Discussion of: AI as "Co-founder": GenAI for Entrepreneurship

By Junhui Jeff Cai, Xian Gu, Liugang Sheng, Mengjia Xia, Linda Zhao, and Wu Zhu

Yilin (David) Yang[1]

ABFER 13th ANNUAL CONFERENCE

May 20, 2026

[1] Carlson School of Management, UMN.

Motivation: a growth perspective (I)

- Classical production combines physical capital, labor, and managerial/specialist talent. These inputs are *complementary*—you cannot run a firm with ideas and labor alone.

- Implication for entry: starting a firm requires *assembling a bundle*—capital, specialist labor (coders, designers, lawyers, marketers), and managerial coordination. The minimum viable bundle defines a **fixed cost of entry**, and potentially good entrepreneurs cannot cross the threshold.
    - Even harder when complementary inputs are concentrated in incumbents.

Motivation: a growth perspective (II)

- A classical macro/growth theme: entry threshold $\Rightarrow$ misallocation $\Rightarrow$ aggregate TFP loss. Foundational evidence:
  - Evans and Jovanovic (1989): liquidity constraints bind on the entry decision. Wealth—not just talent—selects who becomes a founder.
  - Hsieh and Klenow (2009): misallocation of capital and labor across plants would raise manufacturing TFP by 30–50% in China and 40–60% in India.
  - Aghion and Howitt (1992): Schumpeterian creative destruction—growth requires successive waves of innovators displacing incumbents; the speed of displacement determines aggregate progress.
- Brave new world: GenAI—rapid, accessible, broad-task capable. Can GenAI democratize firm creation and promote growth?

Motivation: a growth perspective (III)

- A classical macro/growth theme: entry threshold $\Rightarrow$ misallocation $\Rightarrow$ aggregate TFP loss.
- Brave new world: GenAI—rapid, accessible, broad-task capable. Can GenAI democratize firm creation and promote growth?
- Acemoglu (2024), "The Simple Macroeconomics of AI": task-based model + Hulten's theorem $\Rightarrow$ upper bound of **0.71% TFP gain from AI over 10 years**.
    - But the framework treats AI as cost-saving *within existing firms*—a representative-firm task aggregation.

Motivation: a growth perspective (III)

- A classical macro/growth theme: entry threshold ⇒ misallocation ⇒ aggregate TFP loss.
- Brave new world: GenAI—rapid, accessible, broad-task capable. Can GenAI democratize firm creation and promote growth?
- Acemoglu (2024), "The Simple Macroeconomics of AI": task-based model + Hulten's theorem ⇒ upper bound of **0.71% TFP gain from AI over 10 years**.
  - But the framework treats AI as cost-saving *within existing firms*—a representative-firm task aggregation.
  - It does *not* model: new firm entry, misallocation reduction across heterogeneous firms, or selection of talent into entrepreneurship—channels that matter for *disruptive* growth.

If GenAI also reduces the entry threshold (the channel this paper studies), the misallocation/selection effect *compounds* on Acemoglu's within-firm bound. **The macro number could be substantially larger**—and no one has yet quantified this.

What this paper does

- The paper offers one of the first *large-scale*, *micro-geographic* tests, using:
  - Universal Chinese firm registration data: 12.8M new firms, 2021–2024.
  - 340K AI invention patents (2010–2019) classified via a transformer model.
  - H3 hexagonal grids at $\sim$5 km^2 resolution: 166K cell-grids nationally.
- Headline finding:

> After ChatGPT (Nov 2022), grids with at least one pre-2020 AI patent see ≈ 5 *extra* new firms per quarter $\Rightarrow \approx 6\%$ of national firm entry. Driven by small firms; large-firm entry *declines*.

- Authors' interpretation: GenAI as a *"digital cofounder"* that lowers fixed costs, substitutes for managerial labor, and democratizes founding.

Empirical design

- Within-city DiD comparing high- vs. low-AI grids before/after 2022Q4:

$$Y_{gt} = \beta \left(\text{Post}_t \times \text{HighAI}_g\right) + \mu_{g\times q(t)} + \lambda_{c(g)\times t} + \varepsilon_{gt}.$$

- Fixed effects:
  - $\mu_{g\times q(t)}$: grid $\times$ calendar-quarter (seasonality + grid invariants).
  - $\lambda_{c(g)\times t}$: city $\times$ quarter (all aggregate city-level shocks).

Empirical design

- Within-city DiD comparing high- vs. low-AI grids before/after 2022Q4:

$$Y_{gt} = \beta \left(\text{Post}_t \times \text{HighAI}_g\right) + \mu_{g\times q(t)} + \lambda_{c(g)\times t} + \varepsilon_{gt}.$$

- Fixed effects:
    - $\mu_{g\times q(t)}$: grid $\times$ calendar-quarter (seasonality $+$ grid invariants).
    - $\lambda_{c(g)\times t}$: city $\times$ quarter (all aggregate city-level shocks).
- **Treatment definition.** $\text{HighAI}_g = 1$ if grid g filed ≥ 1 AI patent in 2010–2019. About 6.1% of grids ($\sim 10{,}183$) treated.
- **Firm size.** Small if registered capital $<$ 1M RMB; large otherwise.
- Identification assumption: within a single city, in the same quarter, neighboring grids that differ only in their pre-2020 AI footprint are valid counterfactuals.

Empirical design

- Within-city DiD comparing high- vs. low-AI grids before/after 2022Q4:
$$Y_{gt} = \beta \left(\text{Post}_t \times \text{HighAI}_g\right) + \mu_{g\times q(t)} + \lambda_{c(g)\times t} + \varepsilon_{gt}.$$
- Fixed effects:
  - $\mu_{g\times q(t)}$: grid $\times$ calendar-quarter (seasonality $+$ grid invariants).
  - $\lambda_{c(g)\times t}$: city $\times$ quarter (all aggregate city-level shocks).
- **Treatment definition.** $\text{HighAI}_g = 1$ if grid g filed ≥ 1 AI patent in 2010–2019. About 6.1% of grids ($\sim 10{,}183$) treated.
- **Firm size.** Small if registered capital $<$ 1M RMB; large otherwise.
- Identification assumption: within a single city, in the same quarter, neighboring grids that differ only in their pre-2020 AI footprint are valid counterfactuals.

Main results at a glance

Outcome	Coef	s.e.
Total new firms (Table 2, col. 1)	$+5.04^{***}$	1.40
Small firms (col. 2)	$+7.70^{***}$	1.22
Large firms (col. 3)	-3.12^{***}	0.67
% serial entrepreneurs (Table 6)	-0.41^{**}	0.20
# shareholders (Table 7)	-0.021^{***}	0.007
# executive team members (Table 8)	-0.016^{***}	0.004
Placebo: Non-AI patents (Table 10)	$+0.16$	0.48
Placebo: Residualized AI (Table 11)	$+0.78^{***}$	0.17

- Three things to highlight:
  - (i) Asymmetric size response: small up, large down (not just *more* entry).
  - (ii) Margins of adjustment all move the same way: fewer experienced founders, fewer shareholders, smaller teams.
  - (iii) The non-AI patent placebo is reassuringly null. The residualization placebo loses 85% of the coefficient.

Comment 1: The mechanism could be strengthened

- Preferred story: GenAI as "digital cofounder"—substitutes for labor, lowers fixed costs, enables first-time founders. Evidence: fewer serial entrepreneurs, fewer shareholders, smaller teams in AI-downstream sectors.

- But the same evidence is consistent with several alternatives:

Alternative story	Predicts the same pattern?
(A) Hype-driven overconfidence. Inexperienced founders believe AI makes founding easier and register anyway.	Small, novice, few shareholders.
(B) Speculative shells. Thin vehicles registered to attract VC or AI-themed subsidies.	Small, novice, few shareholders.
(C) Tech-layoff spillovers. 2023 tech layoffs push engineers into solo ventures.	Small, first-time, AI-area.
(D) Local AI subsidies. Zhongguancun, Lingang target AI grids.	Entry concentrated in AI grids.
(E) Strategic registration. Tech-savvy founders optimize across size-contingent benefit cliffs.	Composition shifts to small.

- A direct test would need (i) a measure of *actual GenAI use* by new firms; (ii) heterogeneity that separates GenAI from generic AI-subsidy channels; (iii) survival outcomes 12–24 months out.

Comment 2: Registered capital is not firm size in China

- The headline asymmetry—small up, large down—hinges on the 1M RMB registered-capital cutoff. But registered capital in China is a *strategic choice*, not real scale.
- After the 2014 Company Law revision (the "subscribed capital system"):
  - Founders *subscribe* to capital with no paid-in requirement.
  - A single entrepreneur can register a 100M RMB company without injecting a single yuan.
  - Registered capital reflects licensing thresholds, tax optimization, signaling, shell-company purposes—*not* actual capital.
- Registered capital also interacts with multiple size-contingent benefit cliffs (small-scale VAT taxpayer status, SLPE income tiers, "Specialized SME" / Little Giant eligibility size caps). Tech-savvy founders optimize across all of these.

What looks like "GenAI shifts the optimal scale" may be "the tax code shifted, and AI grids responded faster."

- Table 9 (Section 5.3.1, serial entrepreneur downsizing) is the cleanest piece of evidence—but it still uses registered capital and thus inherits the same measurement problem.

Comment 3: Registered address $\neq$ operating address

- The paper geolocates firms by *registered address* (SAIC data). In China, registered address and operating address are legally and practically distinct.

- Direct policy evidence:
  - "Cluster registration" (集群注册) is a Chinese government policy allowing one address to host hundreds of firms not physically present.
  - Shanghai centralized-registration zones (Chongming, Baoshan, Lin-gang): ~¥2,000/yr per firm for a registered address.
  - Industrial parks, incubators, FTZ-hosted addresses advertise registration without operations: Lin-gang, Qianhai, Hainan, Zhongguancun Sci-Tech Park.

- Direct academic evidence: Hu, Liang, Lu and Chen (2021) (*Growth and Change*, Wiley) document for A-share listed firms:

  > *"Companies with separate headquarters and registered addresses account for **approximately one-third** of all listed companies... usually headquartered in economically developed cities while their registered addresses are mainly located in general cities... leads to **deviations from reality** in the conclusions of many studies that use the registered address as the location of corporate headquarters."*

  Listed firms (highest compliance) already show 1/3 separation. Unlisted firms (this paper's sample): likely more.

The treatment unit (which H3 grid the firm is "in") is partly endogenous to where founders *choose* to register, not where firms operate.

- Suggestion: match SAIC to social-insurance contributions by district, and drop firms registered at known cluster-registration addresses.

Comment 4: What does HighAI capture?

- The conceptual claim: pre-2020 AI patents proxy for "AI-specific human capital that can adopt GenAI tools."
- But the GenAI skill set is largely *disjoint* from the 2010–2019 AI skill set:

2010–2019 AI patents (mostly)	GenAI deployment (mostly)
CNN image classification	Prompt engineering
LSTM / RNN speech	RAG architecture
Recommender systems	Fine-tuning workflows
Classical ML / SVMs	Application-layer integration

- A small-business owner using ChatGPT for marketing copy does not need the human capital that filed a 2015 CNN-based patent.
  - Not necessary to know why, sufficient to know how.

What is the high-AI patent grid *actually* capturing? Tech-firm density, AI-awareness, VPN usage, infrastructure, VC proximity, **industrial-park subsidy**—all plausible mediators, but none specifically about *GenAI human capital*.

Comment 4.1: Who responds is not who has AI human capital

- If pre-2020 AI patents capture "AI-specific human capital," the entry surge should concentrate in technically AI-adjacent sectors. The paper's Table 4 (Panel A, top 15) shows the opposite.

Industry	**Coef.**	**Sig.**
Retail Industry	**1.626**	***
Business Services	0.977	***
Technology Promotion & Application Services	0.871	***
Wholesale Industry	0.455	*
Software and IT Services	**0.056**	**n.s.** $(p = 0.60)$
Broadcasting, Television, Film, and Audio Production	0.044	***

- Two striking facts:
  - **Retail leads at 1.626**—*29× larger* than Software & IT Services.
  - **Software & IT Services** is the *only sector in the top 15 that is statistically insignificant* $(p = 0.60)$. If pre-2020 AI patents capture AI human capital, software firms should respond *most*. They don't.

Comment 4.1: Who responds is not who has AI human capital

- If pre-2020 AI patents capture "AI-specific human capital," the entry surge should concentrate in technically AI-adjacent sectors. The paper's Table 4 (Panel A, top 15) shows the opposite.

Industry	**Coef.**	**Sig.**
Retail Industry	**1.626**	***
Business Services	0.977	***
Technology Promotion & Application Services	0.871	***
Wholesale Industry	0.455	*
Software and IT Services	**0.056**	**n.s. ($p = 0.60$)**
Broadcasting, Television, Film, and Audio Production	0.044	***

- Two striking facts:
  - **Retail leads at 1.626**—*29× larger* than Software & IT Services.
  - **Software & IT Services** is the *only sector in the top 15 that is statistically insignificant* ($p = 0.60$). If pre-2020 AI patents capture AI human capital, software firms should respond *most*. They don't.
- This sets up the next slide: the paper's own AI-upstream/downstream classification (Section 5.2, Table 5) tells the same story.

Comment 4.2: AI upstream vs. downstream industries

Bucket	What firms do	Examples
AI-upstream	*Produce* AI inputs: chip design, foundation models, ML research, training infrastructure	DeepSeek-type firms
AI-downstream	*Use* GenAI as office-productivity tool	Marketing, content, legal drafting, e-commerce, consulting, education

- Paper's Table 5: high-helpfulness coefficient $+5.892$ (s.e. 1.302, $p < 0.01$) vs. low-helpfulness -0.848 (s.e. 0.300). High-upstream is only $+0.846$; low-upstream is $+4.199$.
- AI-upstream entry would require *AI-specific human capital*—consistent with the patent proxy.
- AI-downstream entry requires *access to ChatGPT plus office labor*—available anywhere, not specifically concentrated near AI patent grids.
- The paper's own evidence: GenAI helps *office-labor-intensive, low-capital downstream sectors*. None of those sectors require local AI human capital. So why are AI-patent grids the ones responding? Likely because they coincide with general tech-firm density, VC, infrastructure—*or with subsidy zones*.
- Sharper tests: (i) within HighAI grids, is the downstream effect larger where pre-2020 AI patents were in NLP / language models? (ii) does AI-downstream entry appear in LowAI grids once you control for general tech-firm density?

Comment 5: 2022Q4 is a treacherous event date for China

- The post-ChatGPT window coincides with several large macro shocks that hit *tech-dense urban grids differentially*—exactly the within-city variation the design relies on.

Shock	Timing	Differential effect on AI grids
Zero-COVID exit	Dec 2022	Services / consumer-facing rebound.
Tech crackdown thaw	Early 2023	DiDi relisted; gaming licenses resume.
Property contraction	2023–24	Capital reallocates out of real estate.
Youth unemployment	Jun 2023 ($\rightarrow$ Aug suspension)	21.3% recorded; tech-worker spillovers into self-employment.
SME tax extension	Mar 2023	5% effective CIT for sub-3M RMB income extended to 2027.

Comment 5: 2022Q4 is a treacherous event date for China

- The post-ChatGPT window coincides with several large macro shocks that hit *tech-dense urban grids differentially*—exactly the within-city variation the design relies on.

Shock	Timing	Differential effect on AI grids
Zero-COVID exit	Dec 2022	Services / consumer-facing rebound.
Tech crackdown thaw	Early 2023	DiDi relisted; gaming licenses resume.
Property contraction	2023–24	Capital reallocates out of real estate.
Youth unemployment	Jun 2023 ($\rightarrow$ Aug suspension)	21.3% recorded; tech-worker spillovers into self-employment.
SME tax extension	Mar 2023	5% effective CIT for sub-3M RMB income extended to 2027.

- City $\times$ quarter FE absorbs *aggregate* city shocks. It does *not* absorb within-city heterogeneous responses—and these shocks are inherently heterogeneous across grids (tech-thaw benefits Zhongguancun more than Shijingshan within Beijing).

A clean falsification: use a non-AI "tech hub" indicator (pre-2020 IT employment, software firm density) and show the effect is *specific* to AI patents, not to tech-dense neighborhoods more broadly.

Minor comments

- **Binary HighAI is wasteful.** 6.1% treated. Table B.1 (75th-pctile) finds *larger* effects (11.62 vs. 7.70)—a dose-response. Continuous-treatment should be the main result.

- **Multiple testing in Table 4.** 96 industries, no FDR / Bonferroni correction.

- **Figure 1 is part of the treatment, not evidence for it.** China's AI model releases surged *after* ChatGPT *because of* ChatGPT.

- **Endogenous Post timing.** Response requires awareness, access to Chinese substitute, and time to register. Effective treatment likely begins mid-to-late 2023.

- **Typos.** "mange of these tasks" (p. 21) $\rightarrow$ "many"; "firm models" (p. 6) $\rightarrow$ "business models."

Concluding remarks

- Important, timely question. The data construction—H3 grids + universal firm registration + AI-classified patents—is a contribution in its own right.
- What I like most:
  - Within-city, within-quarter identification at $\sim$5km^2 resolution.
  - Serial-entrepreneur downsizing test (Section 5.3.1)—the cleanest evidence.
  - Non-AI patent placebo (Table 10) is well-designed.
- Where I'd push:
  - **Mechanism.** Several stories fit. Need a direct measure of GenAI use, not just patent proximity.
  - **Spatial unit.** Registered address $\neq$ operating address. Validate with employment/social-insurance data.
  - **HighAI proxy.** Pre-2020 AI patents likely capture tech-district characteristics, not GenAI human capital. Table 4 & 5 confirm.
  - **Confounders.** 2022Q4 has too many co-moving shocks for city-quarter FE alone.

Strong potential. Looking forward to the next draft.

Thank you!

References I

Acemoglu, Daron, "The Simple Macroeconomics of AI," *Economic Policy*, 2024. Forthcoming. April 5, 2024 version.

Aghion, Philippe and Peter Howitt, "A Model of Growth through Creative Destruction," *Econometrica*, 1992, *60* (2), 323–351.

Evans, David S. and Boyan Jovanovic, "An Estimated Model of Entrepreneurial Choice under Liquidity Constraints," *Journal of Political Economy*, 1989, *97* (4), 808–827.

Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 2009, *124* (4), 1403–1448.

Hu, Guojian, Xiaomin Liang, Yuqi Lu, and Yu Chen, "Interprovincial Separation between Headquarters and Registered Addresses: The Case of Listed Companies in China," *Growth and Change*, 2021, *52* (2), 1080–1098.