AI Governance

Information Consumption and Production

Dennis Zhang

Washington University in St. Louis

Singapore
November 2026

What is AI governance?

The rules, institutions, and organizational choices that shape how AI is developed and used

- **Objectives:** Which outcomes should AI systems improve?
- **Decisions:** What should we delegate, and under what conditions?
- **Accountability:** Who evaluates the effects and bears responsibility?

An economic perspective: how do AI systems change behavior, incentives, and outcomes for society?

AI reshapes information consumption and production

Information consumption

AI selects, summarizes, and explains information.

- What people discover and attend to
- What they believe and learn
- How they interact with others

AI reshapes information consumption and production

Information consumption

AI selects, summarizes, and explains information.

- What people discover and attend to
- What they believe and learn
- How they interact with others

Information production

AI lowers the cost of creating content and software.

- Who can produce and enter
- Which skills remain valuable
- Who captures attention and income

AI reshapes information consumption and production

Information consumption
AI selects, summarizes, and explains information.

- What people discover and attend to
- What they believe and learn
- How they interact with others

Information production
AI lowers the cost of creating content and software.

- Who can produce and enter
- Which skills remain valuable
- Who captures attention and income

Consumption allocates attention. Attention shapes incentives to produce.

These changes motivate our study of AI governance: what are the consequences for people and society?

Related work and my papers today

Related research

Information consumption

- Beliefs: Costello et al. (2024), *Science*
- Learning: Bastani et al. (2025), *PNAS*
- Mental health: Heinz et al. (2025), *NEJM AI*

Information production

- Productivity: Brynjolfsson et al. (2025), *QJE*
- Work and income: Hui et al. (2024), *Organization Science*

Related work and my papers today

Related research

Information consumption

- Beliefs: Costello et al. (2024), *Science*
- Learning: Bastani et al. (2025), *PNAS*
- Mental health: Heinz et al. (2025), *NEJM AI*

Information production

- Productivity: Brynjolfsson et al. (2025), *QJE*
- Work and income: Hui et al. (2024), *Organization Science*

My papers today

Consumption **1. The Impact of AI Search on the Online Content Ecosystem**
Peibo Zhang, Ruomeng Cui, and Dennis J. Zhang

Production **2. When Code Gets Cheap**
Bing Bai, Chenshan Hu, Yuval Salant, and Dennis J. Zhang

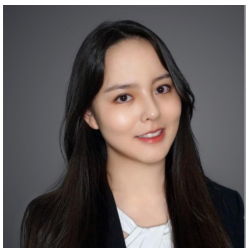
The Impact of AI Search on the Online Content Ecosystem

Evidence from Google and Reddit

Peibo Zhang[1] Ruomeng Cui[1] Dennis J. Zhang[2]

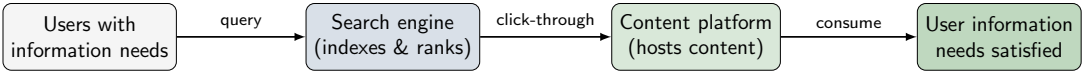
[1]Goizueta Business School, Emory University

[2]Olin Business School, Washington University in St. Louis





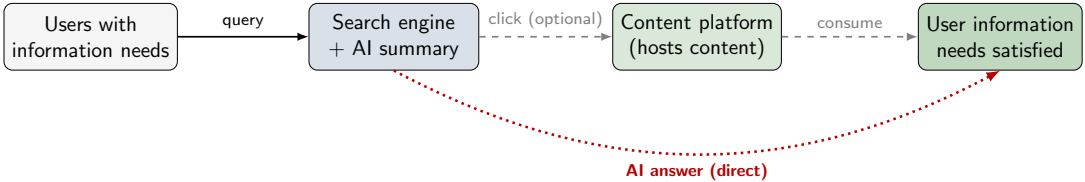
Search engines have long complemented content



- Search drives **discovery**: users reach content through organic links.
- Search and content reinforce each other in a two-sided value chain.

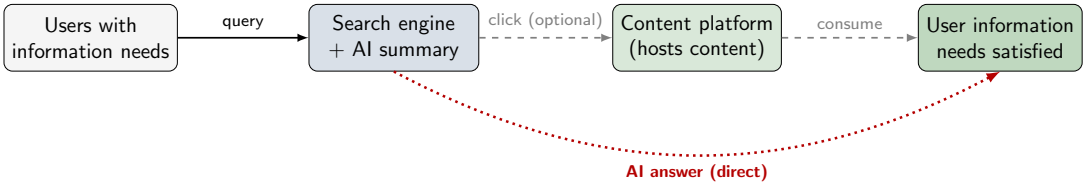
Generative AI changes the relationship

With AI search



Generative AI changes the relationship

With AI search



The governance concern

- AI can answer a question without sending the user to its source.
- This changes traffic, publisher incentives, and the supply of future content.

What does the literature say?

Prior evidence finds substitution in factual information settings

- ChatGPT and Stack Overflow participation.
  Burtch et al. (2023), working paper.
- AI Overviews and Wikipedia referrals.
  Khosravi and Yoganarasimhan (2026), working paper.

What does the literature say?

Prior evidence finds substitution in factual information settings

- ChatGPT and Stack Overflow participation.
  Burtch et al. (2023), working paper.
- AI Overviews and Wikipedia referrals.
  Khosravi and Yoganarasimhan (2026), working paper.

Two questions remain

1. Does AI search affect **other content types** differently?
2. Does a **conversational interface** change the effect?

Research question

After AI enters search,

does search remain a complement to content,

or become a substitute?

Two sub-questions

1. Does the effect vary with the **type of content**?
   Factual information versus experience and discussion.
2. Does the effect vary with the **type of AI interface**?
   Static summaries versus conversational search.

Empirical setting: Google and Reddit

Search engine: Google

- AI Overviews place answers within conventional search.
- AI Mode supports follow-up conversations.

Content platform: Reddit

- Topic-specific communities with public discussion.
- A mix of factual and experience-oriented content.

One setting allows us to compare both content types and AI interfaces.

Two Google AI search products

Date	Event	Interface
May 14, 2024	AI Overviews: US launch	Static summary
Aug. 15, 2024	**AI Overviews: international rollout**	**Static summary**
May 2025	AI Mode: US early access	Conversational
Aug. 21, 2025	**AI Mode: international expansion**	**Conversational**

AI Overviews
A summary above organic results. Users can continue to sources for more detail.

AI Mode
A chat interface within search. Follow-up questions happen within the conversation.

Illustration of Google AI Overviews

Without AI Overview

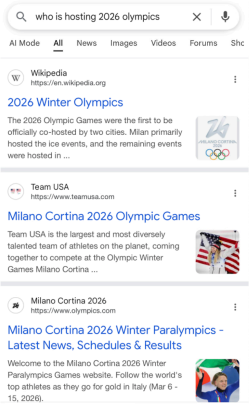
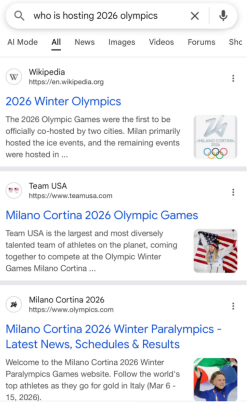


Illustration of Google AI Overviews

Without AI Overview



With AI Overview

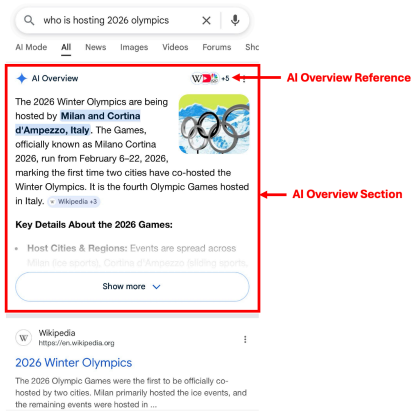
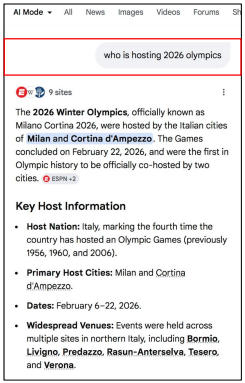


Illustration of Google AI Mode

1. Initial question in AI Mode



2. Follow-up question in AI Mode

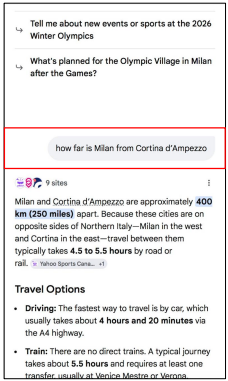


Illustration of Reddit



Illustration of Reddit

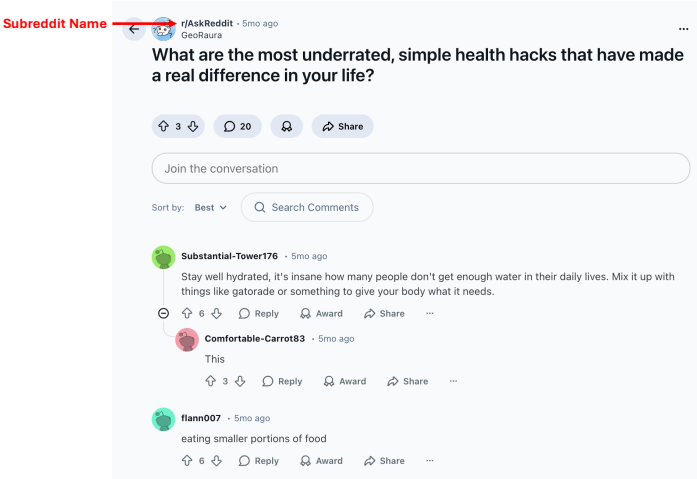


Illustration of Reddit



Illustration of Reddit



Identification: Google's content moderation policy

Google indexes both SFW and NSFW Reddit content. Only SFW is eligible for AI Overviews and AI Mode in the study's design.

Treatment: SFW subreddits
70,683 communities

- Indexed by Google
- **Eligible** for AI search

Control: NSFW subreddits
34,329 communities

- Indexed by Google
- **Excluded** from AI search

Identification: Google's content moderation policy

Google indexes both SFW and NSFW Reddit content. Only SFW is eligible for AI Overviews and AI Mode in the study's design.

Treatment: SFW subreddits

70,683 communities

- Indexed by Google
- **Eligible** for AI search

Control: NSFW subreddits

34,329 communities

- Indexed by Google
- **Excluded** from AI search

Why this comparison is useful

Same platform and common platform-wide shocks. Labels are fixed in December 2023. A manual audit finds essentially no NSFW citations in AI Overviews.

Data: a balanced panel of Reddit communities

Source: Arctic Shift Reddit archive.

Public subreddits with ≥ 10 submissions in Dec. 2023	**105,012**
SFW / NSFW communities	70,683 / 34,329
Period	Jan. 1, 2024–June 30, 2025
Days	547
Subreddit-day observations	**57,441,564**
Outcomes	Comments and comment authors per day

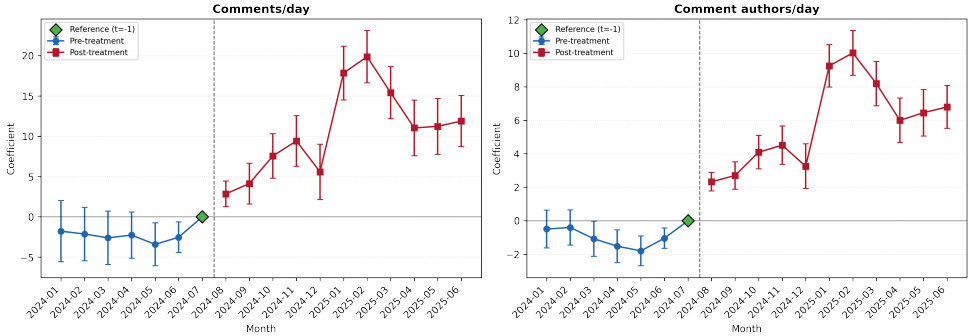
Data: a balanced panel of Reddit communities

Source: Arctic Shift Reddit archive.

Public subreddits with ≥ 10 submissions in Dec. 2023	**105,012**
SFW / NSFW communities	70,683 / 34,329
Period	Jan. 1, 2024–June 30, 2025
Days	547
Subreddit-day observations	**57,441,564**
Outcomes	Comments and comment authors per day

- A **comment author** posts at least one reply that day.
- Treatment date: **August 15, 2024**. The balanced panel includes zero-activity days.

Event study: activity increases after AI Overviews



- Effects emerge after the August 2024 rollout and persist through June 2025.
- July 2024 is the reference month. Standard errors cluster by subreddit.
- The design requires parallel counterfactual trends and no differential concurrent shocks.

Main DiD: AI Overviews increase engagement

$$Y_{it} = \beta(\mathrm{SFW}_i \times \mathrm{Post}_t) + \alpha_i + \gamma_t + \varepsilon_{it}.$$

	Comments/day	**Comment authors/day**
SFW × Post ($\widehat{\beta}$)	**+12.64**	**+6.61**
	(1.22)	(0.49)
Relative effect	**+12.0%**	**+12.3%**

- Both effects have $p < 0.001$.
- Similar growth in comments and authors points to broader participation.

Effect distribution: quantile treatment effects

Changes-in-changes estimates across the outcome distribution.

	Comments/day			Comment authors/day		
Quantile	Baseline	Effect	Relative	Baseline	Effect	Relative
Q20	0.06	0.006	10.9%	0.04	0.006	14.2%
Q30	0.44	0.172	39.4%	0.33	0.134	40.1%
Q40	1.39	0.881	**63.3%**	1.00	0.656	**65.6%**
Q50	3.27	2.009	61.4%	2.27	1.328	58.5%
Q60	7.46	3.572	47.9%	4.96	2.194	44.3%
Q70	17.64	5.625	31.9%	11.35	2.716	23.9%
Q80	48.39	9.294	19.2%	29.20	3.288	11.3%
Q90	171.73	36.169	21.1%	95.99	15.819	16.5%

- All displayed effects have $p < 0.001$.
- Proportional gains peak near the middle of the distribution.

Effect by community size: discovery of hidden gems

Re-estimate the DiD within December 2023 subscriber quartiles.

		Comments/day		**Authors/day**	
Quartile	Subscribers	Effect	Relative	Effect	Relative
Q1	0–238	0.66	44.2%	0.37	55.2%
Q2	239–2,080	5.87	**53.5%**	3.79	**74.5%**
Q3	2,081–12,385	10.75	34.7%	6.59	40.6%
Q4	$> 12,385$	35.84	10.4%	15.29	8.7%

- Absolute gains grow with size. Relative gains peak in Q2.
- Consistent with AI search helping users discover smaller communities.

Mechanism: factual versus experience-oriented content

Factual content

- Information is easier to summarize.
- AI answers can **substitute** for a visit.

Experience-oriented content

- Value comes from discussion and perspectives.
- AI may help users **discover** the community.

Mechanism: factual versus experience-oriented content

Factual content

- Information is easier to summarize.
- AI answers can **substitute** for a visit.

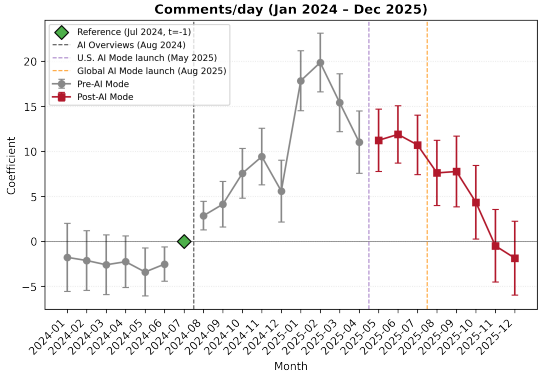
Experience-oriented content

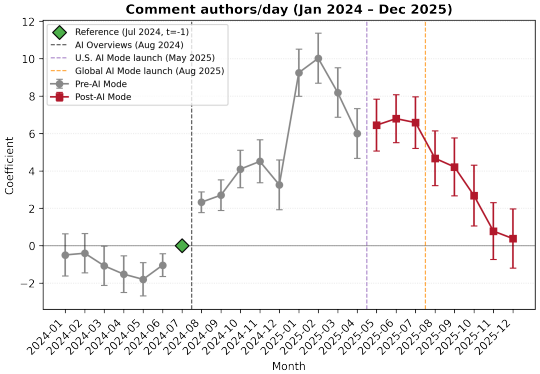
- Value comes from discussion and perspectives.
- AI may help users **discover** the community.

Triple difference	**Comments/day**	**Authors/day**
SFW × Post: factual baseline	7.67	3.51
SFW × Post × Experience	**9.95**	**6.34**
Experience / factual effect	**2.3**×	**2.8**×

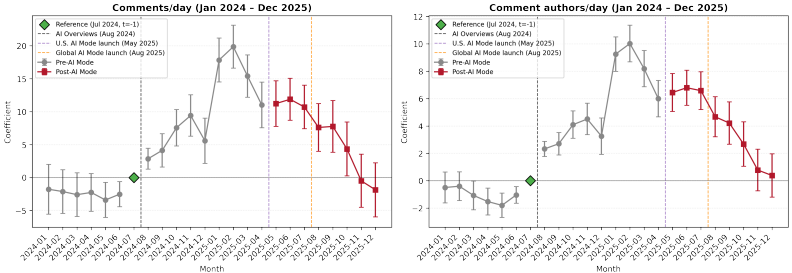
Both types gain, with larger effects for experience-oriented content.

AI Mode largely eliminates the experience premium





AI Mode largely eliminates the experience premium



Experience-content premium	**Comments/day**	**Authors/day**
Under AI Overviews	9.54	6.29
Change after AI Mode	-9.71	-3.72
Implied premium under AI Mode	-0.17	2.57
Attenuation	**102%**	**59%**

Search referrals reproduce the change in the premium

AI Overview period

$+54.4\%$

Estimated Google organic visits

+499.15 visits per subreddit-day
Standard error: 14.72

Experience-content premium

Under Overviews	465.95
Change after Mode	-661.83
Implied under Mode	-195.88

The effect also appears in traffic arriving from Google.

Google and Reddit: what we learn

1. AI Overviews increase engagement: comments **+12.0%**, authors **+12.3%**.
2. Experience-oriented content gains more: effects are **2.3–2.8 times** those for factual content.
3. Conversational AI Mode largely eliminates the experience-content premium.

Complementarity depends on both the content and the AI interface.

When Code Gets Cheap

Agentic AI, Product Entry, and Complement Scarcity

Bing Bai, Chenshan Hu, Yuval Salant, and Dennis J. Zhang



Bing Bai
McGill University



Chenshan Hu
University of Colorado Boulder



Yuval Salant
Northwestern University

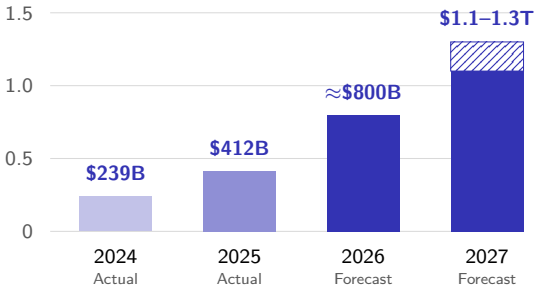
Singapore, November 2026

The demand side: an investment boom

Annual hyperscaler CapEx

Five hyperscalers, companywide spending

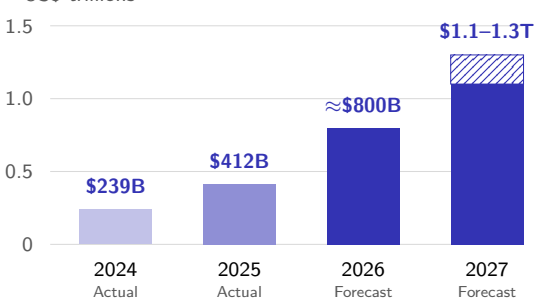
US$ trillions



The demand side: an investment boom

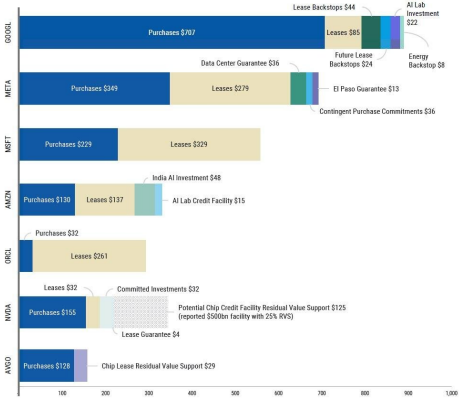
Annual hyperscaler CapEx

Five hyperscalers, companywide spending



>$3.1T in disclosed commitments

Purchases, leases, and guarantees (US$ billions)



The supply side: a task view of AI

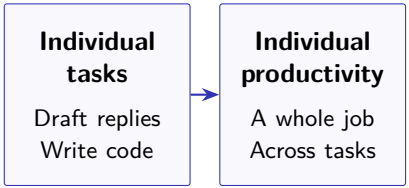
Individual tasks

Draft replies
Write code

Tasks are the starting point

Technology substitutes for some tasks, complements others, and creates new tasks.

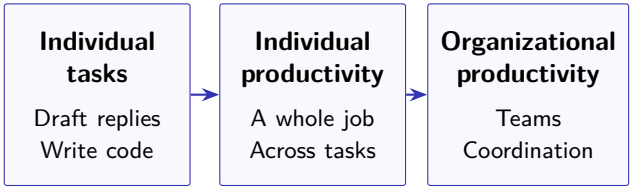
The supply side: a task view of AI



Tasks are the starting point

Technology substitutes for some tasks, complements others, and creates new tasks.

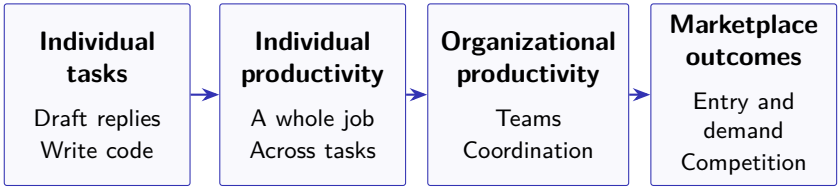
The supply side: a task view of AI



Tasks are the starting point

Technology substitutes for some tasks, complements others, and creates new tasks.

The supply side: a task view of AI



Tasks are the starting point

Technology substitutes for some tasks, complements others, and creates new tasks.

The supply side: a task view of AI



Tasks are the starting point

Technology substitutes for some tasks, complements others, and creates new tasks.

Each step requires evidence

Faster task completion can change production costs. Employment also depends on demand, entry, and adjustment.

AI coding and the 2026 SaaS selloff

AI is changing software production

"Anthropic's viral new work tool wrote itself"
Anthropic says Cowork was built with Claude Code.

"OpenAI launches Codex app to gain ground in AI coding race"
Developers can coordinate multiple coding agents.

AI coding and the 2026 SaaS selloff

AI is changing software production

"Anthropic's viral new work tool wrote itself"
Anthropic says Cowork was built with Claude Code.

"OpenAI launches Codex app to gain ground in AI coding race"
Developers can coordinate multiple coding agents.

Investors fear disruption to SaaS

"US software stocks slammed on mounting fears over AI disruption, lose $1 trillion in week"

$\sim$$1 trillion
S&P 500 software and services market value lost, Jan. 28–Feb. 5

ServiceNow -7.6% Salesforce -4.7%

The App Store: software production at global scale

One of the world's largest software distribution markets

850M+	$149B	>$1.4T
Average weekly users	Digital goods and services billings and sales	Total ecosystem billings and sales

Public launches make product entry observable

- Release dates and producer identities reveal **what enters and who enters**.
- App categories differ in the **non-code inputs** needed to create value.

Research questions

1. Do AI coding agents increase product and producer entry, and by how much?

Measure new app launches and first launches by producer accounts.

2. Which producers account for the change in app launches?

Decompose additional launches between entrant and incumbent producers.

3. How does the entry response differ with reliance on non-code complements?

Compare code-sufficient apps with products requiring content, networks, trust, or other complementary inputs.

The May 2025 coding-agent release window

Date	Release
May 16	OpenAI Codex cloud preview
May 19	GitHub Copilot coding agent
May 20	Google Jules beta
May 22	**Claude 4 and Claude Code general availability**
November 24	Claude Opus 4.5

Main cutoff: May 22, 2025

The design studies the broader release cluster. Later releases can reinforce the response.

Data: building the App Store panel

Charts and keyword searches can miss the little-noticed apps that lower entry costs may bring to market.

Collection step	What we observe
1. Find producer accounts	Enumerate Apple's identifiers and retain *Software Artists*: 1,875,406 producer accounts.
2. Retrieve app portfolios	Query 90 storefronts; 89 return records. Collect 164.7 million app–storefront listings.
3. Build global app records	Merge listings by app ID and keep the earliest release date. Link each app to its producer, genre, description, version history, and ratings.

A global census of App Store products and producers

Full census	
Unique apps	2.52 million
Producer accounts	1.08 million
Country storefronts	89

Release dates, genres, producer identities, version histories, and public ratings.

Estimation sample	
App launches	1,059,543
Producer accounts	564,008
Genres	26

Jan. 2025–Apr. 25, 2026 compared with Jan. 2023–Apr. 25, 2024.

Comparing the same seasonal break across two cycles

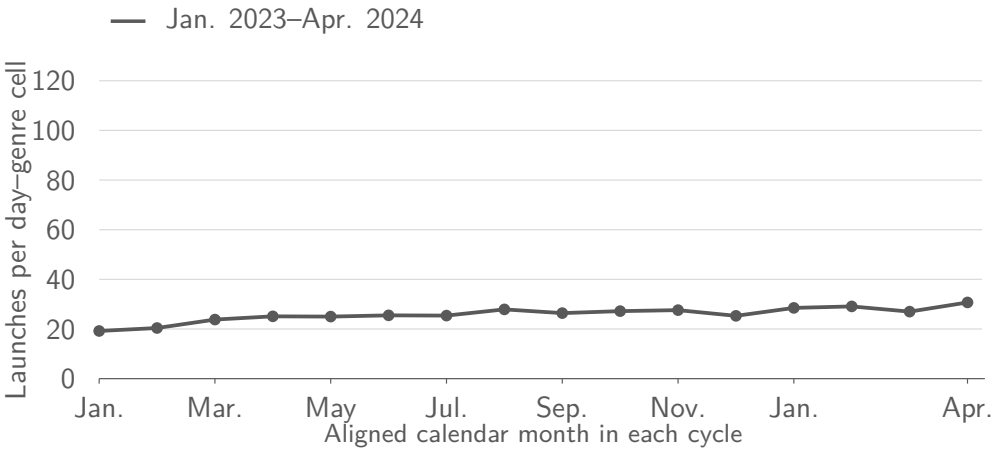
$$y_{dgk} = \beta(\text{Treated}_k \times \text{Post}_d) + \gamma_d + \delta_k + \theta_g + \varepsilon_{dgk}.$$

- d: aligned day within the cycle, g: genre, k: historical or recent cycle.
- γ_d absorbs common seasonality, δ_k cycle levels, θ_g genre differences.
- Standard errors cluster by actual calendar year–week.

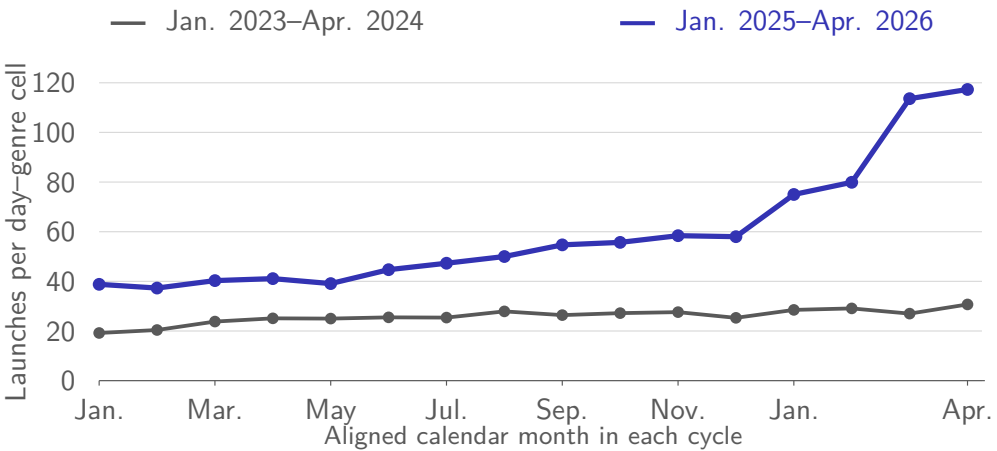
Identifying assumption

Absent the coding-agent releases, the recent cycle would have a comparable post-May seasonal change.

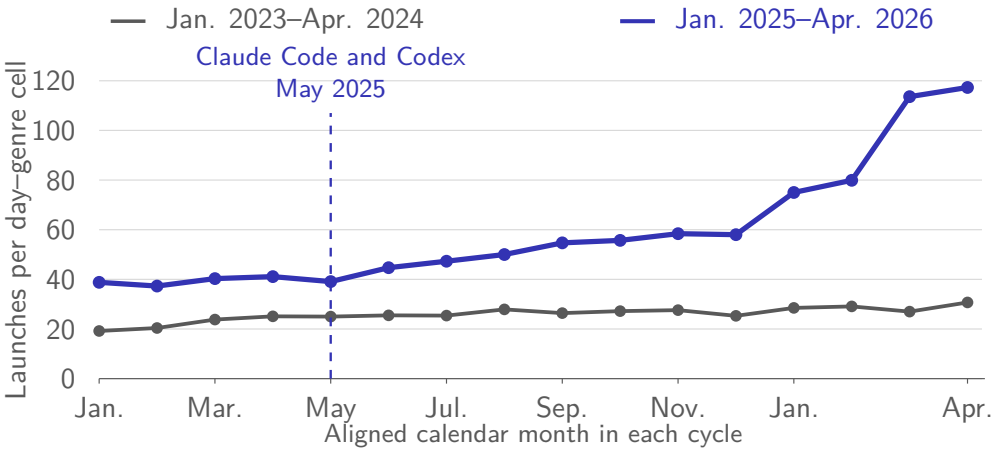
App entry: monthly launches



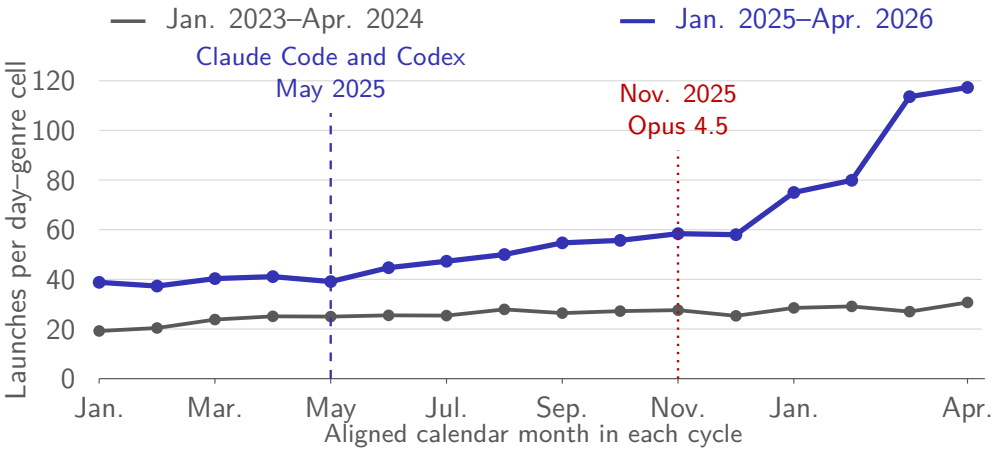
App entry: monthly launches



App entry: monthly launches



App entry: monthly launches



Product and producer entry: main results

Outcome	Estimate	(SE)	Effect (%)
A. Aggregate entry			
App launches	22.634^{***}	(2.553)	51.0
New producer accounts	13.532^{***}	(1.334)	80.6

$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

Product and producer entry: main results

Outcome	Estimate	(SE)	Effect (%)
A. Aggregate entry			
App launches	22.634^{***}	(2.553)	51.0
New producer accounts	13.532^{***}	(1.334)	80.6
B. App launches by producer status			
Entrant producers	21.895^{***}	(2.135)	101.6
Incumbent producers	0.739	(0.575)	3.2

$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

Product and producer entry: main results

Outcome	Estimate	(SE)	Effect (%)
A. Aggregate entry			
App launches	22.634^{***}	(2.553)	51.0
New producer accounts	13.532^{***}	(1.334)	80.6
B. App launches by producer status			
Entrant producers	21.895^{***}	(2.135)	101.6
Incumbent producers	0.739	(0.575)	3.2
C. New producer accounts by name category			
Individual-styled names	9.119^{***}	(0.890)	103.4
Company-styled names	4.414^{***}	(0.452)	55.3

$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

App launches rise by 51.0%, with entrant producer accounts driving 96.7% of the estimated increase.

Testing the task-based view

Prediction

A fall in coding costs should expand entry more where other required inputs are less binding.

Code-sufficient products

Working software delivers the core value.
Utilities, editors, planners, and simple games.

Testing the task-based view

Prediction

A fall in coding costs should expand entry more where other required inputs are less binding.

Code-sufficient products

Working software delivers the core value.
Utilities, editors, planners, and simple games.

Products needing other inputs

Cheaper coding leaves other requirements in place.
Content rights, network participation, trust, integration, local services, or hardware.

Testing the task-based view

Prediction

A fall in coding costs should expand entry more where other required inputs are less binding.

Code-sufficient products

Working software delivers the core value.
Utilities, editors, planners, and simple games.

Products needing other inputs

Cheaper coding leaves other requirements in place.
Content rights, network participation, trust, integration, local services, or hardware.

Test: compare entry across the two groups

Use an LLM to classify app names, genres, and descriptions, without launch timing, treatment status, or outcomes.

App launch responses by code sufficiency

App launches per day	Code-sufficient	Other products
Treated × Post	344.572^{***}	243.913^{***}
Clustered standard error	(48.101)	(48.206)
Effect vs. counterfactual	**126.9%**	**27.6%**
Classified apps	306,383	753,157
Share of apps	28.9%	71.1%
Share of counterfactual launches	23.5%	76.5%
Share of estimated increase	**58.6%**	**41.4%**
Excess share of increase	+29.6 pp	−29.6 pp

$^{***}p < 0.001$

- Entry expands most where working code delivers the product's core value.
- Code-sufficient apps account for **58.6%** of the estimated increase in launches, despite representing only **28.9%** of classified apps.

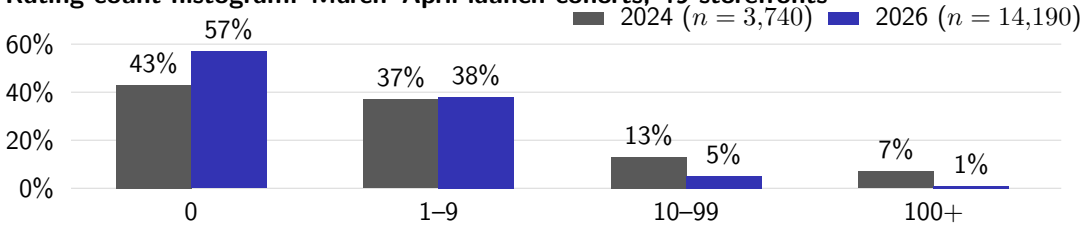
Do the newly produced apps find demand?

Public-reception regressions: full sample, 89 storefronts

	Zero ratings (pp)	≥ 10 ratings (pp)	Ratings/day (winsorized)	Score if rated (stars)
Treated × Post	-1.842^{***}	-1.454^{***}	-0.005	0.075^{***}
Clustered SE	(0.379)	(0.357)	(0.003)	(0.009)
Launch-day, cycle, genre FE	Yes	Yes	Yes	Yes
Observations	1,059,543	1,059,543	1,059,543	583,997

$^{***}p < 0.001$

Rating-count histogram: March–April launch cohorts, 49 storefronts



Newly launched apps receive more updates

Version updates within a fixed window after launch

Follow-up window	30 days	45 days	60 days	90 days
Additional updates	0.163	0.219	0.264	0.327
Clustered SE	(0.019)	(0.023)	(0.027)	(0.032)
Relative effect	14.5%	15.7%	16.2%	15.9%

GitHub provides independent evidence of upstream iOS development

A probability sample of public repository creation

Ten randomly selected five-minute windows per day, or 3.47% of each day, across both cycles.

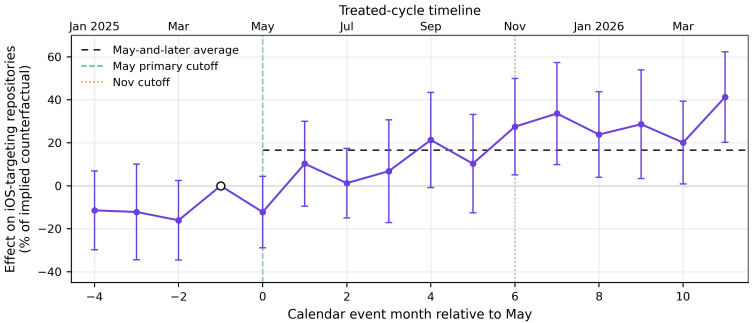
4.88m
Public repositories sampled

77,222
Candidates from a broad metadata screen

22,243
Verified iOS-targeting app repositories

Validation uses repository content, build configuration, and source files. It excludes libraries, tutorials, templates, and forks.

iOS-targeting repository creation also rises on GitHub



+35.2%

All-post effect on iOS-targeting repository creation

$\widehat{\beta} = 6.751$ per sampled day

Clustered SE: 0.945

Observable AI-assisted development also increases sharply

AI-assisted iOS-targeting repositories per sampled day

	Primary definition	**Strict definition**
	(1)	(2)
Treated × Post	2.243***	1.548***
Clustered SE	(0.294)	(0.240)
Treated-post mean	2.587	1.764
Effect vs. counterfactual	**652.3%**	**718.0%**
Day fixed effects	Yes	Yes
Cycle fixed effects	Yes	Yes
Observations	960	960

***$p < 0.001$

Robustness: trends and timing

Check	Robustness results
1. Pre-period trends	No significant pre-trends in app, producer, or GitHub entry.
2. Differential-trend sensitivity	Sensitivity intervals remain above zero under the stated bounds: through $M = 1.5$ for apps and $M = 2.0$ for producers.
3. Placebo cutoffs	None of 156 earlier weekly breaks matches the main effect.
4. Earlier seasonal cycles	Earlier May 22 breaks: -4.4% to 16.7%, compared with 51.0% in the main cycle.

Robustness: sample, timing, and measurement

Check	Robustness results
5. Exclude U.S. storefronts	App-entry effect: 51.3%, compared with the 51.0% baseline.
6. Alternative May cutoffs	App-entry effects range from 49.9% to 53.0%.
7. Classification error	Code-sufficient apps still contribute disproportionately after correcting label errors.

Two papers: information consumption and production

Study	Measured change	Governance question
Reddit	Participation and referrals differ across interfaces and communities	How does discovery support future information supply?
Coding	New producers drive entry, especially where code is sufficient	Which complements determine who benefits?

Thank you

Dennis Zhang

denniszhang@wustl.edu