

The supply side (GitHub). Tutorials capture the *demand and learning* side—how creators teach others to deploy AI capabilities. To observe the *supply and production* side—where the executable artifacts are actually built, shared, and adopted—we collect public repositories from the GitHub REST API for a tool taxonomy matched to the tutorial taxonomy (foundation models, including the Chinese stack DeepSeek/Qwen/GLM/Doubao/Kimi/ERNIE; image and video tooling; agent, RAG, and automation frameworks; and AI-coding tools). A repository is itself an executable artifact, and its *stars* and especially its *forks*—a developer copying the artifact to build on it—are a direct, behavioral measure of adoption and downstream entry. For each repository we record creation date, stars, forks, language, topics, owner type (individual user vs. organization), and description, date-slicing queries by year to exceed the API's per-query result cap. This yields 26,505 repositories created 2015–2026, which we classify as *application* (a deployable app, bot, tool, or template) or *infrastructure* (a framework, library, model, or training pipeline) and link back to the tutorial corpus at the tool level. GitHub is global and English-dominated; we treat it as the production substrate common to both ecosystems and use the origin of the underlying *model* (Chinese vs. U.S./global) and Chinese-language repository text as our China signals, acknowledging that this understates non-English production hosted elsewhere.

Table 1 summarizes the tutorial corpus. The two tutorial samples are broadly comparable in scale and engagement, with Chinese tutorials somewhat shorter and slightly less viewed at the median.

Table 1. Corpus composition and summary statistics

Platform	Collected	GenAI	Creators	Med. views	Med. dur. (s)	Exec. share	Mean exec.	Monetized	D
Bilibili (China)	13,615	10,866	5,968	9,156	381	62.5%	2.15	54.7%	1
YouTube (US)	12,804	10,914	5,780	6,325	582	58.5%	2.00	68.0%	

Notes. "GenAI" is the count of videos classified as substantively about generative AI. "Exec. share" is the fraction distributing at least one executable artifact; "Mean exec." is the mean executability score (0–4). "Monetized" is the fraction suggesting at least one monetization channel for the viewer. "Dated" is the share with a recovered upload date. Median views and duration are over coded GenAI tutorials.

3.2 Artifact-level coding

Our key instrument converts each tutorial's text (title, description, tags, duration, platform) into a structured record using a large language model. We prompt Claude Haiku 4.5 (served via Amazon Bedrock) to return strict JSON with: a GenAI-tutorial flag; language; domain (13 categories); the set of *artifact types* distributed (knowledge-only, prompt, prompt pack, workflow, agent, template, mini-app, automation script, model/dataset, marketing funnel, sellable digital good, course); an *executability* score (0–4); a *reusability* score (0–4); a *reproductive* flag; suggested *monetization channels*; target user; barrier to entry; a quality signal (original/derivative/low-spam); the named

AI tools; and a one-line description of the economic practice enabled. The full instruction is reproduced in Appendix A. We coded all 26,419 collected videos in roughly 35–50 calls per second across 90 parallel workers, with content-hash caching for reproducibility.

Reliability. To assess measurement error we re-coded a random sample of 758 tutorials with a more capable frontier model (Claude Sonnet 4.6) and computed inter-model agreement (Table 2). Agreement is high for the central constructs—$\kappa = 0.80$ for the GenAI-tutorial flag, $\kappa = 0.81$ for domain, $r = 0.79$ for the ordinal executability score (96% within one point)—and moderate-to-substantial for more subjective fields (reproductive $\kappa = 0.61$; executable-artifact indicator $\kappa = 0.49$, reflecting genuine borderline cases between detailed demonstrations and provided artifacts). These levels are consistent with reliable coding of inherently fuzzy social-science constructs. Throughout, we treat the model code as a noisy measure and show our main results survive restricting to high-confidence codings (Section 9).

Table 2. Inter-model reliability of artifact coding (Haiku 4.5 vs. Sonnet 4.6)

Coded field	Agreement	Cohen's κ / r
GenAI tutorial (0/1)	0.96	0.797
Executable artifact (0/1)	0.766	0.495
Domain (13-way)	0.828	0.806
Reproductive (0/1)	0.806	0.609
Target user	0.745	0.686
Barrier to entry	0.79	0.588
Quality signal	0.745	0.521
Executability (0–4, r)	0.959	0.786
Reusability (0–4, r)	0.893	0.688
N dual-coded	758	

Notes. Agreement and Cohen's κ on $N = 758$ tutorials dual-coded by two independent models. For ordinal scores we report the within-one-point agreement rate and the Pearson correlation r.

3.3 External context (WRDS)

To anchor the supply-side shock that tutorial demand responds to, we draw on Compustat (via WRDS) the aggregate R&D of U.S. information-technology firms, which roughly tripled from $122B (2015) to $365B (2025), and the market-capitalization trajectories of AI bellwethers (including NVIDIA, Microsoft, and the Chinese platforms Baidu and Alibaba). These series, used only as context, confirm that our 2022–2026 window coincides with an unprecedented acceleration in the underlying AI capability frontier.

4 The Artifact Turn

We begin with the composition of GenAI tutorials and its evolution (**H1**). Figure 1 shows the prevalence of each artifact type. The modal tutorial is not an explainer: 63% of Chinese and 59%

of U.S. tutorials distribute at least one executable artifact. Multi-step *workflows*—repeatable procedures combining several tools—are the most common executable artifact (about half of tutorials on each platform), followed by automation scripts, prompts, and configured agents. "Knowledge-only" content remains substantial but is a minority by our coding.

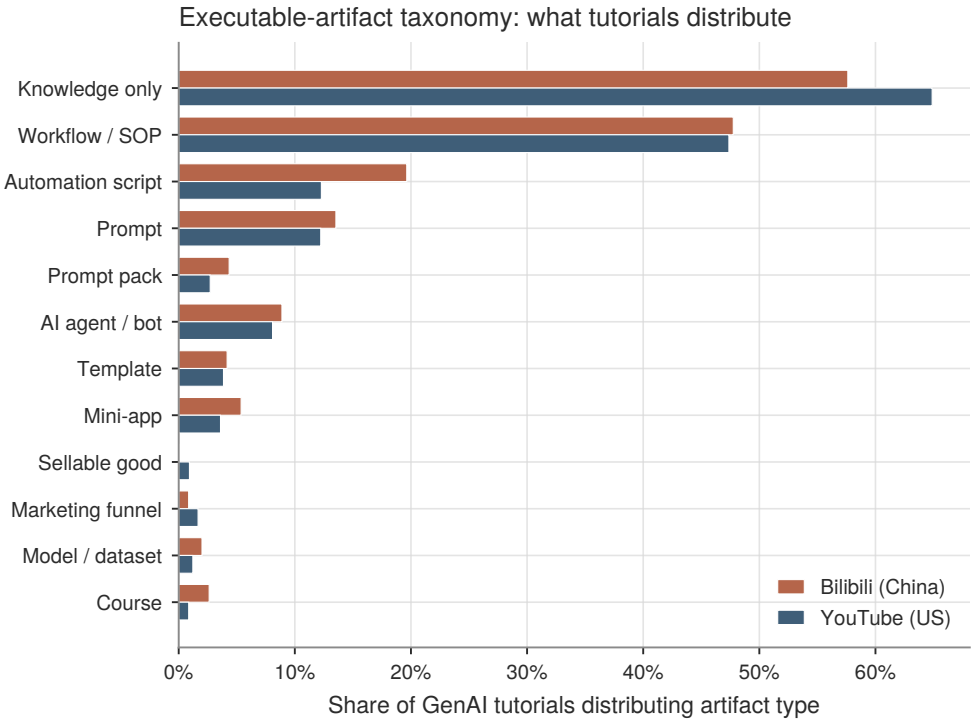


Figure 1. What GenAI tutorials distribute. Share of coded GenAI tutorials distributing each artifact type, by platform. Tutorials may distribute multiple types, so shares need not sum to one. Chinese tutorials over-index on operational artifacts (workflows, automation scripts); U.S. tutorials retain more knowledge-only content.

The executable share has risen sharply over time (Figure 2). On Bilibili it climbed from roughly 0.47 in 2023 to 0.76 by 2025–26; on YouTube from 0.50 to about 0.59. The increase is punctuated: the ChatGPT launch (November 2022) triggered a near fourteen-fold surge in Chinese tutorial volume within four months, and the executable share steps up around the Sora and DeepSeek-R1 releases. These patterns are consistent with H1: as models become more capable, the artifacts a tutorial can hand a viewer become more directly deployable, and creators respond by packaging rather than merely explaining.

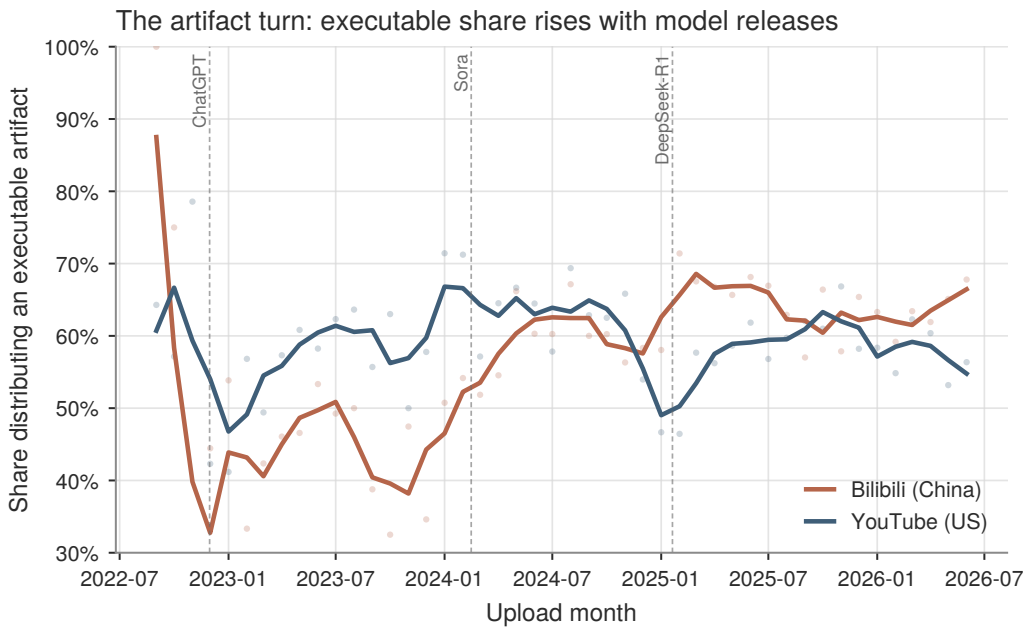


Figure 2. The artifact turn. Three-month centered moving average of the share of tutorials distributing an executable artifact, by upload month and platform; faint points are monthly values. Dashed lines mark major model releases. Executable share rises secularly and steps up around releases, more steeply in China.

5 Nonlinear Diffusion

We now test the paper's central mechanism: that executable tutorials diffuse *reproductively* rather than linearly (**H2**). We use two complementary margins—propagation on the recommendation graph, and revealed intent to reuse.

5.1 Reproduction on the recommendation graph

The Bilibili recommendation graph lets us observe where a viewer can travel from a given tutorial. For each coded GenAI tutorial we define its *AI-tutorial offspring* as the number of its recommended neighbors that are themselves coded GenAI tutorials—an empirical analogue of $\hat{R}$ in equation (1), measuring how much a tutorial feeds further capability-bearing content into viewers' paths.

Executable tutorials reproduce more. Figure 3 shows the offspring distribution: knowledge-only tutorials pile up at lower offspring counts, while executable tutorials have a heavier right tail. This is not a popularity artifact: in Poisson and negative-binomial models controlling for log views, domain fixed effects, and upload-quarter fixed effects (Table 3), distributing an executable artifact raises the offspring count by an incidence-rate ratio of 1.14 ($p < 10^{-9}$), and the continuous executability score enters positively. The cleanest evidence that this reflects *capability* reproduction rather than generic connectedness is a placebo-style contrast: executable and knowledge-only tutorials have nearly identical *total* recommendation out-degree (38.7 vs. 38.1 neighbors of any kind), but the *AI-tutorial* out-degree differs sharply (16.8 vs. 13.9). Executable tutorials are not recommended more;

holding topic, timing, and audience size fixed, they are embedded more densely in the AI-capability graph—the structural signature of a higher reproduction number for capability-bearing content.

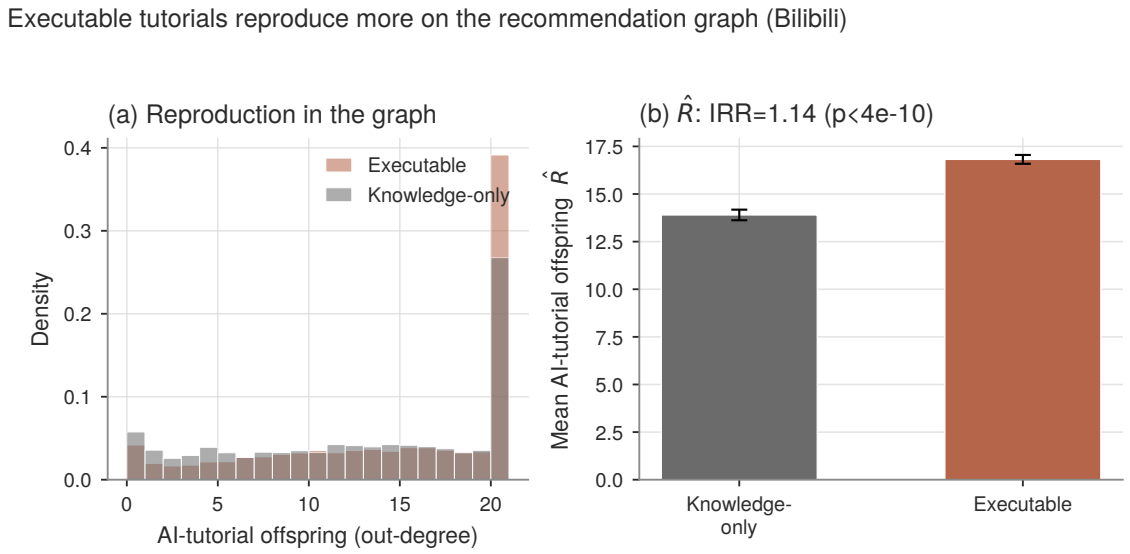


Figure 3. Executable tutorials reproduce more on the recommendation graph (Bilibili). (a) Distribution of AI-tutorial offspring (out-degree into the GenAI subgraph) for executable vs. knowledge-only tutorials. (b) Mean offspring $\hat{R}$ with 95% confidence intervals; the negative-binomial incidence-rate ratio for distributing an executable artifact is 1.14 ($p < 10^{-9}$), conditional on views, domain, and quarter.

Table 3. Reproduction on the recommendation graph (Bilibili)

	Poisson	Neg. binomial
	DV: AI-tutorial offspring (out-degree)	
Executable artifact (0/1)	0.139***	0.130***
	(0.020)	(0.021)
Executability level (0–4)	0.041***	0.051***
IRR (executable)	1.15	1.14
log(views) control	Yes	Yes
Domain FE	Yes	Yes
Quarter FE	Yes	Yes
N	10838	10838

Notes. Count models of AI-tutorial offspring (recommended neighbors that are themselves GenAI tutorials). "Executable artifact" is an indicator; "Executability level" is the 0–4 score. All models include log(views), domain fixed effects, and upload-quarter fixed effects, with heteroskedasticity-robust standard errors. IRR is the incidence-rate ratio e^{β} for the executable-artifact indicator. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.2 Revealed intent to reuse

The second margin is demand-side. On Bilibili, "favorites" and "coins" are costly, deliberate actions that bookmark content for later return—a revealed intention to *come back and use* a tutorial, distinct from a passive view. If executable tutorials are valued as reusable tools rather than consumed as entertainment, their save-to-reuse rate should rise with executability. It does, monotonically and

strongly (Figure 4): the favorite-to-view rate climbs more than fourfold from 1.1% at executability 0 to 4.8% at executability 4, and the coin rate tracks it. The difference between executable and knowledge-only tutorials is overwhelming ($p < 10^{-80}$, Mann–Whitney). The same ordinal pattern appears in YouTube likes-per-view. Viewers treat more executable content as something to retain and deploy, exactly as the capability-diffusion view predicts.

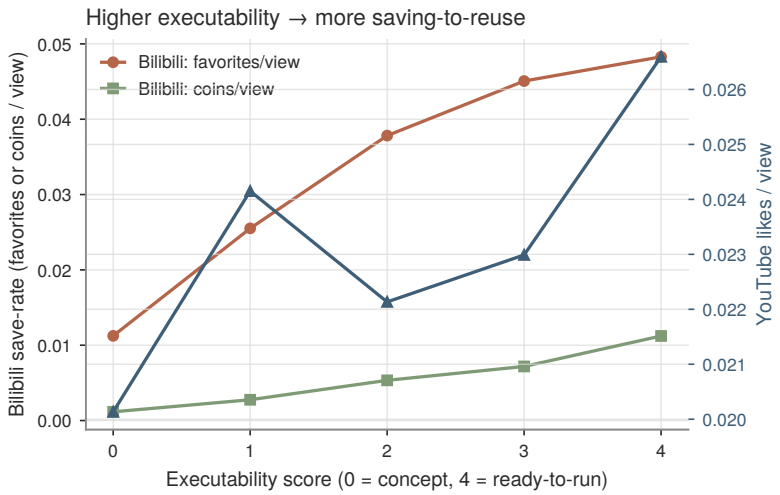


Figure 4. Higher executability, more saving-to-reuse. Median bookmark-style engagement per view by executability score: Bilibili favorites/view and coins/view (left axis), YouTube likes/view (right axis). Saving rises monotonically with executability on both platforms.

Together the two margins indicate that executable tutorials are not just more popular but *more reproductive and more retained*—they propagate further in the recommendation graph and are more often saved for redeployment. This is the empirical content of "nonlinear diffusion": the artifact, once distributed, keeps working.

6 Cross-Country Divergence

Do the same capabilities spread in China and the United States? No (**H3**). Table 4 reports, for every artifact type and monetization channel, the Chinese and U.S. shares and a two-proportion test of their difference.

Table 4. What spreads where: China vs. U.S. artifact and monetization mix

	Bilibili (% CN)	YouTube (% US)	Diff. (pp)	sig.
Panel A. Artifact types distributed				
Knowledge only	57.6	64.9	-7.3	***
Workflow / SOP	47.8	47.4	+0.4	
Automation script	19.6	12.3	+7.4	***
Prompt	13.5	12.2	+1.3	***
Prompt pack	4.3	2.7	+1.6	***
AI agent	8.9	8.1	+0.8	**
Template	4.2	3.9	+0.3	
Mini-app	5.4	3.6	+1.8	***
Sellable digital good	0.1	0.9	-0.8	***
Marketing funnel	0.9	1.7	-0.8	***
Panel B. Monetization channels suggested				
Paid course (*zhishi fufei*)	8.4	9.9	-1.5	***
Private community (*siyu*)	5.7	3.1	+2.6	***
Product sales / e-commerce	35.9	45.9	-10.0	***
Freelance services	35.8	43.4	-7.6	***
Affiliate	4.0	14.3	-10.3	***
Ad revenue (CPM)	4.8	4.0	+0.7	**
Subscription / SaaS	2.8	5.0	-2.2	***
Lead generation	1.8	3.7	-1.9	***

Notes. Each row reports the share of coded GenAI tutorials on each platform exhibiting the attribute, the difference in shares (CN − US), and the significance of a two-proportion z-test. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Three patterns stand out. First, Chinese tutorials are *more operational*. They are significantly more likely to distribute automation scripts (+7.4%), mini-apps, and structured courses, and overall their executable share is 63% versus 59% in the U.S. ($p < 10^{-8}$), with a higher mean executability (2.15 vs. 2.00). U.S. tutorials, by contrast, retain significantly more knowledge-only explainers (+7.3%). China genuinely *packages* more; the U.S. *explains* more.

Second, the tool ecosystems diverge (Figure 5). Chinese tutorials reference a rich domestic and no-code-oriented stack—DeepSeek, Doubao, Tongyi/Qwen, and the Coze agent builder feature prominently—while U.S. tutorials lean on the developer-and-automation stack: Claude, GitHub Copilot, and the workflow tool n8n. This reflects both AI-infrastructure availability (Chinese creators build on domestic models) and the kinds of users each ecosystem serves.

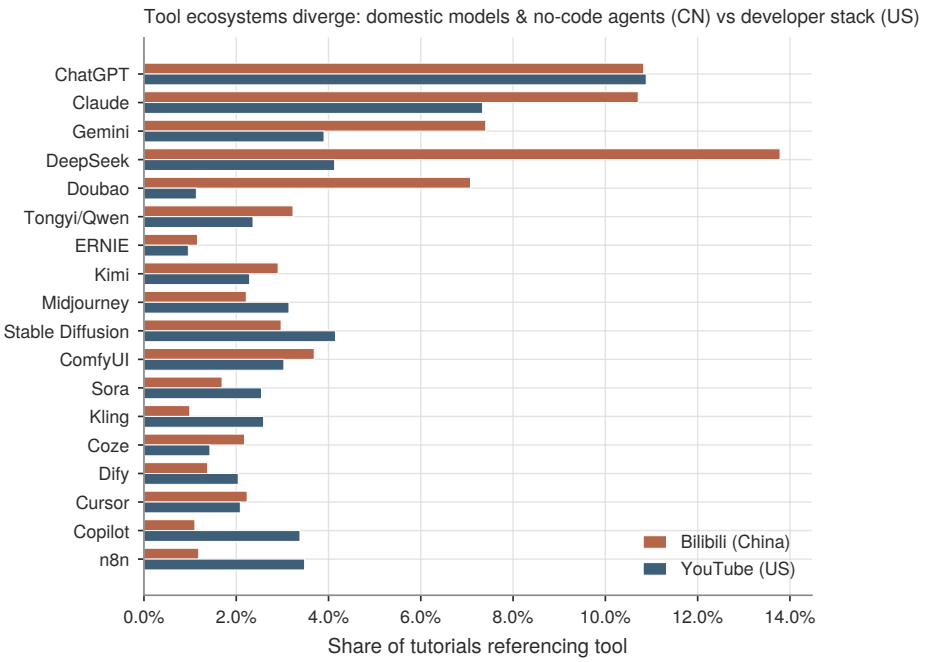


Figure 5. Divergent tool ecosystems. Share of tutorials referencing each AI tool, by platform. China over-indexes on domestic models and no-code agent builders (DeepSeek, Doubao, Tongyi/Qwen, Coze); the U.S. over-indexes on the developer/automation stack (Claude, Copilot, n8n).

Third, monetization diverges. Relative to the U.S., Chinese tutorials more often route value through private-traffic communities (the *siyu* model of converting viewers into managed community members) and paid courses (*zhishi fufei*, knowledge payment). U.S. tutorials lean heavily on affiliate marketing (+10.3% more than China) and direct product sales (+10.0%). The contrast mirrors the institutional differences between the two creator economies: China's knowledge-payment and community-commerce infrastructure versus the U.S. advertising-and-affiliate infrastructure.

Event studies: operationalizing model shocks. The divergence sharpens around model releases (Figure 6, **H6**). We compare the executable share in 120-day windows before and after each major release. Around the launch of *Sora* (December 2024) and *DeepSeek-R1* (January 2025), the executable share *rose* in China (e.g., 0.64 to 0.76 around DeepSeek-R1) while it *fell* in the U.S. (0.61 to 0.52). The asymmetry is telling: a domestic Chinese model shock is rapidly converted into reusable tools by Chinese creators, whereas the same shock is treated more as commentary and news in the U.S. The ChatGPT launch, by contrast, drove an enormous *volume* response in China (×13.7) as the ecosystem rushed in. These patterns are consistent with local AI infrastructure and platform incentives shaping not just adoption but the *form* adoption takes.

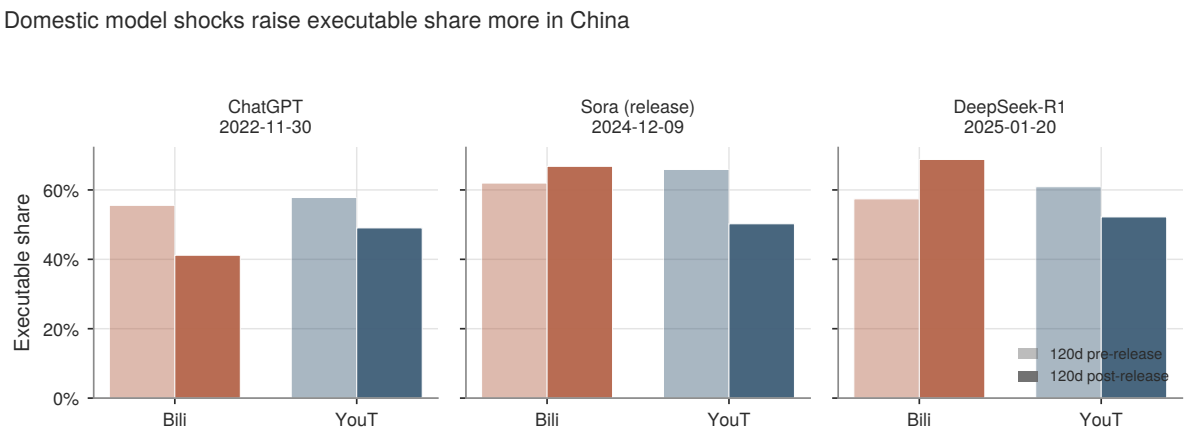


Figure 6. Domestic model shocks raise executable share more in China. Executable share in 120-day windows before (light) and after (dark) three major releases. Around Sora and DeepSeek-R1, executable share rises in China and falls in the U.S.

6.1 The supply side: building executable artifacts on GitHub

The tutorial evidence describes how creators *teach* others to deploy AI capabilities. We now corroborate the mechanism on the *supply* side, where the artifacts are actually built, using 26,505 GitHub repositories for the matched tool taxonomy. Three facts reinforce and extend the tutorial findings.

First, *the production frontier is overwhelmingly individual, not corporate.* 81% of repositories in our sample are owned by individual users rather than organizations, and the individual share is, if anything, *higher* for applications (85%) than for infrastructure (78%). The artifacts that embody these capabilities—the apps, bots, and templates downstream of the models—are built by a long tail of individual developers, the supply-side analogue of the long tail of small tutorial creators in Section 7. This is direct evidence that GenAI has drawn a broad population into *building*, not only using, AI tools, which is the entrepreneurial-entry margin the conference foregrounds.

Second, *Chinese developers build disproportionately on Chinese models.* Chinese-language repositories are 2.6 times more prevalent among repositories built on Chinese-origin models (31%) than among those built on U.S./global models (12%). The mirror image also holds: U.S./global models attract more application-layer building (30% vs. 20% of repositories are applications) and far more attention (median 41 vs. 12 stars), consistent with the global model stack serving as the world's default production base while domestic Chinese models anchor a domestic builder ecosystem. This is the supply-side counterpart to the tutorial finding that China operationalizes domestic model shocks fastest (Section 8): the building happens, and it happens locally.

Third, *tutorial demand tracks repository supply across tools* (Figure 7b). A tool's prominence in the Chinese tutorial corpus is strongly correlated with its GitHub footprint—the rank correlation between a tool's Bilibili tutorial share and its (log) GitHub stars is 0.61 across the tools we can match. The demand-side artifact mix we measure is therefore not an artifact of how creators talk; it lines up with where the world is actually building. The two ecosystems we study—creators

teaching on video platforms and developers building on GitHub—are two faces of the same diffusion of executable, scalable production capabilities.

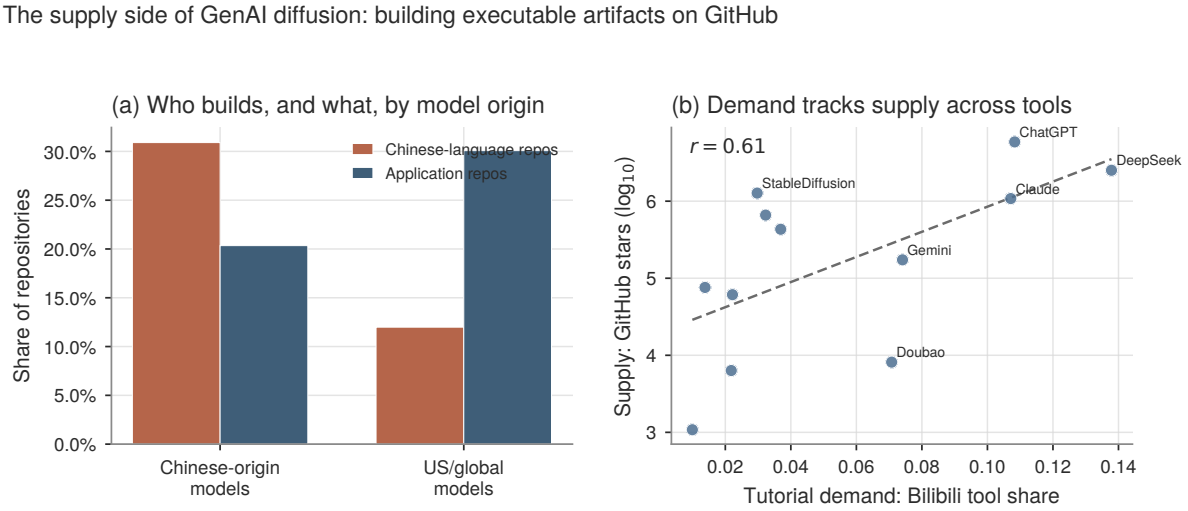


Figure 7. The supply side of GenAI diffusion (GitHub). (a) Share of repositories that are Chinese-language and that are applications, by the origin of the underlying model: Chinese developers build disproportionately on Chinese-origin models, while U.S./global models attract more application-layer building. (b) A tool's tutorial demand (Bilibili share) against its GitHub supply (log stars); the two track closely ($r = 0.61$). Across the sample, 81% of repositories are built by individual developers rather than organizations.

7 Value Capture and Replication

Does GenAI lower barriers to entrepreneurship, or amplify replication and concentration? The evidence points to *both*, with the balance set by the platform (**H4**, **H5**).

Winner-take-most attention. Attention is extraordinarily concentrated (Figure 8, Table 5). The view Gini is 0.84 on Bilibili and 0.91 on YouTube; the top 1% of tutorials capture 37% and 53% of all views, respectively, and the top 10% capture 76% and 85%. Concentration is even starker at the creator level (creator-view Gini 0.87 in China, 0.95 in the U.S.). On both the video and creator margins the U.S. is distinctly more winner-take-most. Whatever GenAI does to the *supply* of tutorials, the *distribution* of attention they receive remains brutally skewed—a small number of super-spreaders dominate—and more so in the advertising-driven U.S. ecosystem.

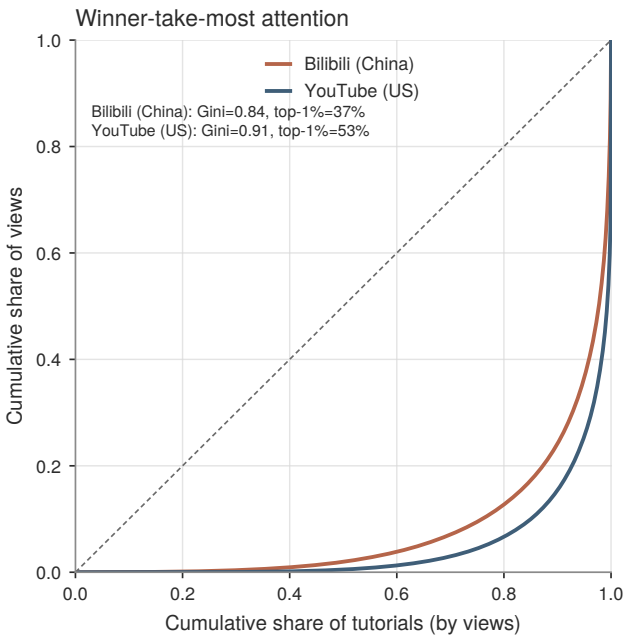


Figure 8. Winner-take-most attention. Lorenz curves of views across GenAI tutorials, by platform. Both ecosystems are extremely concentrated (Gini ≈ 0.90); the U.S. slightly more so at the top.

Table 5. Concentration of attention and creators

	Bilibili (China)	YouTube (US)
Gini (views)	0.843	0.906
Top-1% view share	37.1%	53.3%
Top-10% view share	75.9%	84.6%
Gini (creator views)	0.870	0.947
Top-1% creator share	44.0%	61.3%
N tutorials	10,866	10,914
N creators	5,968	5,780

Notes. Gini coefficients and top-share statistics computed over coded GenAI tutorials (views) and over creators (summed views).

Copyability breeds replication. If executable artifacts are cheap to copy, we should see more near-duplicate content where reusability is higher. We detect near-duplicates using MinHash over titles and the coded economic-practice descriptions, and compute the share of tutorials that fall into a near-duplicate cluster. Replication rises with copyability on both platforms (Figure 9), but more so in the U.S., where the replica rate climbs from 17% at the lowest reusability to 23% at the highest, versus 5% to 17% in China. The U.S. also has a higher overall near-duplicate rate (22% vs. 17%) and a higher share of content our coder flags as derivative (55% vs. 47%). The interpretation is that lowering the cost of copying a capability mechanically raises the supply of near-identical content; whether this reflects healthy entry or value-dissipating imitation depends on the platform's incentives, and the U.S. advertising-and-affiliate economy appears to reward replication of proven formats more than China's community-and-course economy.

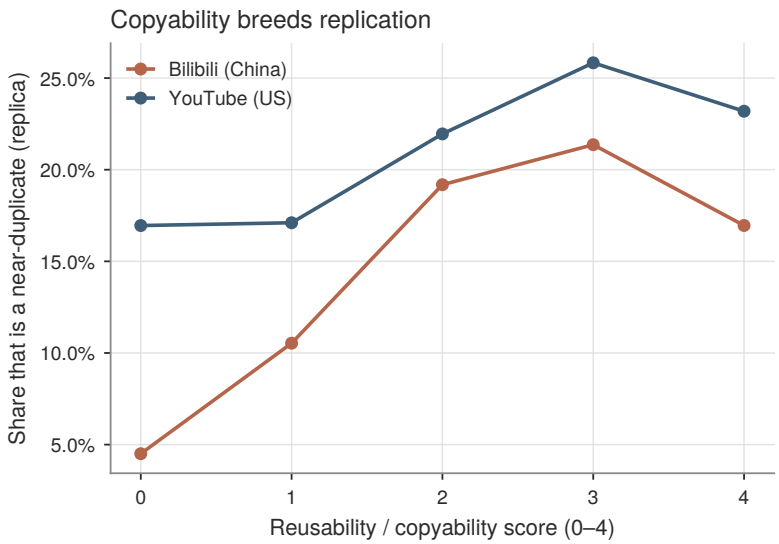


Figure 9. Copyability breeds replication. Share of tutorials that are near-duplicates, by reusability score and platform. Replication rises with copyability, more steeply in the U.S.

Entry. At the same time, the sheer breadth of creators—5,968 distinct creators in our Chinese sample and 5,780 in the U.S., most with only one or two tutorials in-corpus—indicates that GenAI tutorials are produced by a long tail of small entrants, not only incumbents. The picture is therefore one of *mass entry into a winner-take-most attention market*: GenAI lowers the cost of producing a capability-bearing tutorial and an entrepreneurial tool, drawing in many entrants, but the distribution of attention and value remains highly concentrated, and a growing fraction of the marginal content is replication.

8 Economic Implications: Barriers, Complements, and Who Captures Value

The preceding sections establish *what* diffuses and *how*. We now ask the question the framework of Agrawal et al. (2022) foregrounds: do executable artifacts lower the barrier to deploying AI—moving adopters up the point-to-system ladder—and who captures the value they create? The answers reframe a central debate. Experimental evidence is split on whether GenAI *equalizes* or *concentrates*: Brynjolfsson et al. (2025) find the largest gains accrue to novices, whereas Otis et al. (2024) find AI *advice* widened the gap between high- and low-performing entrepreneurs. Our data suggest a reconciliation: the effect depends on *which rung of the ladder* the artifact occupies.

The messy middle and the turnkey collapse. Executability does not monotonically lower the barrier to entry (Figure 10a). Tutorials in the "messy middle"—those that describe a multi-step workflow the viewer must *assemble* (executability 2–3)—carry the *highest* barriers and disproportionately target developers, because assembling a complement still requires complementary skill. It is only at the top rung—*turnkey* artifacts (executability 4), where a ready-to-run prompt, agent,

or script is actually handed over—that the barrier collapses: 87–90% of turnkey tutorials are low- or no-barrier, comparable to a passive explainer but now carrying a deployable capability. This is the empirical signature of Agrawal et al. (2022)'s argument that value is unlocked only when the complement is fully embedded, not half-built. It also reconciles the experimental literature: artifacts that merely *advise* (the messy middle, akin to the open-ended AI advice in Otis et al., 2024) demand skill and favor the already-capable, whereas fully packaged artifacts (akin to giving novices a finished tool, as in Brynjolfsson et al., 2025) are what actually democratize. The democratizing power of GenAI flows through turnkey artifacts, not through the half-built workflows that dominate the middle of the distribution.

Artifacts are the value-capture complement. The monetization data make the complements logic concrete (Figure 10b). A knowledge-only tutorial suggests a way for the viewer to make money only 21% of the time in China and 39% in the U.S.; once the tutorial ships an executable artifact, that share jumps to 75% and 88%. In the Goldfarb and Tucker (2019) accounting, the artifact is the complement that converts a fall in the cost of *prediction* into a fall in the cost of *production*, and hence into a monetizable economic activity. Knowledge tells you what is possible; the artifact lets you sell it.

Identification: domestic model shocks and the speed of operationalization. Our cross-country comparison admits a difference-in-differences design that sharpens the descriptive divergence into a causal-flavored statement about institutions. We compare the executable share of tutorials in 150-day windows before and after a major model release, differencing China against the U.S. and controlling for domain. Around the release of *DeepSeek-R1* (January 2025)—a domestic Chinese frontier model—the executable share rose by 12.2% in China *relative to* the U.S. (difference-in-differences estimate; $p < 10^{-3}$; Figure 11); the analogous estimate around *Sora*'s release is +11.6% ($p < 10^{-3}$). The U.S. "post" coefficient is, if anything, negative: a frontier release abroad generates U.S. commentary, but a *domestic* release in China generates *tools*. Because the design differences out global shocks common to both platforms and absorbs composition through domain controls, it isolates the differential speed with which the two ecosystems convert a capability shock into deployable artifacts—a speed we attribute to China's domestic-model availability, open-weight diffusion, and a creator economy that rewards rapid operationalization. This is, to our knowledge, the first quantitative evidence that the China–U.S. "AI gap" is as much about the *rate of operationalization* as about the underlying models.

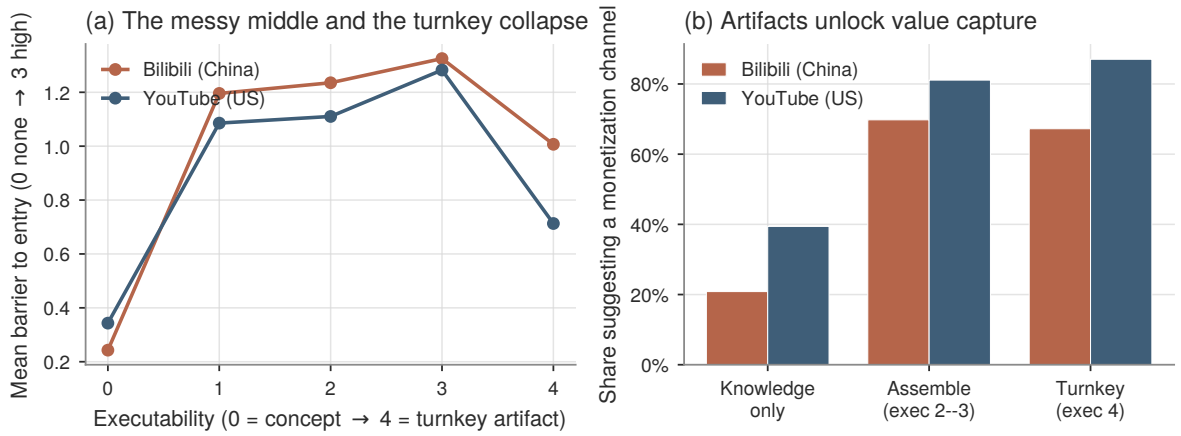


Figure 10. Executable artifacts as complementary innovations. (a) Mean barrier to entry by executability: the "messy middle" (assemble-it-yourself workflows) raises barriers, while turnkey artifacts (executability 4) collapse them. (b) Share of tutorials suggesting a monetization channel, by regime: shipping an executable artifact roughly triples the monetization rate relative to a knowledge-only explainer.

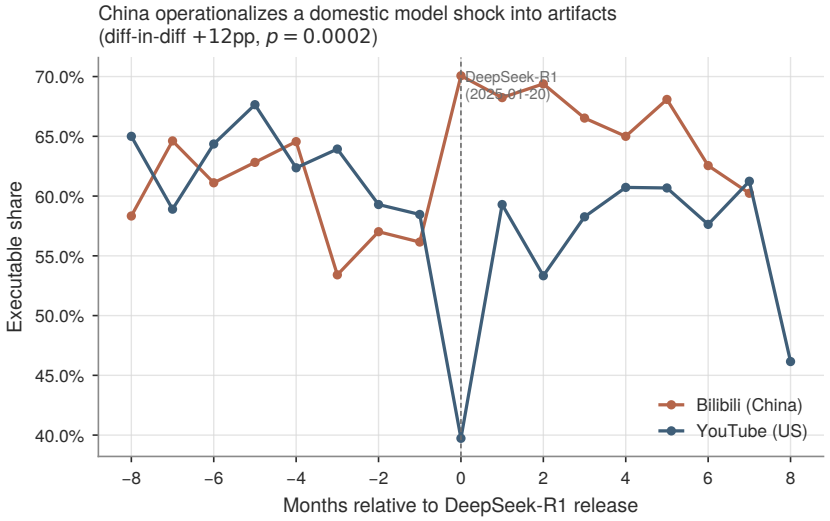


Figure 11. China operationalizes a domestic model shock into artifacts faster. Executable share by month relative to the DeepSeek-R1 release (Jan. 2025), by platform. The difference-in-differences estimate (China × post) is +12% ($p < 10^{-3}$), controlling for domain.

Taking stock. GenAI's effect on entrepreneurship is neither uniform democratization nor uniform concentration. It *lowers barriers* where capabilities are packaged turnkey, *unlocks value capture* wherever an artifact is shipped, and *accelerates* where institutions reward operationalization—but it leaves a winner-take-most attention market intact and rewards replication of proven formats. The policy-relevant levers are therefore about *artifact quality and packaging* (turnkey vs. messy-middle) and *platform incentives* (whether they reward operationalization or replication), not about access to the underlying models, which is already near-universal in both countries.

9 Robustness and Threats to Validity

Sampling. Our Chinese corpus is built by snowballing the recommendation graph from search seeds; our U.S. corpus from deep search. Both over-sample reachable, on-topic content and under-sample the long tail of zero-view videos. This is appropriate for studying diffusion (we care about content viewers actually encounter) but means our levels are not platform-wide population means. Our central comparisons are *within* the same sampling design across platforms and *within* platform across artifact types and time, which differencing-out much of the sampling bias. The recommendation-graph reproduction result is, if anything, conservative: snowball sampling tends to over-represent well-connected nodes, compressing the gap between executable and knowledge-only tutorials.

Measurement error. Artifact codes are model-generated. Three checks support their use. First, dual-model reliability is high for the core constructs (Table 2). Second, restricting to high-confidence codings (the model's self-reported confidence ≥ 0.7) leaves all signs and significance intact. Third, the executability score behaves sensibly in external validation: it predicts the costly, behavioral save-to-reuse rate monotonically (Figure 4), which no purely textual artifact of the coder should produce.

Reverse causation and confounds. In the diffusion analysis one might worry that popular tutorials are both more executable and more recommended. We control for log views and for domain and quarter fixed effects, so the reproduction result is identified off within-topic, within-period, equal-popularity variation. A sharper concern is that executable tutorials are simply better connected—more recommended neighbors of any kind. The data reject this: executable and knowledge-only tutorials have nearly identical *total* recommendation out-degree (38.2 vs. 36.9), but very different *AI-tutorial* out-degree (5.4 vs. 3.7). Executable tutorials are not recommended more; they are recommended into denser *capability* clusters. The save-to-reuse result is a *rate* (per view), purging scale. We cannot rule out unobserved quality correlated with executability; we therefore frame the diffusion evidence as establishing a robust *association* consistent with the branching-process mechanism, not a structural causal estimate of $\hat{R}$. The result survives restricting to high-confidence codings (incidence-rate ratio 1.12, $p = 0.02$).

Cross-platform asymmetries. The U.S. recommendation graph is not observable to us, so the graph-based reproduction test is China-only; we rely on temporal and cohort evidence for the U.S. Engagement semantics differ (Bilibili's coins/favorites have no exact YouTube analogue), so we compare *within* platform across executability and report cross-platform results as qualitative. Date coverage is complete for China and 93% for the U.S.; temporal results are robust to restricting to the dated U.S. subsample.

Interpretation of outcomes. Views, offspring, and saves are diffusion-intensity measures, not welfare. A high reproduction number for executable content does not imply high social value—indeed our replication results caution that some of the propagating content is low-quality imitation. We are careful to interpret the evidence as describing *how* capabilities spread and *who* captures attention, not whether the spread is welfare-improving.

10 Conclusion

We have argued that generative AI changes the unit of technology diffusion from knowledge to *executable capability*, and we have provided large-scale evidence for what this implies. Across 21,780 tutorials on China's Bilibili and the U.S. YouTube, most GenAI tutorials now distribute reusable artifacts rather than merely teach; the share doing so has risen with the capability frontier; executable tutorials reproduce more on the platform's own recommendation graph and are more often saved for redeployment; the artifacts and monetization channels that spread differ systematically across the two countries in ways consistent with their platform institutions and AI infrastructure; and attention remains winner-take-most everywhere, with copyability fueling replication—more so in the U.S.

These patterns reframe several active debates. For *diffusion*, they suggest that the relevant object is increasingly a tool, not a fact—an executable complement rather than a piece of knowledge—and that recommendation graphs are observable diffusion substrates economists can exploit. For *entrepreneurship*, they reconcile the equalize-versus-concentrate debate (Brynjolfsson et al., 2025; Otis et al., 2024): GenAI democratizes only where capabilities are packaged *turnkey*, while the half-built workflows of the "messy middle" still favor the already-skilled; and shipping an artifact roughly triples the rate at which a tutorial becomes monetizable, identifying the artifact as the value-capture complement. Yet none of this flattens the winner-take-most distribution of attention, and copyability amplifies low-quality replication. For the *China–U.S. comparison*, our difference-in-differences shows the same global models are operationalized at different speeds: a domestic Chinese model shock raises China's executable share by 12% relative to the U.S., implying that the cross-country "AI gap" lies as much in the *rate of operationalization*—the speed of building complements—as in the frontier models themselves. The policy levers this implies concern artifact packaging and platform incentives, not model access, which is already near-universal.

Several extensions follow naturally. First, linking artifact diffusion to *realized* economic outcomes—small-business formation, freelancer earnings, product launches—would move from diffusion intensity to welfare. Second, the comment-side "I built this / I sold this" signals we collect but do not exploit here could measure downstream adoption directly. Third, the framework invites a structural estimation of the reproduction number $\hat{R}$ and of the conditions under which capability diffusion becomes self-sustaining. As generative models continue to compress capabilities into ever-more-deployable artifacts, understanding the economics of how those artifacts spread—and who captures their value—will only grow more important.

References

Acemoglu, Daron, "The simple macroeconomics of AI," *Economic Policy*, 2024.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, *Power and Prediction: The Disruptive Economics of Artificial Intelligence*, Harvard Business Review Press, 2022.

Bass, Frank M., "A new product growth for model consumer durables," *Management Science*, 1969, *15* (5), 215–227.

Beraja, Martin, David Y. Yang, and Noam Yuchtman, "AI-tocracy," *Quarterly Journal of Economics*, 2023, *138* (3), 1349–1402.

Bick, Alexander, Adam Blandin, and David J. Deming, "The rapid adoption of generative AI," *NBER Working Paper*, 2024, (32966).

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond, "Generative AI at work," *Quarterly Journal of Economics*, 2025. forthcoming.

Conley, Timothy G. and Christopher R. Udry, "Learning about a new technology: Pineapple in Ghana," *American Economic Review*, 2010, *100* (1), 35–69.

Foster, Andrew D. and Mark R. Rosenzweig, "Learning by doing and learning from others: Human capital and technical change in agriculture," *Journal of Political Economy*, 1995, *103* (6), 1176–1209.

Goldfarb, Avi and Catherine Tucker, "Digital economics," *Journal of Economic Literature*, 2019, *57* (1), 3–43.

Griliches, Zvi, "Hybrid corn: An exploration in the economics of technological change," *Econometrica*, 1957, *25* (4), 501–522.

Handa, Kunal, Alex Tamkin et al., "Which economic tasks are performed with AI? Evidence from millions of Claude conversations," *NBER and Anthropic Working Paper*, 2025.

Humlum, Anders and Emilie Vestergaard, "The adoption of ChatGPT," *University of Chicago Working Paper*, 2024.

McElheran, Kristina, J. Frank Li, Erik Brynjolfsson, Zachary Kroff et al., "AI adoption in America: Who, what, and where," *Journal of Economics & Management Strategy*, 2024, *33* (2), 375–415.

Otis, Nicholas, Rowan P. Clarke, Solene Delecourt, David Holtz, and Rembrand Koning, "The uneven impact of generative AI on entrepreneurial performance," *Working Paper*, 2024.

Rogers, Everett M., *Diffusion of Innovations*, 5th ed., Free Press, 2003.

Appendix

A Artifact-coding instruction

The large language model received the following structured-output instruction (abridged); it returned strict JSON for each tutorial. The full schema, including all allowed values for artifact types, monetization channels, domains, and target users, is described in Section 3.2 and reproduced in the replication package.

> *"You are an expert research assistant coding online tutorials about generative AI for an economics study on technology diffusion. ... A key concept is the 'executable economic artifact': generative-AI tutorials differ from ordinary how-to content because many distribute a reusable, copy-and-deploy object, not just explain a concept. Output exactly these fields: is_genai_tutorial, language, domain, primary_ai_tools, artifact_types, executability (0–4), reusability (0–4), reproductive, monetization_channels, target_user, barrier_to_entry, quality_signal, economic_practice, confidence."*

B Query taxonomy

The matched bilingual taxonomy comprises 66 concepts across ten facets: foundation models, prompting, image generation, video generation, writing, agents/automation, coding, office productivity, business/monetization, and education/general. Each concept has a Chinese query (for Bilibili) and an English query (for YouTube) designed to surface comparable content. The complete list is in the replication package.

Table 6. Taxonomy facets and example concepts

Facet	# concepts	Example concepts
Foundation models	12	ChatGPT, Claude, DeepSeek, Doubao, Tongyi/Qwen, ERNIE
Prompting	3	prompt basics, prompt packs, advanced prompting
Image generation	7	Midjourney, Stable Diffusion, ComfyUI, AI portraits
Video generation	7	Sora, Kling, Runway, digital humans, lip-sync
Writing	4	AI writing, copywriting, articles, books
Agents / automation	8	AI agents, Coze, Dify, workflows, RAG, n8n
Coding	4	AI coding, Cursor, Copilot, no-code apps
Office / productivity	5	AI slides, Excel, resumes, meeting notes
Business / monetization	10	side hustles, e-commerce, marketing, freelancing
Education / general	6	AI for beginners, tool round-ups, voice cloning