

AI as “Co-founder”: GenAI for Entrepreneurship *

Junhui Jeff Cai [†] Xian Gu [‡] Liugang Sheng [§]
Mengjia Xia [¶] Linda Zhao ^{||} Wu Zhu ^{**}

December 6, 2025

Abstract

This paper studies whether, how, and for whom generative artificial intelligence (GenAI) facilitates firm creation. Our identification strategy exploits the November 2022 release of ChatGPT as a global shock that lowered start-up costs and leverages variations across geo-coded grids with differential pre-existing AI-specific human capital. Using high-resolution and universal data on Chinese firm registrations by the end of 2024, we find that grids with stronger AI-specific human capital experienced a sharp surge in new firm formation—driven entirely by small firms, contributing to 6.0% of overall national firm entry. Large-firm entry declines, consistent with a shift toward leaner ventures. New firms are smaller in capital, shareholder number, and founding team size, especially among small firms. The effects are strongest among firms with potential AI applications, weaker financing needs, and among first-time entrepreneurs. Overall, our results highlight that GenAI serves as a pro-competitive force by disproportionately boosting small-firm entry.

JEL Classification: L26, O33, L53, G30

Keywords: Generative AI; Entrepreneurship; Firm Entry; Human Capital; Technological Change; Spatial Economics

*We are grateful to Hui Chen, Petra Todd, Bernard Yeung, and Bohui Zhang for their insightful comments and suggestions. We thank Yao Li, Xingyu Lu, Fangzhuo Yang, Dibo Peng, and Zhihan Zhou for outstanding research assistance. Zhu acknowledges financial support from the Tsinghua University Initiative Scientific Research Program (No. 2022Z04W02016) and the Tsinghua University School of Economics and Management Research Grant (No. 2022051002). All remaining errors are our own.

[†]Mendoza College of Business, University of Notre Dame, jcai2@nd.edu

[‡]Department of Finance, Durham University Business School, xian.gu@durham.ac.uk

[§]Department of Economics, The Chinese University of Hong Kong, Hong Kong SAR, China, lsheng@cuhk.edu.hk

[¶]Department of Economics, The University of Pennsylvania, xiax@sas.upenn.edu

^{||}Wharton School, The University of Pennsylvania, lzhaio@wharton.upenn.edu

^{**}Department of Finance, SEM, Tsinghua University, zhuwu@sem.tsinghua.edu.cn

1 Introduction

Technological evolution underpins long-run growth, driving entrepreneurship as “creative destruction” (Schumpeter, 1943) that turns new technologies into new firms and markets (Aghion et al., 2018). The recent meteoric rise of *generative artificial intelligence* (GenAI) was unprecedented: OpenAI’s ChatGPT (GPT-3.5), released in November 2022, reached roughly 1 million users in five days and about 100 million monthly active users within two months, outpacing any prior platforms.¹ Crucially for entrepreneurship, GenAI differs from earlier waves because it combines broad task generality with easy accessibility via a universal natural-language interface (Eloundou et al., 2024). In practice, it automates language-based information-processing (reasoning, synthesis, coding, and generative), alongside creative and managerial tasks, thereby lowering the barriers to capabilities that were once capital-intensive or concentrated among specialized labor and well-resourced organizations. GenAI thus can serve as a co-founder and enables entrepreneur to function as a small, multi-skilled team. As a result, teams shrink, fixed costs fall, and entry accelerates. “AI-native” ventures such as Midjourney², Cursor³, and Perplexity⁴ exemplify this shift. Against this backdrop, we ask: *does the diffusion of GenAI increase firm creation, through which mechanisms, and who benefits most?*

In this paper, we study these questions by exploiting the November 2022 release of ChatGPT as a sharp shock to the salience and availability of GenAI. Motivated by recent evidence that human capital plays a central role in the diffusion and adoption of AI technologies (Babina et al., 2023; Gofman and Jin, 2024; Andreadis et al., 2025; Cong and Zhu, 2024), our identification strategy leverages *within-city* variation across geocoded grids with differential pre-existing AI-specific human capital, proxied by the number of

¹TikTok reached 100 million users in around 9 months and Instagram in around 2.5 years. [Yahoo Finance](#).

²<https://www.midjourney.com/home>

³<https://cursor.com>

⁴<https://www.perplexity.ai>

AI invention patents filed between 2010 and 2019. To implement this design, we combine two comprehensive datasets: (i) universal firm registration records in China covering more than 12 million newly established firms from 2021–2024, and (ii) the complete set of AI invention patents filed between 2010–2019. We geolocate each new firm to a hexagonal H3 grid cell—at roughly 5 km^2 spatial resolution—and aggregate pre-2020 AI patents to the same cells using assignee addresses. This procedure yields a novel, nationwide H3 grid-by-quarter panel comprising 166,156 spatial cells, enabling a highly granular comparison of how GenAI reshaped firm creation across grids with varying stocks of AI-relevant human capital.

We focus on China for both scale and relevance: it is a major AI hotbed for entrepreneurship and research, and a major economy with detailed nationwide geo-coded firm registration microdata. We identify over 12 million new firm registrations from 2021–2024, including over 5.5 million *small* firms (registered capital < 1 million RMB) and over 6.5 million *large* firms, and classify 340,771 AI-related patents from 2010 to 2019 following the methodology of [Fang et al. \(2025a\)](#). Our key finding is a sizable surge in new firm formation following the release of ChatGPT in 2022, concentrated in areas with more AI-specific human capital, and driven primarily by small, resource-constrained, AI downstream firms. Our results suggest that GenAI can function as a “co-founder,” providing domain knowledge as well as managerial and operational know-how, which reduces entry barriers and favors smaller, leaner organizational forms.

Our baseline empirical identification relies on a novel *difference-in-differences* (DiD) framework that exploits fine-grained spatial heterogeneity in AI-specific human capital within cities before 2020. Specifically, we define a grid with high AI-specific human capital (high AI-exposure) if this grid has at least one AI-related invention patent filed between 2010 and 2019. Next, we compare changes in firm entry before and after the release of ChatGPT between grid cells with high- and low-AI human capital within the same city. The specification includes both *grid-by-calendar-quarter* and *city-by-quarter*

fixed effects. The former absorbs seasonal patterns in firm registrations as well as any time-invariant unobservables specific to each grid, while the latter controls for all *time-varying city-level factors*, including local policy initiatives, business environment shifts, industrial restructuring, and macroeconomic fluctuations. This fixed-effect structure ensures that identification arises solely from *within-city, within-quarter* variation across neighboring grids that differ in pre-existing exposure to AI.

This design provides several advantages for causal inference. First, by exploiting micro-geographic variation, it mitigates concerns that our estimates are driven by cross-city differences in economic development, infrastructure, or administrative quality. Second, because the diffusion of ChatGPT represents a global, near-synchronous technology and information shock, the temporal variation in exposure is plausibly exogenous to local entrepreneurial environment and activities. Third, our approach isolates the role of *AI-specific human capital*—rather than general innovative capacity—by comparing otherwise similar grids within the same urban environment. Remaining identification concerns, such as local spillovers, differential pretrend, confoundedness, and robustness are addressed through a series of placebo and robustness tests discussed later in the paper.

We find that the diffusion of GenAI triggered a pronounced surge in new firm creation across China. Grids with stronger pre-existing AI-specific human capital experienced a sharp and persistent rise in firm entry following the release of ChatGPT. The magnitude of this effect is economically meaningful: high-AI grids recorded roughly five additional new firms per grid-quarter relative to low-AI grids. Aggregated across all high-AI grids, this translates into approximately 51,000 additional new firms per quarter, accounting for about 6.0% of total nationwide firm entry after ChatGPT’s release. Event-study estimates show no evidence of differential pre-trends between high- and low-AI grids prior to the shock, supporting a causal interpretation of the results.

A defining feature of this entrepreneurial boom is its asymmetry by firm size. The increase in entry is driven entirely by small, resource-constrained firms, while entry by

larger firms declines. This pattern is consistent with evidence that GenAI substantially lowers fixed organizational and managerial costs and increases efficiency, thereby effectively lowering the minimum viable scale of new ventures and allowing them to operate with less labor and external finance.

Further evidence shows that the largest increases in firm entry arise in AI-downstream sectors that apply AI in core production or service—such as retail, business services, online platforms and other digital-service industries—where GenAI tools can be readily applied to a broad set of product-developing, information-processing, customer-facing, and back-office activities. By contrast, the ChatGPT shock slightly reduces firm entry in traditional and capital-intensive industries including construction and manufacturing, indicating that the sectors whose core of business activities rely on knowledge work, creativity, marketing, and digital service are most responsive to GenAI.

We further investigate the underlying mechanisms. First, the share of serial entrepreneurs who have started at least one firm in the previous three years declines in AI-intensive areas after the release of ChatGPT, indicating that GenAI facilitates entry by first-time founders who traditionally lack experience and managerial or operational know-how. Second, new firms are created with fewer shareholders, reflecting reduced financial constraints at entry as founders rely less on external capital or co-investors. Third, founding teams become smaller, consistent with AI substituting for early-stage managerial and specialized labor and enabling much faster prototyping, thereby allowing leaner organizational forms. Evidence from repeat founders further corroborates this substitution effect in firm scale: serial entrepreneurs deliberately launch smaller ventures in the post-ChatGPT period, suggesting that GenAI not only expands the extensive margin of entry but also shifts the optimal scale of new firm creation towards leaner organizations.

We conduct a series of placebo exercises to verify that the entrepreneurial response to GenAI is driven by *AI-specific human capital* rather than existing local innovation capability or entrepreneurial trends. To separate AI-specific from general innovative capabilities,

we replace AI-related patenting with non-AI patenting, a proxy for general innovation capabilities. The estimated treatment effects largely vanish: grids with strong non-AI innovation capacity do not experience comparable post-ChatGPT surges in new firm formation.

To address the concern that our estimates might simply reflect pre-existing differences in entrepreneurial activity—rather than AI-specific human capital—we orthogonalize pre-2019 firm entry with respect to AI patenting at the grid level. Specifically, we regress pre-2019 entrepreneurial activity on historical AI patents and construct residuals that purge the variation associated with AI-specific human capital. We then re-estimate our design using these residuals as the placebo “treatment.” The resulting effects shrink by roughly 80–85%, indicating that our baseline findings are not mechanically driven by pre-existing entrepreneurial vibrancy, but rather by the component of AI exposure that reflects AI-specific human capital.

As an additional falsification check, we randomly reassign the “high-AI” label across spatial grids and re-estimate our baseline specification for 100 times. The distribution of the estimated coefficients centers tightly around zero, indicating that our baseline results are due to chance spatial correlations or model mis-specification.

Finally, robustness checks confirm that our results are not driven by regional composition, sample design, or measurement choices. Excluding first-tier provinces such as Beijing, Shanghai, and Guangdong yields similar results, indicating that the effects extend well beyond China’s major innovation hubs. Restricting the sample to AI-active grids, matching AI-intensive locations with nearby non-AI counterparts within the same city, and varying the threshold for small firm definition all produce consistent results: grids with stronger pre-existing AI capabilities continue to experience large and statistically significant increases in small-firm entry, alongside corresponding declines in large-firm entry, after the release of ChatGPT. We further use a matching estimator that matches AI-grids with non-AI grids within the same city and obtain similar results, suggesting that the

effect is not driven by the national or local government industrial policies that support AI as a strategic industry.

Taken together, our results provide consistent evidence that GenAI relaxes key frictions that traditionally constrain new business formation, including prior entrepreneurial experience, financing, and managerial labor. These effects are concentrated in regions with stronger AI-specific human capital, where potential founders are in a better position to recognize, adopt, and deploy GenAI, highlighting the importance of human capital and technology-skill complementarities in shaping the real economic implications of general-purpose technologies. With GenAI serving as a “co-founder” that allows small teams to replicate capabilities once requiring larger organizations, our findings suggest that GenAI has already begun to reshape entrepreneurial landscape: by democratizing who can start a firm, how firms are organized, and the scale of resources they require to operate.

Contribution to the literature. Our study contributes to several strands of literature at the intersection of AI, firm dynamics, and labor markets. The first is on firm entry. Firm entry is a vital engine for economic growth, fueling innovation, competition, job creation, and the efficient reallocation of resources. The theoretical importance of this process is rooted in the Schumpeterian concept of “creative destruction,” which posits that long-run economic progress stems not from gradual optimization but from a disruptive process where new firms—armed with innovative products, technologies, or firm models—displace less efficient incumbents ([Schumpeter, 1943](#)).⁵ A vast body of empirical work, leveraging large-scale firm- and plant-level data, has confirmed that the process of entry and exit is a primary driver of aggregate productivity gains, as resources are reallocated from exiting firms to more productive entrants ([Foster et al., 2001](#); [Brandt et al., 2012](#)). Furthermore, it is young firms, not simply small ones, that account for a dispropor-

⁵Modern economic theory has formalized this dynamic, with foundational models demonstrating how entrants act as the primary agents of this disruption, driving a selection process that filters out unproductive firms and sustains aggregate growth ([Jovanovic, 1982](#); [Hopenhayn, 1992](#); [Aghion and Howitt, 1992](#)).

tionate share of net job creation ([Haltiwanger et al., 2013](#)).⁶

A key question in entrepreneurship research is the extent to which technological breakthroughs contribute to firm creation. While they are widely recognized as powerful catalysts, their precise quantitative impact remains difficult to measure. For example, the diffusion of the personal computer fostered the rise of pioneers like Microsoft and Apple, and the internet’s commercialization in the mid-1990s triggered a wave of entrepreneurship that produced today’s tech giants, including Google, Amazon, and later, Facebook. The primary obstacle to measurement has been methodological: the slow and endogenous diffusion of these past technologies presents significant identification challenges for isolating their causal impact on startup activity.

The launch of ChatGPT, however, offers a stark contrast and a unique opportunity. Its arrival was largely unexpected, and its adoption has been unprecedentedly rapid, creating a natural quasi-experimental setting. Moreover, since the technology originated in the US, its arrival in China—and the subsequent boom in domestic large language model (LLM) development—represents a plausibly exogenous shock for local entrepreneurs. This context allows researchers to more cleanly identify the causal effects of a major technological breakthrough on firm creation. Our paper addresses this critical gap by providing one of the first pieces of systematic evidence on how technological breakthroughs, such as GenAI, facilitate firm entry, using the universe of newly registered firms in China between 2021 and 2024.

The second body of literature explores the impact of AI on firm dynamics and identifies distinct channels through which it enhances growth and productivity. For instance, [Babina et al. \(2024\)](#) find that AI-investing firms in the US experience higher growth in sales, employment, and market valuations, driven primarily by increased product innovation. Complementary evidence from European firm-level data shows that AI innovation

⁶Conversely, a decline in entrepreneurship can signal economic stagnation. For instance, [Decker et al. \(2014\)](#) document a decline in firm dynamism in the US and suggest this trend—characterized by lower startup rates and weaker growth among young firms—is a likely contributor to the nation’s recent sluggish productivity growth.

and adoption also directly promote firm output and productivity ([Czarnitzki et al., 2023](#); [da Silva Marioni et al., 2024](#)). More recently, the boom in GenAI has spurred a new wave of research, with several studies demonstrating that it increases worker productivity across various occupations and promotes sales and productivity of online retail ([Fedyk et al., 2022](#); [Brynjolfsson et al., 2025](#); [Kanazawa et al., 2025](#); [Fang et al., 2025b](#)), while others find that the release of GPT boosted the market valuation of firms with high exposure to the technology ([Eisfeldt et al., 2025](#); [Bertomeu et al., 2025](#)).

While existing research focuses on incumbent firms and worker-level effects, the impact of AI on market entry remains unexplored. Moreover, the net effect of GenAI on firm creation is theoretically ambiguous. On one hand, GenAI could strengthen market concentration by favoring large, resource-rich incumbents, thereby stifling new entry. On the other hand, it could lower entry barriers by democratizing access to sophisticated tools. Our study addresses this critical question by providing the first large-scale, systematic evidence on how GenAI affects firm entry, drawing on the universe of newly registered firms in China between 2021 and 2024. Contrary to concerns about the increasing market power of large incumbents ([Babina et al., 2024](#)), we find that GenAI predominantly encourages the entry of small firms. This finding has important implications for market dynamics, suggesting that GenAI may serve as a pro-competitive force that reduces concentration rather than reinforcing the dominance of incumbents.

Third, we also engage with the long-standing debate on the “race” between machine and human, and in particular the tension between technological progress enhancing labor productivity and displacing human labor. While previous waves of automation, such as industrial robotics, primarily impacted low- and middle-skilled blue-collar workers in routine manual tasks ([Acemoglu and Restrepo, 2020](#); [Acemoglu et al., 2022](#)), the advent of GenAI has extended this fear to high-skilled and white-collar professions, though the impact may take time to materialize. By performing sophisticated cognitive and knowledge work, GenAI’s impact is likely to be highly concentrated in skilled occupations and

high-pay jobs (Eloundou et al., 2024; Hampole et al., 2025). Indeed, emerging evidence indicates that this substitution effect often materializes not as layoffs, but as reduced hiring for entry-level skilled roles, creating significant barriers for new labor-market entrants (Lichtinger and Hosseini Maasoum, 2025). In contrast to this focus on the substitution effect, our study investigates the role of GenAI in facilitating entrepreneurship in boosting startup and firm creation. We argue that GenAI can reduce the costs associated with starting a firm by automating tasks like coding, marketing content creation, legal document drafting, and firm plan development. This may reduce the initial need for large teams and managerial capital. Thus, GenAI can serve as a powerful engine for entrepreneurship, particularly for inexperienced founders and firms in AI-downstream industries. This entrepreneurial channel is crucial, as new firm creation is a primary engine for net job growth and innovation, offering a vital counterweight to the potential for AI-driven job displacement (Haltiwanger et al., 2013). Moreover, we also find that geo-coded areas with high AI-specific human capital, not the general capacity, benefit the most from the positive impact of GenAI in boosting firm entry.

Finally, we also contribute to the booming literature exploring the broad implications of AI on academic research (Korinek, 2023; Cong and Zhu, 2024), financial market (Cong et al., 2025; Croom, 2025; Ashraf, 2025; Xue et al., 2025), innovation (Wu et al., 2025; Wang and Wu, 2025), and international trade (Antoniades et al., 2025). Our study focuses on the impact of GenAI on entrepreneurship and startup, the engine of economic growth.

The remainder of this paper proceeds as follows. Section 2 provides background on ChatGPT and the development of GenAI in China. Section 3 describes the firm registration and patent data, the construction of our fine-grained geo-coded panel as well as the summary statistics. Section 4 outlines our empirical strategies, and Section 5 presents the main results and explores mechanisms. Section 6 discusses the placebo and robustness tests. Section 7 concludes.

2 Release of ChatGPT and the Rise of AI Activity in China

The public release of ChatGPT by OpenAI in November 2022 represented a global turning point in the diffusion of GenAI. Unlike earlier waves of automation confined to coding or data analytics, ChatGPT demonstrated that large language models (LLMs) could perform a broad spectrum of cognitive and creative tasks—including text generation, document drafting, coding assistance, marketing content, and basic customer interaction—at negligible marginal cost. Its intuitive interface via natural languages and rapid performance improvements signaled a step-change in the accessibility of advanced AI. Within months, ChatGPT surpassed 100 million users worldwide and dominated global media attention, generating an unprecedented awareness shock that transformed how entrepreneurs and small firms perceived the cost and feasibility of launching new ventures. A growing cohort of “AI-native” start-ups built and run by “tiny teams”—often just a handful of engineers using tools like ChatGPT—are already scaling to millions of users and substantial revenue with far less labor and capital than previous generations, illustrating how GenAI can radically change the cost and organizational structure of building new firms.⁷

The diffusion effects were particularly visible in China, major technology firms rapidly accelerated domestic LLM development.⁸ Between early 2023 and mid-2024, Baidu launched ERNIE Bot⁹, Alibaba introduced Tongyi Qianwen (Qwen)¹⁰, iFlytek released

⁷<https://www.nytimes.com/2025/02/20/technology/ai-silicon-valley-start-ups.html>

⁸Although ChatGPT itself was officially inaccessible in mainland China, the technological and informational spillovers were immediate, disseminated through media coverage and professional networks. Within weeks of its release, Chinese technology firms, developers, and professional users accessed the model through API integrations, enterprise accounts registered abroad, and proxy connections. Moreover, domestic platforms rapidly introduced close substitutes such as Baidu’s ERNIE Bot and Alibaba’s Tongyi Qianwen (Qwen), which replicated many core functionalities of ChatGPT. As a result, both the awareness and practical use of GenAI tools spread widely across China’s entrepreneurial ecosystem.

⁹<https://ernie.baidu.com>

¹⁰<https://qwen.ai/home>

SparkDesk¹¹, Tencent rolled out Hunyuan¹², and ZhiPu AI debuted ChatGLM¹³, as well as DeepSeek¹⁴ that has grabbed global headlines. Most of these systems are released as an open-source or open-weight models. This wave of indigenous innovation coincided with a surge of public interest, experimentation, and investment in GenAI applications for marketing, e-commerce, education, and professional services. For self-employed individuals and small entrepreneurs, these tools could dramatically reduce start-up costs by substituting for specialized labor (e.g., designers, coders, or translators) and by enabling professional-grade outputs without prior technical expertise. Collectively, the ChatGPT shock thus acted as a global productivity catalyst—lowering entry barriers and stimulating new firm creation across sectors and regions.

3 Data and Variables

We use the administrative firm registration data to measure new firm entry and patent data to proxy AI human capital, and we construct fine-grained geo-coded grids based on firm and patent addresses. This section describes our data sources, sample construction, and main variables along with their corresponding summary statistics. Variable definitions and additional details including their construction are provided in Appendix A.

3.1 Data Source and Sample Construction

Firm Registration Data. We obtain firm registration data from the State Administration for Industry and Commerce (SAIC) of the People’s Republic of China. The dataset provides detailed firm-level information, including the company name, registered address, contact information, legal representative, industry classification, firm description, regis-

¹¹<https://xinghuo.xfyun.cn>

¹²<https://hunyuan.tencent.com>

¹³<https://chatglm.cn>

¹⁴<https://www.deepseek.com/en>

tered capital, registration date, operating status, as well as information on executives and both firm and individual shareholders.¹⁵

Our sample covers all firms registered between January 2021 and December 2024. We exclude individual industrial and commercial households and remove duplicate entries. The resulting dataset contains 12,820,211 newly established firms. Following China’s Company Law (2014), we classify a firm as “small” if its registered capital is below RMB 1 million and “large” otherwise. This yields 5,536,247 small firms and 6,562,515 large firms in our sample.¹⁶ Before ChatGPT released in November 2022, a total of 5.95 million new firms were established, including 2.01 million (33.8%) small firms and 3.62 million (60.8%) large firms. In the post-ChatGPT period, new firm registrations rose to 6.87 million, driven by a sharp increase in small-firm entries to 3.53 million (51.4%), while large-firm entries declined to 2.94 million (42.8%).

Patent Data and AI Classification. To measure AI-related innovation in the pre-ChatGPT period and to proxy AI human capital, we obtain patent data from the China National Intellectual Property Administration (CNIPA) and focus on invention patents filed between 2010 and 2019. Following the classification approach of Fang et al. (2025a), we apply the AIPatentSBerta model, a transformer-based algorithm pretrained on the full corpus of U.S. patent documents and fine-tuned on labeled Chinese AI and non-AI patents, to identify AI-related patents. This procedure yields 340,771 AI-related invention patents over the period, which we use to construct measures of local AI innovation intensity and AI-specific human capital.

Spatial Grids and Panel Construction. To spatially index both patents and firm registrations, we construct a fine spatial grid covering mainland China using the H3 geospatial indexing system, a hierarchical framework that tessellates the Earth into hexago-

¹⁵The SAIC dataset has been recently used by Shi et al. (2020); Allen et al. (2024, 2025); Cai et al. (2025).

¹⁶Among the 12,820,211 newly registered firms, 721,449 (5.6%) have missing information on registered capital. As a result, the total number of firms does not equal the sum of small and large firms.

nal grid cells with uniform area and adjacency properties.¹⁷ China spans approximately 18°N–54°N latitude and 73°E–135°E longitude; applying the H3 resolution-7 grid over this range yields 166,156 unique hexagonal cells.

Each firm in our registration dataset is geolocated based on its registered firm address, with latitude and longitude coordinates retrieved from Baidu Maps. Similarly, we extract the geographic location of each patent assignee. We then assign both firms and patents to their respective H3 grid cells, indexed by g . For each cell, we compute the total number of AI-related patents filed between 2010 and 2019, denoted by $AIpat_g$, capturing the intensity of AI innovation activity prior to the emergence of GenAI.

Based on the spatial grid and the spatial indexing of firms and patents, we construct a grid-by-quarter panel spanning 2021Q1 to 2024Q4. This allows us to consistently relate pre-existing AI human capital, as proxied by $AIpat_g$, to subsequent patterns of firm formation at a fine spatial resolution.

3.2 Main Variables

Our analysis is conducted at the H3 grid cell g level. The key variables are: (i) local firm entry, measured by the number of new firm registrations in grid g and quarter t ; (ii) AI-specific human capital, proxied by pre-ChatGPT AI patenting activity in grid g ; (iii) entrepreneurial characteristics, including indicators for serial entrepreneurs; and (iv) organizational structure at founding, captured by the number of shareholders and executive team members.

New Firm Creation. Our main dependent variable is the number of new firm registrations in each grid and quarter. For each grid–quarter, we calculate both the total number of newly registered firms and subcategories by firm size.

¹⁷For details, see <https://h3geo.org/>. We use resolution 7, which corresponds to hexagons with a radius of approximately 2 km and an area of roughly 5 km², comparable to a 2×2 km square.

Pre-ChatGPT Grid-Level AI Human Capital. Our identification strategy hinges on the fact that human capital plays a central role in the diffusion and adoption of AI technologies (Babina et al., 2023; Gofman and Jin, 2024; Andreadis et al., 2025; Cong and Zhu, 2024). To proxy pre-ChatGPT grid level AI-relevant human capital, we use the number of AI-related invention patents filed between 2010 and 2019 in H3 grid cell g , denoted by $AIpat_g$, based on the classification in Section 3.1. This measure captures the ex-ante local stock of AI-specific technological capability and serves as our main measure of AI human capital. Our hypothesis is that areas with higher pre-existing AI expertise are better positioned to absorb and deploy GenAI, with entrepreneurs and workers in these regions more likely to recognize, experiment with, and incorporate GenAI into new firm formation.

For our baseline regression analysis, we construct a binary exposure indicator $HighAI_g$, equal to 1 if $AIpat_g > 0$, i.e., grid g has at least one AI-related patent in 2010–2019, and 0 otherwise. Under this definition, 6.1 percent of grids (10,183 out of 166,156) are classified as high AI exposure. We assess robustness to alternative exposure definitions in Section 6.2.

Serial Entrepreneur. To examine heterogeneity in founder experience, we construct a measure of serial entrepreneurship as a proxy for local entrepreneurial capital. For each newly registered firm, we trace the legal representative across historical records to determine whether the same individual had founded any other firm in the preceding three years. A match indicates prior entrepreneurial experience, classifying the individual as a *serial entrepreneur*.

We then aggregate this information to the grid–quarter level and define *Pct Serial Entrep* as the share of new firm founders in a given cell and quarter who are serial entrepreneurs. To investigate differential patterns by firm size, we separately compute this share for small and large firms, denoted as *Pct Serial Entrep Small* and *Pct Serial Entrep Large*, respectively. These measures allow us to examine whether the rise in firm forma-

tion associated with GenAI is driven by novice founders or those with prior experience navigating the startup process.

Organizational Structure at Entry. To characterize the structure of firms at entry, we extract detailed information from registration records on founding teams. Specifically, for each newly registered firm, we record: (i) the number of shareholders, (ii) the share of individual (i.e., natural-person) shareholders, and (iii) the number of listed executive members. These metrics respectively capture organizational complexity, ownership composition, and managerial team size.

We aggregate these variables to the grid–quarter level by taking simple averages across all newly registered firms in a given cell and quarter. The resulting panel allows us to analyze how firm-level organizational structure varies across grid and time. In extended analyses, we use these measures to examine whether and how GenAI changes how firms are structured at entry—potentially enabling leaner startups.

3.3 Descriptive Statistics

Table 1 reports summary statistics for the main variables used in our analysis. The sample covers 16 quarters from 2021Q1 to 2024Q4 and includes all H3 grid–quarters with at least one firm registration in our sample period, resulting in a balanced panel of approximately 2,658,496 observations. On average, each grid hosts around 4.8 new firms per quarter, with substantial spatial heterogeneity. Roughly 46% of the new entrants are small firms with registered capital below RMB 1 million, while large firms account for the remaining 54%. The average grid records 2 new small firms and 2 new large firms per quarter (mean value 2.1 and 2.5 respectively), indicating that entry is not dominated by micro-enterprises. Figure 1 further reports the quarterly trends of average new firm formation per grid. The figure shows that prior to 2022Q4, the average number of new small firms per grid is stable at around 1.4–1.7 per quarter, while large firms average about

2.6–3.0. After 2022Q4, small firm entry increases to roughly 2.5–3.2 per quarter, whereas large firm entry declines to about 1.3–1.5. This pattern indicates a structural shift in the composition of new firm formation toward smaller entrants.

For variables related to entrepreneur characteristics, we have to restrict attention to grid–quarter observations with at least one new firm. Consequently, the usable sample size declines by around 60.1% to 1.04 million. Summary statistics indicate that the prevalence of serial entrepreneurs is modest: the mean value of *Pct Serial Entrep* is 26.7%, suggesting that just over one in four new founders have previous firm-creation experience. For small firms, *Pct Serial Entrep Small* averages slightly lower at 16.3%, consistent with a greater presence of first-time entrepreneurs in smaller ventures.

The average number of shareholders per firm is 1.5 (1.4 among small firms), indicating that start-ups are typically founded by a very limited number of equity holders. Correspondingly, individual shareholders account for 93.2% of the total on average, and an even higher share of 96.1% among small firms, implying that most start-ups are closely held by a few private owners. The mean size of the executive team is about two people, with limited variation across firm size categories, suggesting that founding teams tend to remain relatively small in both small and larger start-ups.

Turning to the key explanatory variable, AI exposure exhibits pronounced spatial variation and concentration. Figure 2 shows that the distribution of *AIpat_g* is highly right-skewed, with a mean of 2.1 and a median of 0. Most grids host only a handful of AI patents, while a small subset exhibits much higher patent intensity, leading to a long right tail and indicating strong spatial concentration of AI innovation. Under our baseline definition, approximately 6% of grids are classified as high-AI regions (*HighAI_g*=1). These high-AI grids are primarily located in major metropolitan areas such as Beijing, Shanghai, Shenzhen, and Hangzhou.

Overall, the summary statistics confirm that the sample captures substantial variation in firm entry, entrepreneurial composition, and local AI exposure across grid and time. The

data also show that regions with stronger pre-existing AI activity tend to be more urban and economically developed, a pattern we explicitly control for in the regression analysis.

[Insert Table 1 Here]

[Insert Figure 1 Here]

[Insert Figure 2 Here]

Geographic Concentration and Heterogeneity of AI Innovation Activity. We document the spatial distribution of pre-ChatGPT AI patenting. Across grids, we show both regional variation and diversification, as well as pronounced within-city heterogeneity in AI patents. These patterns provide the foundation for our identification strategy, in which local AI patenting serves as a proxy for AI-specific human capital.

Figure 3 visualizes the spatial distribution of AI-related patents and inventors across China at the H3 grid-cell level. Each point represents a 2km-resolution hexagonal cell, with intensity power-normalized to reflect the local concentration of AI activity. As expected, both AI patents and inventors are concentrated in major innovation hubs such as Beijing, Shanghai, Shenzhen, Hangzhou, and Guangzhou. We also find the following features.

First, although AI innovation is disproportionately concentrated in China’s southeastern corridor, it is not narrowly confined to top-tier cities. We observe meaningful *regional diversification*: inland growth poles and homes to dense cluster of high-education institutions such as Chengdu, Xi’an, Wuhan, and Hefei, as well as many second- and third-tier cities, exhibit substantial AI activity. Overall, the spatial spread of AI patents and inventors aligns closely with the economically active zone *east of the Hu Line* (also known as the Heihe–Tengchong Line), a well-known geographic demarcation separating China’s densely populated and industrialized eastern region from its sparsely populated west.¹⁸

¹⁸The Hu Line traces a line from Heihe in the northeast to Tengchong in the southwest. It roughly divides China into a highly urbanized, economically developed eastern region (comprising only around 40% of land area but home to around 94% of the population) and a less-developed western hinterland.

Second, and more critically for our identification strategy, there exists substantial *within-city heterogeneity* in AI exposure. Even within the same metropolitan area, certain H3 grids exhibit intense innovation activity, while nearby grids remain largely inactive. To illustrate this variation, we zoom in on three representative innovation hubs—Beijing, Shanghai, and Shenzhen—and plot the spatial distribution of AI patents at the grid level. As shown in Figure B.1, even within these leading technology centers, some neighborhoods exhibit dense clusters of AI patenting prior to the GenAI era, while adjacent grids display little to no AI activity. This fine-grained variation provides the foundation for our empirical strategy: it allows us to flexibly absorb unobserved city-level shocks while retaining variation in AI activity, and it serves as a valid source of cross-grid variation in AI human capital. By comparing grids that share similar institutional, regulatory, and macroeconomic environments but differ sharply in pre-existing AI intensity, we can identify the local entrepreneurial response to the diffusion of GenAI.

[Insert Figure 3 Here]

4 Empirical Strategy

4.1 Baseline Specification

As detailed in Section 3.1, we construct a quarterly panel at the grid-cell level, measuring the number of newly registered firms in each hexagonal H3 cell. To estimate the effect of GenAI on firm formation, we exploit cross-sectional variation in pre-2020 exposure to AI-specific human capital, as captured by the intensity of AI patents filed by firms or inventors within each grid. Grids with high concentrations of AI patents serve as treated units, while those with low or no prior exposure form the control group.

The November 2022 release of ChatGPT provides a sharp and plausibly exogenous technological shock that dramatically increased public awareness and experimentation

with GenAI tools. Our empirical strategy compares changes in firm formation between high- and low-AI (exposure) grids before and after this event, relying on fine-grained *within-city* variation to isolate local treatment effects. The baseline difference-in-differences specification is as follows:

$$Y_{gt} = \beta (\text{Post}_t \times \text{HighAI}_g) + \mu_{g \times q(t)} + \lambda_{c(g) \times t} + \varepsilon_{gt}, \quad (1)$$

where Y_{gt} denotes the number of newly registered firms in grid g at time t ; HighAI_g equals one for grids with high pre-2020 AI patenting intensity; and Post_t equals one for quarters following ChatGPT’s release (2022Q4 onward). The function $q(t)$ maps time t into its calendar quarter, and $c(g)$ denotes the city containing grid g . The coefficient of the interaction term β measures how much more firm creation changes in AI-intensive grids relative to less exposed ones within the same city after ChatGPT’s release.

The specification includes both *grid-by-calendar-quarter* fixed effects ($\mu_{g \times q(t)}$) and *city-by-quarter* fixed effects ($\lambda_{c(g) \times t}$). This high-dimensional fixed-effect structure absorbs (i) persistent grid-specific seasonality and any time-invariant unobservables within each grid, and (ii) any time-varying city-level shocks, such as local policy initiatives, changes in the business environment, industrial restructuring, or macroeconomic fluctuations. Standard errors are clustered at the city level. Identification thus arises exclusively from *within-city*, *within-quarter* differences across neighboring grids that differ in pre-existing AI exposure.

This spatially granular design offers several identification advantages. By comparing grids within a common institutional and economic context, we eliminate bias from unobserved regional heterogeneity and policy differences across cities. The grid-level fixed effects control for micro-spatial factors such as persistent differences in local infrastructure or land use, while the city-by-quarter fixed effects purge all contemporaneous shocks at the city level. The global and unanticipated nature of the ChatGPT release strengthens the quasi-experimental setting, providing a sharp and plausibly exogenous timing for

treatment.

To evaluate remaining identification concerns, we perform event-study tests for parallel pre-trends in Section 4.3 and a battery of robustness and placebo tests in Section 6 including placebo regressions using non-AI or residualized patent measures, and random reassignment of treatment labels, all of which consistently yield null or attenuated effects. These exercises confirm that the estimated coefficients capture the causal impact of GenAI through localized AI-specific human capital, rather than reflecting general entrepreneurial pattern or spurious spatial correlations.

4.2 Heterogeneous Effects by Firm Size

As we argue, GenAI, as a general-purpose technology, redefines the economics of entrepreneurship. It automates language-based information-processing (reasoning, synthesis, coding, and content generation), alongside creative, managerial, and operational tasks, thereby lowering the barriers to capabilities that were once capital-intensive or concentrated among specialized labor and well-resourced organizations. In practice, GenAI allows a very small founding team, or even a single entrepreneur, to handle functions that previously required multiple specialists in product development, marketing, operations, and customer support. This transformation shifts the production frontier of what lean ventures can achieve and challenges traditional assumptions about firm size, scale, and the division of labor. A growing wave of *GenAI-native* firms—such as *Midjourney*, *Cursor*, and *Perplexity*—illustrates this paradigm shift. These startups achieve rapid growth with minimal staff and modest capital by embedding generative models directly into design, engineering, and marketing workflows. In this sense, GenAI operates as a “digital co-founder,” lowering the minimum viable scale of new ventures.

Building on this intuition, we hypothesize that GenAI disproportionately lowers entry barriers for small and resource-constrained firms relative to larger incumbents. The mechanism rests on differences in fixed costs and internal capabilities at the early stages

of firm formation. Traditional firm creation requires access to specialized expertise, such as in software development, marketing, legal, and operational management, creating a fixed cost structure that weighs most heavily on startups. GenAI relaxes these constraints by automating or augmenting many of these tasks, allowing small teams or individual entrepreneurs to substitute AI capabilities for costly human capital.

This mechanism is consistent with broader theoretical insights from the entrepreneurship literature. Specifically, in standard models of occupational choice under frictions (e.g., Evans and Jovanovic, 1989; Kerr and Nanda, 2015), reductions in entry costs expand the extensive margin of entrepreneurship by allowing financially or skill-constrained agents to overcome fixed investment thresholds. If GenAI compresses these thresholds, we should observe greater firm entry in regions with higher effective access to GenAI, particularly among smaller firms and first-time entrepreneurs.

We test this prediction empirically by estimating heterogeneous treatment effects by firm size. Specifically, we extend our baseline difference-in-differences specification 1 to separate small and large entrants:

$$Y_{gt}^{\text{size}} = \beta^{\text{size}} (\text{Post}_t \times \text{HighAI}_g) + \mu_{g \times q(t)} + \lambda_{c(g) \times t} + \epsilon_{gt}^{\text{size}}, \quad (2)$$

where Y_{gt}^{size} the number of new firms of a given size category (small or large) in grid g and time t . We classify firms as small or large based on the RMB 1 million registered-capital cutoff described in Section 3.1, and we report robustness checks using alternative capital thresholds in Section 6.2.4. If GenAI primarily reduces fixed costs or substitutes for managerial labor, we would expect $\beta^{\text{small}} > 0$ and $\beta^{\text{large}} \leq 0$.

4.3 Dynamic Model

To examine the temporal evolution of treatment effects and validate the parallel trends assumption, we estimate an event-study specification of the following form:

$$Y_{gt} = \sum_{k \neq -1} \beta_k \cdot 1\{t - t_0 = k\} \times \text{HighAI}_g + \gamma_g + \lambda_{c(g) \times t} + \varepsilon_{gt}, \quad (3)$$

where t_0 denotes the reference quarter (2022Q4), and each coefficient β_k captures the difference-in-differences-style effect: the difference in outcomes between high- and low-AI exposure grids k quarters away from the baseline, relative to the reference period $k = -1$ (the quarter immediately preceding ChatGPT’s release). Grid fixed effects γ_g control for time-invariant heterogeneity across locations, while city-by-quarter fixed effects $\lambda_{c(g) \times t}$ flexibly absorb common time-varying shocks at the city level.

We plot the estimated event-time coefficients $\hat{\beta}_k$ with 99% confidence intervals in Figure 4. Under the parallel trends assumption, pre-treatment estimates ($k < 0$) should be statistically indistinguishable from zero, while significant post-treatment effects ($k \geq 0$) would indicate a divergence in firm formation between high- and low-exposure grids following the release of GenAI technologies.

[Insert Figure 4 Here]

5 Empirical Results and Mechanisms

5.1 Baseline results

Table 2 presents the baseline difference-in-differences estimates of the impact of genAI on new firm creation based on Equations (1) and (2). Across all specifications, the coefficient on the interaction term $Post \times HighAI$ is positive and statistically significant, indicating that firm entry increased more strongly in grids with high pre-existing AI activity following the release of ChatGPT in November 2022. The magnitude of the effect is economically meaningful. Column 1 shows the results based on Equation (1) (controlling for grid-by-calendar-quarter and city-by-quarter fixed effects), and the estimated coefficient on $Post \times HighAI$ is 5.038 (s.e.=1.395), significant at the 1% level. This suggests

that, all else equal, high-AI grids experienced an average increase of about five additional new firms per grid-quarter following the introduction of ChatGPT. Aggregating across all 10,183 high-AI grids and 8 quarters after the release of ChatGPT, this effect implies approximately $5.038 \times 10,183 \times 8 \approx 0.41$ million additional firm entries post ChatGPT, accounting for around 6.0 percent of the total 6.87 million new firms established nationwide during the same period.

Columns 2-3 separate the estimates by firm size based on Equation (2). Consistent with the hypothesis that GenAI primarily reduces fixed costs and substitutes for scarce managerial labor, the positive effect is entirely driven by small firms. The estimated coefficient for small-firm entry is positive and significant, whereas the corresponding coefficient for large firms is negative and significant. Correspondingly, the coefficient on the interaction term for small firms is 7.704, suggesting that all else equal, high-AI grids experienced an average increase of approximately eight additional firms per quarter. This divergence suggests that GenAI tools, by automating core tasks like marketing, customer interaction, and software development, compress the scale required to operate a business. Entrepreneurs who might previously have needed larger teams and capital to launch may now find it feasible to enter as smaller, more agile firms. The decline in large firm entry in high-AI grids may reflect this substitution toward leaner organizational forms, rather than a decline in overall entrepreneurial activity. Since all regressions exploit within-city variation and absorb city-level shocks, the results are not confounded by local policy interventions or macroeconomic trends. Instead, they capture how heterogeneous AI exposure shapes local responses to a frontier general-purpose technology. Overall, these results suggest that the diffusion of GenAI disproportionately stimulated entrepreneurship among small and resource-constrained firms, supporting the view that GenAI functions as a “digital cofounder” that lowers entry barriers.

[Insert Table 2 Here]

Table 3 reports robustness checks that winsorize the dependent variable to mitigate the

influence of extreme values. Columns 1–3 winsorize within quarter, while Columns 4–6 winsorize across the full sample. Across all specifications, the interaction term $Post_t \times HighAI_g$ remains statistically significant and economically sizable. For total new firm creation, the coefficient is 3.692 (s.e. = 0.826) in Column 1 and 2.518 (s.e. = 0.667) in Column 4, suggesting increase of approximately three to four additional firms (52.2 to 76.5 percent relative to the sample mean of new firms).

Columns 2–3 and 5–6 report results separately by firm size. Consistent with the mechanism highlighted in Table 2, the post-ChatGPT rise in firm entry is concentrated among small firms: the coefficients are 6.170 (s.e. = 0.800) with within-quarter winsorization and 4.911 (s.e. = 0.528) with full-sample winsorization. By contrast, large-firm entry declines: the coefficients are -2.928 (s.e. = 0.332) and -2.922 (s.e. = 0.336). Taken together, these robustness checks confirm that the baseline effects are not artifacts of outliers or specific fixed-effects choices: GenAI is associated with a sizeable increase in small-firm formation and a corresponding decline in large-firm entry within the same city and quarter.

[Insert Table 3 Here]

5.2 Industry Heterogeneity

Table 4 reports industry-specific difference-in-differences estimates of the interaction term $Post \times HighAI$, examining how the impact of GenAI on firm entry varies across industries. The specification is identical to our baseline model (1) but estimated separately by industry. Panel A reports the top fifteen industries ranked by the estimated coefficients. The results reveal substantial heterogeneity in the responsiveness of firm entry to GenAI. The largest effects are observed in the *Retail*, *Business Services*, and *Technology Promotion and Application Services* sectors, with coefficients of 1.625 (s.e. = 0.545), 0.977 (s.e. = 0.292), and 0.871 (s.e. = 0.231), respectively, all statistically significant at the 1% level. These magnitudes indicate that grids with higher pre-existing AI activity ex-

perience stronger post-ChatGPT increases in firm formation within commercially oriented and service-intensive industries.

Several additional sectors also exhibit positive and significant coefficients, including *Wholesale, Entertainment, Catering, and Culture and Arts*, suggesting that GenAI facilitated new firm creation in customer-facing and creative domains where automation and content generation tools directly augment production. By contrast, Panel B reports the bottom fifteen industries ranked by the estimated coefficients where traditional and capital-intensive industries including construction and manufacturing exhibit small or even negative effects. These patterns reinforce the interpretation that the entrepreneurial impact of GenAI operates primarily through demand-side adoption and creative applications. Overall, Table 4 underscores that the sectors whose core of business activities rely on knowledge work, creativity, marketing, and digital service are most responsive to GenAI.

[Insert Table 4 Here]

Sectoral heterogeneity in AI-Relevance Scores. We further exploit the rich information in the firm registration data to study industry heterogeneous effects. Since GenAI affects industries differently, we first use the official 96 industry labels to classify firms by sector. To obtain a more AI-specific measure of exposure at sectoral level, we then analyze firms’ business descriptions: we extract keywords using topic methods and adopt a state-of-the-art language model to assign three continuous scores to each firm: an *AI-upstream* score (the extent to which the firm’s activities relate to upstream AI, such as foundation model development, data infrastructure, or computing resources), an *AI-downstream* score (the extent to which the firm operates in downstream AI uses, such as application development, product integration, or end-user services) and an *AI entrepreneurship helpfulness* score, which captures the degree to which GenAI can assist and support a firm’s core business activities or lower key entrepreneurial frictions, for example by speeding up product prototyping and iteration, helping inexperienced founders perform specialized tasks (e.g.,

basic coding, legal/financial drafting, market research), and facilitating content creation, personal branding, low-cost experimentation, or digital marketing. All three scores range from -100 to 100 , with higher values indicating stronger relevance or complementarity with GenAI.

The construction proceeds in two stages. First, we represent each firm’s business scope as a weighted combination of latent *textual factors* derived from its official business description following the approach by Cong et al. (2025). These topics summarize the underlying business activities and provide a less noisy representation of its underlying business than the raw text. Second, we employ the state-of-the-art large language model GPT-4o (Hurst et al., 2024) to evaluate each topic’s relevance to AI across multiple dimensions. Each firm’s AI relevance score is then computed as a weighted average of these topic-level evaluations, where the weights correspond to the firm’s topic loadings. We validate the measure by manually reviewing a random subset of firms, i.e., examining their business websites, product descriptions, and public filings, and find a close correspondence between these qualitative assessments and our model-based scores. Full construction details and validation procedures are provided in Appendix A.2.

As Table 1 shows, around 20% of the new entrants have high (compared to low) AI upstream scores; 52% of the new entrants have high (compared to low) AI downstream scores; 75% of the new entrants have high (compared to low) AI entrepreneurship scores. By examining how the effects vary across sectors along these three AI-related dimensions, we characterize industry-level heterogeneity in exposure to GenAI.

Table 5 presents the results. Columns 1–4 use alternative dependent variables: the number of high-upstream, low-upstream, high-downstream, and low-downstream new firms in each grid-quarter, respectively, where “high” denotes a score > 0 and “low” a score ≤ 0 . Columns 5–6 similarly use as dependent variables the number of new firms with high versus low AI-entrepreneurship helpfulness (score > 0 vs. ≤ 0). Across specifications, the interaction term $Post \times HighAI$ is positive and statistically significant in AI-

relevant firms. The coefficients are much larger for low-upstream (4.199 with s.e. = 1.225) and high-downstream (3.323 with s.e. = 0.789) firms than for high-upstream (0.846 with s.e. = 0.388) and low-downstream (1.722 with s.e. = 0.753) firms, suggesting that the diffusion of GenAI has primarily stimulated entry in industries that adopt or apply AI rather than those that develop it. Upstream industries (such as semiconductors, cloud infrastructures, and data centers) remain capital- and expertise-intensive, limiting new entry even after the release of ChatGPT. By contrast, downstream industries—such as digital marketing, content generation, and business software integration—can more readily embed off-the-shelf generative models into their products and workflows.

Columns 5–6 examine heterogeneity by AI entrepreneurship helpfulness. The coefficient for high-entrepreneurship sectors is 5.892 (s.e. = 1.302), significant at the 1% level, while that for low-entrepreneurship industries is negative (−0.848, s.e. = 0.300). Recall that the entrepreneurship helpfulness score captures the extent to which GenAI can support a firm’s core activities or relax key entrepreneurial frictions. This contrast between high- and low-entrepreneurship sectors suggests that GenAI most strongly stimulates entrepreneurial activity in sectors that are both technologically close to AI and naturally suited to individual or small-team creative work, such as digital marketing, software development, and data analytics/ By contrast, sectors with weaker entrepreneurial or technological complementarities to AI exhibit little or even negative response. Overall, these results reinforce the view that the diffusion of GenAI operates through local technological and human-capital complementarities, amplifying entrepreneurial activity precisely where AI capabilities are most easily integrated into new firm creation.

[Insert Table 5 Here]

5.3 Mechanisms

The results above show that the diffusion of GenAI substantially increased new firm formation, particularly among small and resource-constrained entrants. In this subsection,

we examine the mechanisms behind this effect. GenAI may stimulate entrepreneurship by relaxing key constraints that traditionally hinder new entry: (i) *experience-related constraints*, as AI tools substitute for the skills, and managerial experience typically accumulated by older or serial founders; (ii) *financial constraints*, by reducing the need to pool capital from multiple shareholders to launch a new firm; and (iii) *labor constraints*, by automating managerial and operational tasks that would otherwise require hiring additional executives. We test these channels using proxies for founders’ prior experience in Table 6, number of shareholders in Table 7, and executive team size in Table 8.

Table 6 presents evidence on the entrepreneurial experience channel. The dependent variable is the average share of serial entrepreneurs defined in Section 3.2—legal representatives who had established at least one other firm within the preceding three years—within each grid-quarter. Column 1 reports estimates for the full sample, while Columns 2–3 distinguish small and large firms. Across specifications, the interaction term $Post \times HighAI$ is negative and statistically significant in the aggregate, indicating that the post-ChatGPT increase in firm entry was primarily driven by first-time entrepreneurs rather than experienced repeat founders. In Column 1, the coefficient of -0.405 (s.e. = 0.201) suggests that the share of serial founders declined by roughly 1.5 percent relative to its pre-ChatGPT mean of 27.5 percent.

Columns 2–3 further reveal that this decline is concentrated entirely among small firms: the coefficient for small-firm founders is -2.456 (s.e. = 0.266), highly significant, whereas the effect for large firms is positive at 0.648 (s.e. = 0.294). These patterns suggest that GenAI has lowered entry barriers associated with prior entrepreneurial experience, enabling individuals without a founding track record to establish new firms. In high-AI regions, ChatGPT and related tools appear to substitute for experiential knowledge—offering guidance, drafting documents, and automating early-stage tasks that would otherwise require seasoned founders or extensive professional networks.

[Insert Table 6 Here]

Table 7 reports estimates for the financing channel, examining whether GenAI reduces the need to pool capital among multiple shareholders at entry. Panel A focuses on the number of shareholders per new firm. Across all specifications, the coefficient on $Post \times HighAI$ is negative and statistically significant, indicating that firms founded in high-AI grids after the release of ChatGPT are created with fewer shareholders. In Column 1, the coefficient of -0.0211 (s.e. = 0.00747) implies a decline of roughly 1.4% relative to the mean number of shareholders (1.51). When splitting by firm size in Columns 2-3, the decline is present for both small and large firms, consistent with the idea that GenAI lowers the initial financing and coordination burden regardless of firm scale. These findings suggest that founders in AI-exposed regions require fewer co-investors or partners to mobilize the minimum viable capital necessary to launch a business.

Panel B examines whether the composition of shareholders shifts toward individual owners. Across all columns, the estimated effects are small in magnitude and statistically insignificant. The lack of meaningful change in ownership composition indicates that the decline in the number of shareholders is not driven by substitution between individual and corporate shareholders. Instead, the results point to a broader reduction in financial frictions at entry: firms can be launched by fewer equity founders without altering the underlying source of ownership capital. Taken together, Table 7 supports the mechanism that GenAI relaxes financing constraints by lowering the amount of human and financial capital required to start a firm. This evidence complements the experience channels documented earlier, suggesting that GenAI enables more streamlined and capital-lean firm formation.

[Insert Table 7 Here]

Table 8 evaluates the labor substitution channel by examining whether GenAI reduces the number of executive members required to launch a new firm. The dependent variable is the average size of the executive team in each grid-quarter. Across all specifications, the coefficient on $Post \times HighAI$ is negative and statistically significant, indicating that

firms established in high-AI grids after the release of ChatGPT rely on fewer executives at entry. In Column 1, the estimate of -0.0161 (s.e. = 0.00392) represents a decline of around 0.8% relative to the pre-ChatGPT mean executive team size of 2.032.

Columns 2-3 reveal that the decline is concentrated among small firms. The coefficient for small-firm executive team size is -0.0187 (s.e. = 0.00401), highly significant, whereas the effect for large firms is small in magnitude and statistically insignificant. These patterns suggest that GenAI substitutes for early-stage managerial labor, enabling new ventures—particularly smaller ones—to operate with leaner founding teams. Along with the declines in serial entrepreneurship (Table 6) and the number of shareholders (Table 7), the results provide consistent evidence that GenAI relaxes multiple complementary input constraints in firm formation, allowing less-experienced and capital-constrained founders to enter with smaller founding teams and lower organizational complexity.

[Insert Table 8 Here]

5.3.1 Serial Entrepreneurs and the Substitution Towards Smaller Firms

Our baseline analysis shows that the diffusion of GenAI is associated with a significant rise in the entry of small firms alongside a significant decline in the number of newly established large firms. This shift in the composition of new ventures is consistent with a *substitution effect* in entrepreneurial strategy.

One natural interpretation is that GenAI lowers fixed costs and substitutes for creative, information-processing, and managerial labor, enabling entrepreneurs to launch viable businesses with fewer resources as documented in previous section. Even experienced founders—those with the capital and capacity to start larger firms—may instead choose to establish smaller, more agile ventures in the ChatGPT era. By automating key functions such as marketing, customer engagement, and software development, GenAI reduces the minimum efficient scale required to operate a business. In this subsection, we directly test this substitution hypothesis by examining whether serial entrepreneurs downsize the

scale of their new ventures following ChatGPT’s release, particularly in regions with high pre-2019 AI-patent activity, where local founders possess more AI-relevant human capital and are more likely to effectively leverage GenAI tools.

For each new firm established by a serial entrepreneur, we compare its scale with that of the entrepreneur’s most recent prior venture, using *registered capital*—a standardized and regulatorily disclosed measure of initial firm size—as our proxy for entry scale. We then assess whether the relative firm size of new establishments declines more sharply after ChatGPT’s release in high-AI exposure grids. We estimate the same difference-in-differences specification as in Equation (1), but with alternative dependent variables that capture the *relative* scale of new firms founded by serial entrepreneurs in grid g and quarter t . Specifically, in *Panel A* of Table 9, the dependent variable Y_{gt} denotes the percentage of new firms established by a serial entrepreneur in grid g and quarter t whose registered capital exceeds that of the entrepreneur’s previous firm. In *Panel B* of Table 9, Y_{gt} denotes the grid-quarter average ratio of registered capital in the new firm established by a serial entrepreneur in grid g and quarter t to that in the previous firm, trimmed at the top 1% and bottom 1% to mitigate outlier influence. The interaction term $Post_t \times HighAI_g$ captures the differential change in new-firm scale among serial entrepreneurs in AI-intensive grids following ChatGPT’s release.

Table 9 shows the results. In *Panel A*, the coefficient on $Post \times HighAI$ is consistently negative and statistically significant, indicating that serial entrepreneurs in high-AI regions tend to reduce firm size after the release of ChatGPT. The effect is particularly strong among small new firms below the threshold of 1 million RMB: the estimate of -5.416 (s.e. = 0.483) in Column 2 implies that the share of new firms whose registered capital exceeds that of the entrepreneur’s previous firm decreases by approximately 5.4 percentage points. Even for large firms, the coefficient remains negative (-0.917 , s.e. = 0.241), suggesting a modest but statistically significant tendency toward smaller scale.

Panel B corroborates this pattern using the average ratio of new-to-old registered cap-

ital as a continuous measure of downsizing. The estimated coefficient on the interaction term $Post \times HighAI$ is -1.983 (s.e. = 0.560) for the overall sample, and the magnitude becomes substantially larger for small firms, at -7.100 (s.e. = 1.148). This result indicates that, following the release of ChatGPT, serial entrepreneurs in high-AI grids started new firms with substantially smaller capital relative to their previous firms, compared with entrepreneurs in low-AI grids. On average, new firms are about seven times smaller than the previous ones in high-AI regions. The negative and statistically significant effect indicates a clear contraction in startup scale among repeat founders during the era of GenAI. Even among large firms, the estimated coefficient remains negative (-0.180 , s.e. = 0.023), suggesting that the downscaling tendency extends across the size distribution—albeit more modestly for capital-intensive ventures.

[Insert Table 9 Here]

Overall, these findings provide direct evidence of a substitution mechanism: experienced entrepreneurs are recalibrating the optimal scale of new ventures in response to the productivity gains afforded by GenAI. Rather than reducing overall entrepreneurial activity, GenAI appears to *reshape* the landscape—enabling smaller, more agile firms to enter markets once dominated by larger, capital-intensive entities and shifting new firm creation toward leaner organizations, especially in high-AI regions where entrepreneurs are likely to adopt the GenAI tools. This response among repeat founders underscores the role of GenAI as a “digital cofounder” that lowers the managerial and financial thresholds for market entry, thereby compressing the scale distribution of new firms while amplifying the overall rate of business creation.

6 Placebo and Robustness Tests

6.1 Placebo Tests

While our baseline specification absorbs grid-level and city-by-quarter fixed effects, we implement a set of placebo tests designed to further probe the interpretation of our treatment variable as capturing *AI-specific human capital*—rather than existing local dynamism in innovation and entrepreneurship in this section.

6.1.1 AI-Relevant Human Capital vs. General Innovation Capacity

One hypothesis in our empirical design is that pre-existing AI patent activity proxies for localized *AI-relevant human capital*—that is, the presence of technical talent and domain-specific expertise capable of adopting and productively leveraging GenAI tools. This design underpins our interpretation of the treatment effect: the post-ChatGPT surge in firm entry in high-AI grids arises not merely from general innovation intensity, but from the complementarity between GenAI and pre-existing AI-specific capabilities.

To evaluate this channel more directly, we conduct a placebo test using non-AI patenting activity. The objective is to distinguish the effects of AI-specific knowledge from those driven by broader innovation capacity. If our proposed mechanism is correct, then a similar difference-in-differences design based on non-AI patent strength should yield attenuated or null effects—since general patenting captures innovation capacity not specifically related to AI.

We implement the test in two steps. First, we regress the log number of non-AI patents (filed between 2010 and 2019) on the log number of AI patents at the grid level to remove the shared variation between general and AI-specific innovation activity:

$$\log(\text{nonAI_patents}_g + 1) = \alpha + \beta \log(\text{AI_patents}_g + 1) + \varepsilon_g. \quad (4)$$

We then classify grids into a *High nonAI* group if their residual $\hat{\varepsilon}_g$ lies above the 75th percentile—capturing locations with stronger-than-expected non-AI patenting conditional on AI intensity. In other words, these are regions with technological strength *unrelated* to AI. We re-estimate our baseline DiD model with $Post \times High\ nonAI$ interaction term instead.

While this placebo test helps isolate the role of AI-specific human capital, we interpret the results conservatively. If our baseline effects are indeed driven by entrepreneurs with AI-relevant expertise—those best positioned to adopt GenAI tools—then the DiD coefficient based on *High nonAI* should shrink substantially relative to the baseline. While we do not expect it to fall exactly to zero—owing to potential measurement error or minor alternative channels—the attenuation itself provides meaningful support for our proposed mechanism. This mirrors the logic of our complementary placebo test using residualized firm counts, where removing the component correlated with AI innovation allows us to better identify the contribution of AI-specific capabilities rather than general entrepreneurial dynamism.

Table 10 presents the results. Consistent with our hypothesis, the coefficient on $Post \times High\ nonAI$ is statistically indistinguishable from zero in Columns 1-2, which examine total firm entry and small firm formation. For example, in Column 1, the coefficient is 0.159 (s.e. = 0.478), while in Column 2 it is 0.558 (s.e. = 0.485)—both near zero and not significant. Column 3 shows a marginally significant effect for large firms (−0.427, s.e. = 0.249), but its magnitude is an order of magnitude smaller than the effect observed for AI-intensive grids.

[Insert Table 10 Here]

These findings reinforce the interpretation that the economic impact of GenAI operates primarily through AI-specific human capital rather than general innovation capacity. In particular, the muted response in regions with strong non-AI innovation capacity suggests that it is not general innovation capacity *per se*, but rather the relevance of that capacity to

AI, that governs local adoption and entrepreneurial response. This placebo test thus adds credibility to our identification strategy and supports the mechanism of AI-human capital complementarity as the primary channel through which GenAI affects firm formation.

6.1.2 Residualizing AI Exposure to Address Pre-Existing Entrepreneurial Activity

Another potential concern is that grid cells with higher AI patent activity may have had higher levels of entrepreneurship prior to the diffusion of GenAI. In that case, our baseline estimates might partially reflect pre-existing local dynamism or unobserved fundamentals, rather than the causal effect of AI-relevant human capital.

To address this concern, we construct a residualization-based placebo test that isolates the component of entrepreneurship uncorrelated with AI-related innovation. Specifically, we regress the logarithm of the average number of new firms in each grid prior to 2019 on the logarithm of the number of AI invention patents filed in 2010–2019:

$$\log(\text{FirmEntry}_g^{\text{pre}} + 1) = \alpha + \delta \cdot \log(\text{AI.Patents}_g + 1) + \varepsilon_g, \quad (5)$$

where the residual $\hat{\varepsilon}_g$ captures grid-level entrepreneurial intensity that is linearly uncorrelated with local AI patent activity. We then define a binary indicator, HighResid_g , equal to one for grid cells with above-median residuals, and re-estimate our main difference-in-differences specification, replacing HighAI_g with HighResid_g .

Table 11 presents the results. We find that the interaction term $\text{Post} \times \text{HighResid}$ remains statistically significant—particularly for small firms (Column 2)—but its magnitude is substantially attenuated relative to the baseline. For example, the coefficient for small-firm entry declines from 7.704 in the main specification in Table 2 to 1.216 here, representing an approximate 85% reduction in estimated effect size. Similarly, the coefficient for large-firm entry shrinks in magnitude from -3.120 to -0.481 , also reflecting a roughly 85% attenuation.

[Insert Table 11 Here]

One thing worth mentioning is that we do not expect this placebo test to yield a null effect. Several factors explain why some residual significance may persist. First, AI patents are a proxy: residual entrepreneurial vibrancy may still embed unmeasured AI-relevant human capital (e.g., in IT services or software startups). Second, the number of pre-2019 firms in a grid—especially in digitally intensive industries—may correlate with latent technological capabilities. Third, residuals from linear regression cannot fully purge nonlinear relationships.

Nevertheless, the attenuation of treatment effects, combined with consistent findings across other placebo and robustness tests, suggests our interpretation that *AI-specific human capital*, not general local vitality, drives the observed post-ChatGPT surge in business formation. This result strengthens the mechanism-based reading of our findings: GenAI facilitates new firm entry most where complementary technical expertise is already in place.

6.1.3 Random Assignment of AI Exposure Labels

To further probe the validity of our identification strategy, we implement a placebo test based on random assignment. The goal of this test is to assess whether our main results could be spurious—i.e., whether the observed post-ChatGPT differences in firm formation across high- and low-AI grids are merely artifacts of sample variation, unrelated to actual AI exposure.

Specifically, we randomly assign the *HighAI* label to a set of grids such that the distribution of treated and control observations matches the original sample. We then re-estimate our baseline difference-in-differences specification using these placebo labels. Since this artificial assignment severs any real link between local AI capabilities and GenAI adoption, we would expect no systematic post-period difference in firm entry between the pseudo-treated and control grids.

This placebo test offers a strong falsification check: if our baseline results were driven by chance correlations, model overfitting, or latent spatial trends, we would observe spurious treatment effects even under random assignment. Conversely, the absence of significant coefficients under placebo supports the claim that the original results are driven by meaningful variation in AI-relevant human capital.

We re-estimate our baseline difference-in-differences specification 100 times. Figure 5 plots the distribution of the resulting $\text{Post}_t \times \text{HighAI}_g$ coefficients. The placebo estimates are tightly centered around zero and statistically insignificant across all outcomes—including total firm entry, small business formation, and large firm entry. The mean coefficient across simulations is 0.064 with a standard deviation of 0.166, indicating no systematic treatment effect under random assignment. This absence of spurious significance reinforces the credibility of our baseline findings and underscores that the observed entrepreneurial response is driven by genuine pre-ChatGPT AI exposure—captured by local AI patenting activity—rather than by chance correlations or unobserved spatial trends.

[Insert Figure 5 Here]

6.2 Additional Robustness Tests

6.2.1 Excluding First-Tier Provinces

A potential concern with our baseline results is that they may be disproportionately driven by a few highly developed regions with concentrated AI activity and vibrant entrepreneurial ecosystems—particularly Beijing, Shanghai, and Guangdong. These provinces account for a large share of China’s innovation output and firm registrations and may exhibit dynamics that are not representative of the broader national landscape.

To address this concern, we re-estimate our difference-in-differences specification after excluding all grid cells located in these three first-tier provinces. Table 12 presents the

results. Column 1 replicates our baseline specification using the full sample of firm entry as the dependent variable, with city-by-quarter and grid-by-calendar-quarter fixed effects. Columns 2-3 examine entry by small and large firms respectively.

Reassuringly, the estimated coefficients on the interaction term $Post \times HighAI$ remain positive, statistically significant, and similar in magnitude to those in the full-sample analysis. For example, Column 1 reports a coefficient of 4.0 (s.e. = 1.54), while Column 2 shows that the effect on small firm formation remains sizable at 6.3 (s.e. = 1.17) and significant. The effect on large firms, reported in Column 3, remains negative and significant (−2.7, s.e. = 0.76), consistent with our earlier findings. These results confirm that our main findings are not driven solely by the behavior of a few high-profile innovation hubs. Instead, the observed post-ChatGPT surge in new firm formation—particularly among small entrants in AI-exposed areas—represents a broader pattern that extends beyond China’s most developed provinces.

[Insert Table 12 Here]

6.2.2 Restricting to AI-Active Grids

A natural concern with our baseline difference-in-differences design is that the contrast between high- and low-AI exposure grids may in practice be driven by whether a grid has any AI activity at all. In the pre-treatment period, a large share of grids report zero AI patenting activity. As a result, our treatment indicator $HighAI_g$ may conflate intensity of AI exposure with a binary distinction between treated and untreated regions, potentially introducing unobserved differences in baseline entrepreneurship potential.

Table B.1 presents a robustness check where we restrict the sample to grids that had at least one AI patent prior to the release of ChatGPT. This addresses the concern that our baseline difference-in-differences estimates may conflate the entry response in AI-active areas with structurally different grids that had no prior AI activity. By limiting the sample to AI-active grids, we focus the comparison on areas with varying intensities of prior AI

exposure—distinguishing between “high” versus “moderate” AI patenting activity. In this specification, $HighAI_g$ is defined as an indicator equal to 1 for grids with more than eight AI patents (the 75th percentile).

The treatment effect remains economically and statistically significant. Column 1 shows that total firm entry in high-AI grids increased by 6.568 firms per grid–quarter relative to moderately exposed grids after the release of ChatGPT (s.e. = 1.962, $p < 0.01$). Column 2 reveals that this effect is even larger for small firms, with an estimated coefficient of 11.62 (s.e. = 2.815, $p < 0.01$), indicating that the acceleration in firm formation is particularly concentrated among smaller, more agile entrants. In contrast, Column 3 shows a statistically significant and economically meaningful decline in large firm formation, with a coefficient of -5.859 (s.e. = 1.942, $p < 0.01$). These effects remain robust across specifications with city-by-quarter and grid-by-calendar-quarter fixed effects.

Overall, this exercise confirms that our main results are not driven by differences between AI-active and AI-inactive areas. Even among locations that had already participated in AI innovation prior to the diffusion of GenAI, those with more intense AI activity experienced significantly stronger boosts in small business formation—and larger declines in large firm entry—following ChatGPT’s release. This pattern reinforces the interpretation that GenAI serves as a complementary force to localized AI knowledge capital, especially in enabling smaller firms to overcome startup frictions and scale new ventures.

Moreover, the dynamic analysis presented earlier (Figure 4) suggests no evidence of pre-treatment divergence between high- and low-AI exposure grids. This lends further support to the identification strategy and mitigates concerns that results are mechanically driven by initial differences in innovation activity or omitted local fundamentals. Overall, this robustness check demonstrates that our core findings are not artifacts of treatment definition or driven by the presence of non-AI grids in the control group.

6.2.3 Matched Comparison with Nearby Non-AI Grids

One concern is that the effect might be potentially driven by the national and local government supporting policies in AI as a strategic industry. Throughout 2023, the Chinese government sought to integrate GenAI into broader digital-economy and entrepreneurship policies. Cities such as Beijing, Shanghai, Shenzhen, Hangzhou, and Chengdu introduced subsidy schemes for cloud computing and digital start-ups, promoted “AI + Industry” integration, and established public data-service platforms to support small firms. These initiatives were consistent with China’s long-standing policy orientation toward digital upgrading and innovation-driven growth, but the post-ChatGPT period saw an explicit shift toward the use of AI as an enabling infrastructure for productivity and cost reduction across all sectors rather than as a stand-alone high-tech industry.

To address this and further assess the robustness of our main finding, we implement a geographically matched comparison design to tackle the systematic differences between AI-active and AI-inactive grids. Specifically, for each grid with at least one AI patent filed between 2010 and 2019, we identify the five geographically nearest grids that had no AI patent activity in the same period, allowing for replacement. This matching strategy ensures a more comparable control group based on spatial proximity, holding constant unobservable regional characteristics that may correlate with economic development, industrial composition, or access to infrastructure.

Table B.2 presents the difference-in-differences estimates from this matched sample. Column 1 shows that following the release of ChatGPT, new firm entry in AI-active grids increased by 4.345 firms per grid–quarter relative to their spatially matched counterparts (s.e. = 1.354, $p < 0.01$). Column 2 demonstrates that this effect is again concentrated among small firms, with an estimated increase of 6.584 firms (s.e. = 1.484, $p < 0.01$). In contrast, Column 3 reveals a statistically significant decline in the formation of large firms (coefficient = -2.621 , s.e. = 0.693, $p < 0.01$), consistent with our previous findings. All specifications include city-by-quarter fixed effects to absorb aggregate shocks,

and grid-by-calendar-quarter fixed effects to control for fine-grained temporal and spatial heterogeneity.

6.2.4 Capital Threshold for Firm Size

Our baseline definition of small versus large firms relies on a registered capital threshold of one million RMB. While this benchmark is commonly used in administrative classifications, it may not fully capture meaningful differences in firm scale across industries or regions. To ensure that our findings are not mechanically driven by this cutoff, we conduct a robustness check in which we vary the definition of small business across a range of alternative capital thresholds.

Table B.3 presents the results using thresholds of 2 million, 3 million, and 5 million RMB. In each case, small firms are defined as those below the specified cutoff, and all specifications include city-by-quarter and grid-by-calendar-quarter fixed effects. Across all three panels, we continue to find strong and statistically significant treatment effects in line with our baseline results. Specifically, the estimated coefficients for $Post \times HighAI$ on small firm formation are 7.680 (s.e. = 1.368) at the 2M cutoff, 7.380 (s.e. = 1.384) at 3M, and 7.157 (s.e. = 1.386) at 5M, all significant at the 1% level. Meanwhile, the corresponding effects for large firms remain significantly negative and decline slightly in magnitude as the threshold increases, ranging from -3.096 (s.e. = 0.473) to -2.573 (s.e. = 0.421).

These patterns confirm that our central conclusion—that GenAI disproportionately stimulates the entry of smaller, more agile firms while displacing or discouraging larger entrants—is robust to alternative classifications of firm size. They also suggest that the observed asymmetry is not an artifact of arbitrary thresholds but rather reflects fundamental differences in how the new AI tools facilitate firms of varying scale.

7 Conclusion

This paper provides the first large-scale evidence on how the release of GenAI affects entrepreneurial activity in the real economy. Using comprehensive business registration records for over 12.8 million new firms in China, combined with high-resolution spatial measures of pre-existing AI innovation, we document a substantial increase in new firm formation following the launch of ChatGPT. The effect is concentrated in regions with stronger AI-relevant human capital, driven entirely by small and first-time entrepreneurs. These patterns suggest that GenAI has lowered key barriers to entry, enabling individuals with fewer resources or limited experience to launch new ventures.

Exploring mechanisms, we show that new firms in locations with stronger AI presence become increasingly “lightweight” after the diffusion of GenAI: founders have less prior experience, establish ventures with fewer shareholders, and rely on smaller executive teams. These changes, along with a demonstrated downsizing among repeat founders, suggest that GenAI serves as a substitute for key complementary inputs to entrepreneurship, acting as a “digital cofounder” that enables individuals to launch viable firms with minimal resources. The strongest effects arise in downstream and adoption-oriented sectors—where AI tools can be readily adopted for commercialization and customer-facing activities—rather than in capital- or invention-intensive upstream AI industries.

Taken together, our evidence highlights a new pathway through which general-purpose technologies stimulate economic dynamism—not by expanding frontier innovation, but by democratizing who can start a firm and how ventures are organized at inception. As GenAI tools continue to diffuse and improve, understanding their long-run implications for job creation, entrepreneurial quality, and industrial structure is a central direction for future research. Our findings show that the rise of GenAI represents not only a technological breakthrough, but also an institutional shift in the accessibility, inclusiveness, and scale of entrepreneurship.

References

- Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.
- Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244.
- Aghion, P., Bechtold, S., Cassar, L., and Herz, H. (2018). The causal effects of competition on innovation: Experimental evidence. *The Journal of Law, Economics, and Organization*, 34(2):162–195.
- Aghion, P. and Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, pages 323–351.
- Allen, F., Cai, J., Gu, X., Qian, J., Zhao, L., and Zhu, W. (2024). Centralization or decentralization? the evolution of state-ownership in China. *The Evolution of State-Ownership in China (October 20, 2024)*.
- Allen, F., Cai, J., Gu, X., Qian, J., Zhao, L., and Zhu, W. (2025). Ownership networks and firm growth: what do forty million companies tell us about the chinese economy? *Available at SSRN 3465126*.
- Andreadis, L., Kalotychou, E., Chatzikonstantinou, M., Louca, C., and Makridis, C. A. (2025). Local heterogeneity in artificial intelligence jobs over time and space. In *AEA Papers and Proceedings*, volume 115, pages 29–34. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.
- Antoniades, A., He, C., Liang, Z., and Xu, M. (2025). Learning by AI: the impact of AI hiring on chinese exporters. *Available at SSRN 5361273*.
- Ashraf, M. (2025). Does automation improve financial reporting? evidence from internal controls. *Review of Accounting Studies*, 30(1):436–479.
- Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745.
- Babina, T., Fedyk, A., He, A. X., and Hodson, J. (2023). Firm investments in artificial intelligence technologies and changes in workforce composition. *National Bureau of Economic Research*, 31325.

- Bertomeu, J., Lin, Y., Liu, Y., and Ni, Z. (2025). The impact of generative AI on information processing: Evidence from the ban of ChatGPT in Italy. *Journal of Accounting and Economics*, page 101782.
- Brandt, L., Van Biesebroeck, J., and Zhang, Y. (2012). Creative accounting or creative destruction? firm-level productivity growth in Chinese manufacturing. *Journal of Development Economics*, 97(2):339–351.
- Brynjolfsson, E., Li, D., and Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2):889–942.
- Cai, J. J., Gu, X., Zhao, L., and Zhu, W. (2025). *State Ownership in China: An Equity Network Perspective*, page 156–191. Cambridge University Press.
- Chen, J., Xiao, S., Zhang, P., Luo, K., Lian, D., and Liu, Z. (2024). M3-embedding: Multi-linguality, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 2318–2335.
- Cong, L. W., Liang, T., Zhang, X., and Zhu, W. (2025). Textual factors: A scalable, interpretable, and data-driven approach to analyzing unstructured information. *Management Science*.
- Cong, L. W. and Zhu, W. (2024). Divergent LLM adoption and heterogeneous convergence paths in research writing. *Available at SSRN 5170513*.
- Croom, J. (2025). Interactivity and illusions of ability: How using generative AI affects investor judgments. *Journal of Accounting Research*.
- Czarnitzki, D., Fernández, G. P., and Rammer, C. (2023). Artificial intelligence and firm-level productivity. *Journal of Economic Behavior & Organization*, 211:188–205.
- da Silva Marioni, L., Rincon-Aznar, A., and Venturini, F. (2024). Productivity performance, distance to frontier and AI innovation: Firm-level evidence from Europe. *Journal of Economic Behavior & Organization*, 228:106762.
- Decker, R., Haltiwanger, J., Jarmin, R., and Miranda, J. (2014). The role of entrepreneurship in US job creation and economic dynamism. *Journal of Economic Perspectives*, 28(3):3–24.
- Eisfeldt, A. L., Schubert, G., Zhang, M. B., and Taska, B. (2025). Generative AI and firm values. *Journal of Finance*, forthcoming.

- Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702):1306–1308.
- Evans, D. S. and Jovanovic, B. (1989). An estimated model of entrepreneurial choice under liquidity constraints. *Journal of Political Economy*, 97(4):808–827.
- Fang, H., Gu, X., Yan, H., and Zhu, W. (2025a). AI inventions and AI talent migrations: US and China.
- Fang, L., Yuan, Z., Zhang, K., Donati, D., and Sarvary, M. (2025b). Generative AI and firm productivity: Field experiments in online retail. *arXiv preprint arXiv:2510.12049*.
- Fedyk, A., Hodson, J., Khimich, N., and Fedyk, T. (2022). Is artificial intelligence improving the audit process? *Review of Accounting Studies*, 27(3):938–985.
- Foster, L., Haltiwanger, J. C., and Krizan, C. J. (2001). Aggregate productivity growth: Lessons from microeconomic evidence. In *New developments in productivity analysis*, pages 303–372. University of Chicago Press.
- Gofman, M. and Jin, Z. (2024). Artificial intelligence, education, and entrepreneurship. *The Journal of Finance*, 79(1):631–667.
- Haltiwanger, J., Jarmin, R. S., and Miranda, J. (2013). Who creates jobs? small versus large versus young. *Review of Economics and Statistics*, 95(2):347–361.
- Hampole, M., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2025). Artificial intelligence and the labor market. *National Bureau of Economic Research*.
- Hopenhayn, H. A. (1992). Entry, exit, and firm dynamics in long run equilibrium. *Econometrica: Journal of the Econometric Society*, pages 1127–1150.
- Hurst, A., Lerer, A., Goucher, A. P., Perelman, A., Ramesh, A., Clark, A., Ostrow, A., Welihinda, A., Hayes, A., Radford, A., et al. (2024). Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Jovanovic, B. (1982). Selection and the evolution of industry. *Econometrica*, pages 649–670.
- Kanazawa, K., Kawaguchi, D., Shigeoka, H., and Watanabe, Y. (2025). AI, skill, and productivity: The case of taxi drivers. *Management Science*.
- Kerr, W. R. and Nanda, R. (2015). Financing innovation. *Annual Review of Financial Economics*, 7(1):445–462.

- Korinek, A. (2023). Generative AI for economic research: Use cases and implications for economists. *Journal of Economic Literature*, 61(4):1281–1317.
- Lichtinger, G. and Hosseini Maasoum, S. M. (2025). Generative AI as seniority-biased technological change: Evidence from US resume and job posting data. *Available at SSRN*.
- McInnes, L., Healy, J., Astels, S., et al. (2017). HDBSCAN: Hierarchical density based clustering. *J. Open Source Softw.*, 2(11):205.
- McInnes, L., Healy, J., and Melville, J. (2018). UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.
- Schumpeter, J. A. (1943). *Capitalism, socialism and democracy*. New York: Harper and Brothers.
- Shi, Y., Townsend, R. M., and Zhu, W. (2020). Tiered intermediation in equity-holding networks. *Available at SSRN 3722173*.
- Wang, X. and Wu, L. (2025). Artificial intelligence, lean startup method, and product innovations. *Management Science*.
- Wu, L., Lou, B., and Hitt, L. M. (2025). Innovation strategy after IPO: How AI analytics spurs innovation after IPO. *Management Science*, 71(3):2360–2389.
- Xiao, S., Liu, Z., Zhang, P., Muennighoff, N., Lian, D., and Nie, J.-Y. (2024). C-pack: Packed resources for general chinese embeddings. In *Proceedings of the 47th international ACM SIGIR conference on research and development in information retrieval*, pages 641–649.
- Xue, J., Zhang, Q., and Zhu, W. (2025). Generative AI for analysts. *SSRN Electronic Journal*:10.2139/ssrn.5821142.

Figures & Tables

Figure 1: Quarterly Trends in Average New Firm Formation by Grid

This figure shows the quarterly trend in new firm formation. The x-axis indicates the quarter, with 2022Q4 marking the release of ChatGPT. The blue solid line represents the average number of new small businesses across all grids in each quarter, while the red dashed line represents the corresponding average number of new large businesses.




Figure 2: Histogram Distribution of AI Patent

This figure shows the distribution of the number of AI patents across grids that have at least one AI patent.




Figure 3: Geographical Distribution of AI Patents and Inventors






Figure 4: Dynamic Effect of the Release of ChatGPT on Firm Creation


This figure illustrates the dynamic evolution of new firm formation in the quarters surrounding the release of ChatGPT. The X-axis represents the quarter relative to the release of ChatGPT, where 0 corresponds to 2022Q4. The empirical specification is based on regression Equation (3). Panels A–C display results for the following dependent variables: (A) the number of new firms, (B) the number of small firms, and (C) the number of large firms. The Y-axis plots the estimated coefficients, β_k .



(a) Panel A: Entry of New Firms



(b) Panel B: Entry of Small Firms



(c) Panel C: Entry of Large Firms

Figure 5: Random Assignment of AI Exposure Labels

This figure shows the distribution of estimated coefficients for the interaction term $Post \times HighAI$ across 100 simulations, in which 10,183 grids are randomly assigned to $HighAI = 1$. The mean coefficient is 0.064, with a standard deviation of 0.166.




Table 1: Summary Statistics

This table shows the summary statistics of the variables we are interested in at the grid-quarter level. For the definition of the variables, see the Appendix [A.1](#).

Variable	N	Mean	sd	Min	Median	Max
Num new firms	2,658,496	4.822	39.626	0	0	9,678
Num small firms	2,658,496	2.082	20.536	0	0	9,672
Num large firms	2,658,496	2.469	22.179	0	0	4,262
Num high upstream	2,658,496	0.969	11.010	0	0	1,772
Num low upstream	2,658,496	3.850	30.386	0	0	9,677
Num high downstream	2,658,496	2.494	24.761	0	0	3,949
Num low downstream	2,658,496	2.325	18.700	0	0	9,652
Num high entrep	2,658,496	3.617	34.748	0	0	8,504
Num low entrep	2,658,496	1.202	6.997	0	0	1,856
Pct serial entrep	1,041,042	26.743	33.560	0	12.5	100
Pct serial entrep small	1,041,042	16.346	30.013	0	0	100
Pct serial entrep large	1,041,042	19.936	32.136	0	0	100
Tot shareholders	1,002,793	1.502	0.852	1	1.333	139
Tot shareholders small	670,018	1.389	0.693	1	1.103	49
Tot shareholders large	743,846	1.624	0.991	1	1.5	139
Pct indiv shareholders	1,002,793	93.205	18.650	0	100	100
Pct indiv shareholders small	670,018	96.057	14.851	0	100	100
Pct indiv shareholders large	743,846	90.705	21.807	0	100	100
Exec team size	934,016	2.031	0.636	1	2	22
Exec team size small	595,785	1.965	0.608	1	2	19
Exec team size large	714,361	2.095	0.662	1	2	32
AIpat	2,658,496	2.051	58.812	0	0	11,279
HighAI	2,658,496	0.061	0.240	0	0	1

Table 2: Number of New Firms

This table reports DiD estimates from Equations (1) and (2), examining the impact of GenAI on new firm formation. The variable *Post* equals one for quarters after 2022Q4, and *HighAI* is an indicator for grid cells with at least one AI-related patent filed between 2010 and 2019. Columns (1)–(3) use the following dependent variables: (1) number of new firms; (2) number of small firms; and (3) number of large firms. All regressions include city-by-quarter fixed effects and grid-by-calendar-quarter fixed effects to absorb seasonal patterns. Standard errors are clustered at the city level. Variable definitions are provided in Section A.1. Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1) Num new firms	(2) Num small firms	(3) Num large firms
Post × HighAI	5.038*** (1.395)	7.704*** (1.222)	-3.120*** (0.671)
Constant	4.668*** (0.0427)	1.847*** (0.0374)	2.564*** (0.0205)
Observations	2,658,304	2,658,304	2,658,304
R-squared	0.810	0.625	0.790
City × Quarter FE	Yes	Yes	Yes
Grid × Cal QTR FE	Yes	Yes	Yes

Table 3: Robustness: Number of New Firms

This table reports DiD estimates from Equation (1), examining the impact of GenAI on new firm formation. The variable *Post* equals one for quarters after 2022Q4, and *HighAI* is an indicator for grid cells with at least one AI-related patent filed between 2010 and 2019. Columns (1)–(6) use the following dependent variables, respectively: (1) and (4) the number of new firms; (2) and (5) the number of small firms; and (3) and (6) the number of large firms. To mitigate the influence of extreme values, we apply a 0.1% winsorization to the dependent variable. Columns (1)–(3) are winsorized within quarter, while Columns (4)–(6) are winsorized within the full sample. All regressions include city-by-quarter fixed effects and grid-by-calendar-quarter fixed effects to absorb seasonal patterns. Standard errors are clustered at the city level. Variable definitions are provided in Section A.1. *Notes*: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	Winsorize by quarter			Winsorize across all data		
	(1) Num new firms	(2) Num small firms	(3) Num large firms	(4) Num new firms	(5) Num small firms	(6) Num large firms
Post × HighAI	3.692*** (0.826)	6.170*** (0.800)	-2.928*** (0.332)	2.518*** (0.667)	4.911*** (0.528)	-2.922*** (0.336)
Constant	4.275*** (0.0253)	1.677*** (0.0245)	2.321*** (0.0102)	4.304*** (0.0204)	1.694*** (0.0162)	2.307*** (0.0103)
Observations	2,658,304	2,658,304	2,658,304	2,658,304	2,658,304	2,658,304
R-squared	0.876	0.771	0.856	0.885	0.811	0.875
City × Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Grid × Cal QTR FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 4: Industry-Specific DID Estimates

This table presents the estimated coefficients of the interaction term from Equation (1), assessing the impact of GenAI on new firm formation across 96 industries. Industries are sorted in descending order by the estimated coefficient. Panel A presents the top 15 industries, and Panel B presents the bottom 15 industries.

Panel A: Top 15

Industry Name	Count	Coef	SE	p-value
Retail Industry	2,019,644	1.6256	0.5447	0.0030
Business Services	1,135,654	0.9771	0.2921	0.0009
Technology Promotion and Application Services	1,301,263	0.8712	0.2316	0.0002
Wholesale Industry	2,061,046	0.4547	0.2466	0.0661
Entertainment Industry	137,133	0.2223	0.0534	0.0000
Catering Industry	199,082	0.1970	0.0389	0.0000
Culture and Arts Industry	289,160	0.1806	0.0579	0.0020
Resident Services Industry	169,229	0.1464	0.0294	0.0000
Internet and Related Services	118,450	0.0963	0.0277	0.0006
Software and IT Services	503,900	0.0556	0.1064	0.6016
Broadcasting, Television, Film, and Audio Production	48,885	0.0443	0.0155	0.0045
Electricity, Heat, Gas, and Water Supply	42,173	0.0400	0.0055	0.0000
Leasing Industry	152,152	0.0388	0.0215	0.0720
Agriculture	233,704	0.0360	0.0172	0.0370
Chemical Raw Materials and Chemical Products Manufacturing	28,633	0.0340	0.0149	0.0232

Panel B: Bottom 15

Industry Name	Firm Count	Coefficient	SE	p-value
Building Construction	217,478	-0.1588	0.1063	0.1363
Civil Engineering Construction	256,719	-0.0858	0.0372	0.0218
Decoration, Renovation, and Other Construction	235,693	-0.0565	0.0148	0.0002
Construction Installation	106,877	-0.0194	0.0172	0.2591
Real Estate Industry	243,165	-0.0136	0.0166	0.4108
Other Manufacturing	25,799	-0.0096	0.0067	0.1492
Education	82,810	-0.0089	0.0119	0.4569
Metal Products Manufacturing	67,641	-0.0079	0.0038	0.0381
General Equipment Manufacturing	55,450	-0.0066	0.0036	0.0665
Computer, Communication, and Electronics Manufacturing	23,854	-0.0054	0.0095	0.5653
Ecological and Environmental Governance	17,471	-0.0053	0.0021	0.0102
Nonmetallic Mineral Products	48,849	-0.0041	0.0028	0.1460
Monetary and Financial Services	8,630	-0.0032	0.0013	0.0169
Specialized Equipment Manufacturing	42,793	-0.0029	0.0054	0.5898
Instrumentation Manufacturing	9,486	-0.0020	0.0014	0.1536

Table 5: Industry Heterogeneity: Relevance to AI

This table reports estimates of changes in the number of new firms across different industries following the release of ChatGPT. *Post* is an indicator equal to one for quarters after 2022Q4, and *HighAI* is a dummy equal to one for grids with at least one AI patent filed between 2010 and 2019. Columns (1)–(6) use the following dependent variables, respectively: (1) number of new firms in industries more likely to be upstream of AI; (2) number of new firms in industries less likely to be upstream of AI; (3) number of new firms in industries more likely to be downstream of AI; (4) number of new firms in industries less likely to be downstream of AI; (5) number of new firms in industries with high entrepreneurship scores; and (6) number of new firms in industries with low entrepreneurship scores. All specifications include city-by-quarter fixed effects and grid-by-calendar-quarter fixed effects. Standard errors are clustered at the city level. Variable definitions are provided in Section A.1. *Notes*: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1) Num high upstream	(2) Num low upstream	(3) Num high downstream	(4) Num low downstream	(5) Num high entrep	(6) Num low entrep
Post × HighAI	0.846** (0.388)	4.199*** (1.225)	3.323*** (0.789)	1.722** (0.753)	5.892*** (1.302)	-0.848*** (0.300)
Constant	0.944*** (0.0119)	3.721*** (0.0375)	2.393*** (0.0242)	2.272*** (0.0231)	3.437*** (0.0399)	1.228*** (0.00918)
Observations	2,658,304	2,658,304	2,658,304	2,658,304	2,658,304	2,658,304
R-squared	0.834	0.770	0.857	0.638	0.797	0.746
City × Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Grid × Cal QTR FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 6: ChatGPT and Serial Entrepreneurship

This table reports estimates of changes in the ratio of serial entrepreneurs following the release of ChatGPT. *Post* is an indicator equal to one for quarters after 2022Q4, and *HighAI* is a dummy equal to one for grids with at least one AI patent filed between 2010 and 2019. The ratio of serial entrepreneurs is defined as the percentage of new firm founders (legal representatives) who had established at least one other firm within the three years preceding the current firm's founding, excluding the month of establishment. Columns (1)–(3) use the following dependent variables, respectively: (1) the ratio of serial entrepreneurs in the full sample; (2) the ratio among small firms; and (3) the ratio among large firms. All specifications include city-by-quarter fixed effects and grid-by-calendar-quarter fixed effects. Standard errors are clustered at the city level. Variable definitions are provided in Section A.1. *Notes*: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1) Pct serial entrep	(2) Pct serial entrep small	(3) Pct serial entrep large
Post × HighAI	-0.405** (0.201)	-2.456*** (0.266)	0.648** (0.294)
Constant	27.47*** (0.0116)	17.82*** (0.0154)	21.25*** (0.0170)
Observations	863,869	863,869	863,869
R-squared	0.433	0.420	0.425
City × Quarter FE	Yes	Yes	Yes
Grid × Cal QTR FE	Yes	Yes	Yes

Table 7: ChatGPT and Shareholder Numbers

This table reports estimates of changes in the number of shareholders following the release of ChatGPT. *Post* is an indicator equal to one for quarters after 2022Q4, and *HighAI* is a dummy equal to one for grids with at least one AI patent filed between 2010 and 2019. Panel A reports results for the total number of shareholders (including both individual and corporate shareholders), while Panel B focuses on the share of individual shareholders. Column (1) presents results for the full sample, Column (2) for small firms, and Column (3) for large firms. All specifications include city-by-quarter fixed effects and grid-by-calendar-quarter fixed effects. Standard errors are clustered at the city level. Variable definitions are provided in Section A.1. Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Panel A: Total Shareholders			
	(1) Tot shareholders	(2) Tot shareholders small	(3) Tot shareholders large
Post × HighAI	-0.0211*** (0.00747)	-0.0134* (0.00755)	-0.0173** (0.00781)
Constant	1.513*** (0.000444)	1.412*** (0.000646)	1.639*** (0.000598)
Observations	827,286	512,172	585,030
R-squared	0.423	0.434	0.427
City × Quarter FE	Yes	Yes	Yes
Grid × Cal QTR FE	Yes	Yes	Yes

Panel B: Percentage of Individual Shareholders			
	(1) Pct indiv shareholders	(2) Pct indiv shareholders small	(3) Pct indiv shareholders large
Post × HighAI	0.0578 (0.161)	-0.185 (0.132)	-0.0774 (0.169)
Constant	93.29*** (0.00955)	96.22*** (0.0113)	90.55*** (0.0130)
Observations	827,286	512,172	585,030
R-squared	0.456	0.434	0.472
City × Quarter FE	Yes	Yes	Yes
Grid × Cal QTR FE	Yes	Yes	Yes

Table 8: ChatGPT and Executive Member

This table reports estimates of changes in the number of executive members following the release of ChatGPT. *Post* is an indicator equal to one for quarters after 2022Q4, and *HighAI* is a dummy equal to one for grids with at least one AI patent filed between 2010 and 2019. Columns (1)–(3) use the following dependent variables, respectively: (1) number of executive members; (2) number of executive members in small firms; and (3) number of executive members in large firms. All specifications include city-by-quarter fixed effects and grid-by-calendar-quarter fixed effects. Standard errors are clustered at the city level. Variable definitions are provided in Section A.1. Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1) Exec team size	(2) Exec team size small	(3) Exec team size large
Post × HighAI	-0.0161*** (0.00392)	-0.0187*** (0.00401)	-0.00373 (0.00441)
Constant	2.032*** (0.000249)	1.976*** (0.000382)	2.091*** (0.000349)
Observations	765,404	452,023	560,617
R-squared	0.568	0.611	0.546
City × Quarter FE	Yes	Yes	Yes
Grid × Cal QTR FE	Yes	Yes	Yes

Table 9: Serial Entrepreneurship and Substitution Towards Smaller Firms

This table reports estimates of changes in the relative size of new firms established by serial entrepreneurs (legal representatives who had founded at least one other firm within the three years preceding the current firm's establishment, excluding the month of founding). *Post* is an indicator equal to one for quarters after 2022Q4, and *HighAI* is a dummy equal to one for grids with at least one AI patent filed between 2010 and 2019. Panel A reports the grid-quarter level percentage of new firms established by a serial entrepreneur whose registered capital exceeds that of the entrepreneur's previous firm. Panel B reports the grid-quarter average ratio of registered capital in the new firm established by a serial entrepreneur to that in the previous firm. To mitigate the impact of outliers, the top 1% and bottom 1% of values are trimmed. Columns (1) presents the result for all firms, Column (2) for small firms, and Column (3) for large firms. All specifications include city-by-quarter fixed effects and grid-by-calendar-quarter fixed effects. Standard errors are clustered at the city level. Variable definitions are provided in Section A.1. Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Panel A: Percentage of New Firms Established by Serial Entrepreneurs with Larger Registered Capital

	(1) All firms	(2) Small firms	(3) Large firms
Post × HighAI	-0.629* (0.320)	-5.416*** (0.483)	-0.917*** (0.241)
Constant	36.29*** (0.0348)	40.95*** (0.0525)	17.73*** (0.0262)
Observations	375,548	375,548	375,548
R-squared	0.411	0.438	0.422
City × Quarter FE	Yes	Yes	Yes
Grid × Cal QTR FE	Yes	Yes	Yes

Panel B: Average Capital Ratio of New to Prior Firms for Serial Entrepreneurs

	Ratio of New-to-Previous Firm Registered Capital		
	(1) All firms	(2) Small firms	(3) Large firms
Post × HighAI	-1.983*** (0.560)	-7.100*** (1.148)	-0.180*** (0.0230)
Constant	9.601*** (0.0612)	16.64*** (0.124)	1.200*** (0.00249)
Observations	366,604	370,610	370,861
R-squared	0.375	0.359	0.405
City × Quarter FE	Yes	Yes	Yes
Grid × Cal QTR FE	Yes	Yes	Yes

Table 10: Placebo Tests Using NonAI Patents

This table reports estimates of changes in the number of new firms following the release of ChatGPT. *Post* is an indicator equal to one for quarters after 2022Q4. In this table, we focus on grids with non-AI patents and construct *High nonAI* based on the residuals from regressing the log number of non-AI patents on the log number of AI patents. *High nonAI* equals one for grids whose residual exceeds the 75th percentile. Columns (1)–(3) use the following dependent variables: (1) number of new firms; (2) number of small firms; and (3) number of large firms. All regressions include city-by-quarter fixed effects and grid-by-calendar-quarter fixed effects to absorb seasonal patterns. Standard errors are clustered at the city level. Variable definitions are provided in Section A.1. *Notes*: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1) Num new firms	(2) Num small firms	(3) Num large firms
Post × High nonAI	0.159 (0.478)	0.558 (0.485)	-0.427* (0.249)
Constant	16.69*** (0.0589)	7.063*** (0.0599)	8.721*** (0.0308)
Observations	546,768	546,768	546,768
R-squared	0.818	0.645	0.791
City × Quarter FE	Yes	Yes	Yes
Grid × Cal QTR FE	Yes	Yes	Yes

Table 11: Residualizing AI Exposure to Remove Entrepreneurial Correlation

This table reports estimates of changes in the number of new firms following the release of ChatGPT. *Post* is an indicator equal to one for quarters after 2022Q4. *HighResid* is constructed from the residuals of a regression of the log number of prior firms (plus one) on the log number of AI patents (plus one). *HighResid* equals one for grid cells with residuals above the median. Columns (1)–(3) use the following dependent variables: (1) number of new firms; (2) number of small firms; and (3) number of large firms. All regressions include city-by-quarter fixed effects and grid-by-calendar-quarter fixed effects to absorb seasonal patterns. Standard errors are clustered at the city level. Variable definitions are provided in Section A.1. Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1) Num new firms	(2) Num small firms	(3) Num large firms
Post × HighResid	0.782*** (0.173)	1.216*** (0.143)	-0.481*** (0.0601)
Constant	4.629*** (0.0427)	1.782*** (0.0352)	2.587*** (0.0149)
Observations	2,658,304	2,658,304	2,658,304
R-squared	0.810	0.623	0.790
City × Quarter FE	Yes	Yes	Yes
Grid × Cal QTR FE	Yes	Yes	Yes

Table 12: Excluding First-Tier Provinces

This table reports estimates of changes in the number of new firms following the release of ChatGPT. *Post* is an indicator equal to one for quarters after 2022Q4, and *HighAI* is a dummy equal to one for grids with at least one AI patent filed between 2010 and 2019. To mitigate the influence of highly active regions, we exclude grids located in Beijing, Shanghai, and Guangdong. Columns (1)–(3) use the following dependent variables: (1) number of new firms; (2) number of small firms; and (3) number of large firms. All regressions include city-by-quarter fixed effects and grid-by-calendar-quarter fixed effects to absorb seasonal patterns. Standard errors are clustered at the city level. Variable definitions are provided in Section A.1. Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1) Num new firms	(2) Num small firms	(3) Num large firms
Post × HighAI	3.979** (1.537)	6.288*** (1.165)	-2.684*** (0.761)
Constant	3.512*** (0.0383)	1.352*** (0.0291)	1.942*** (0.0190)
Observations	2,477,248	2,477,248	2,477,248
R-squared	0.758	0.510	0.802
City × Quarter FE	Yes	Yes	Yes
Grid × Cal QTR FE	Yes	Yes	Yes

A Appendix A

A.1 Variables and Definitions

Variable	Definition
Num new firms	Number of new firms in grid g during quarter t .
Num small firms	Number of small firms (registered capital < 1 million RMB) in grid g during quarter t .
Num large firms	Number of large firms (registered capital ≥ 1 million RMB) in grid g during quarter t .
Num high upstream	Number of new firms with AI upstream score > 0 in grid g during quarter t .
Num low upstream	Number of new firms with AI upstream score ≤ 0 in grid g during quarter t .
Num high downstream	Number of new firms with AI downstream score > 0 in grid g during quarter t .
Num low downstream	Number of new firms with AI downstream score ≤ 0 in grid g during quarter t .
Num high entrep	Number of new firms with Entrepreneurship Helpfulness score > 0 in grid g during quarter t .
Num low entrep	Number of new firms with Entrepreneurship Helpfulness score ≤ 0 in grid g during quarter t .
Pct serial entrep	Average percentage of serial entrepreneurs (legal representatives who had founded at least one other firm within the three years preceding the current firm's establishment, excluding the month of founding) in grid g during quarter t .
Pct serial entrep small	Average percentage of serial entrepreneurs in small businesses in grid g during quarter t .
Pct serial entrep large	Average percentage of serial entrepreneurs in large businesses in grid g during quarter t .
Tot shareholders	Average number of shareholders of new firms in grid g during quarter t .
Tot shareholders small	Average number of shareholders of small businesses in grid g during quarter t .
Tot shareholders large	Average number of shareholders of large businesses in grid g during quarter t .
Pct indiv shareholders	Average number of individual shareholders divided by total shareholders $\times 100$ for new firms in grid g during quarter t .
Pct indiv shareholders small	Average number of individual shareholders divided by total shareholders $\times 100$ for small businesses in grid g during quarter t .
Pct indiv shareholders large	Average number of individual shareholders divided by total shareholders $\times 100$ for large businesses in grid g during quarter t .
Exec team size	Average number of executive members of new firms in grid g during quarter t .
Exec team size small	Average number of executive members of small businesses in grid g during quarter t .
Exec team size large	Average number of executive members of large businesses in grid g during quarter t .
AIpat	Number of AI-related patents filed between 2010 and 2019 in grid g
HighAI	Binary indicator equal to 1 if grid g has at least one AI-related patent

A.2 AI relevance Score

This section outlines the construction and validation of the firm-level AI relevance scores that quantify (i) how tightly a firm’s business activities connect to AI technologies and applications (*upstream* and *downstream*) and (ii) the extent to which AI can enable *entrepreneurship*. The construction proceeds in two steps. First, we represent each firm’s business scope as a combination of latent *textual factors*, i.e., topics extracted from firm’s business description that capture the semantic structure and economic meaning of its activities (see Appendix A.2.1). In the second step, we employ the state-of-the-art large language model GPT-4o (Hurst et al., 2024) to assign topic-specific AI relevance scores. A firm’s AI relevance scores are then computed as a weighted average of these topic scores, with weights given by the firm’s topic loadings (Appendix A.2.2).

We check the validity of the AI relevance scores in Appendix A.2.3. We randomly sample a subset of firms, review their business websites and publicly available information, and compare these qualitative assessments with our constructed AI-relevance scores. Importantly, we do not reuse the business-scope text in construction, thereby avoiding circularity. The clear correspondence between the two provides supporting evidence for validity of our constructed AI relevance measures.

A.2.1 Business Topics

Let N denote the set of all firms that entered during the sample period, and let N_t represent the set of entrants in quarter t . The full sample of entrants can therefore be expressed as the union across all quarters:

$$N = \bigcup_t N_t.$$

We denote $|N|$ as the total number of new firms in the sample and $|N_t|$ as the number of entrants in quarter t .

Each firm in our dataset is associated with a paragraph describing its business scope, which we treat as a short textual document. Let V denote the set of all unique vocabulary terms appearing across these documents after standard text preprocessing (e.g., removal of stop words), and let $|V|$ be its cardinality. To extract structured information from this corpus, we adopt the topic modeling framework developed by Cong et al. (2025), which represents each document as a mixture of latent topics and each topic as a distribution over vocabulary terms. This approach enhances the interpretability of traditional topic models while maintaining scalability for large corpora. The construction proceeds in three steps:

- **Word Embedding.** We encode the semantic meaning of each term using the pre-trained Chinese embedding model `bge-base-zh-v1.5`, developed by the Beijing Academy of Artificial Intelligence (Xiao et al., 2024). This model maps each word into a dense vector space that captures both syntactic and semantic relationships.
- **Term Clustering.** We then apply Locality-Sensitive Hashing (LSH) to cluster these word vectors into K mutually exclusive groups, each containing at most 50 terms. Every word is assigned to exactly one cluster. This clustering step groups semantically similar words—based on their embedding proximity—into coherent economic concepts, improving the interpretability of the resulting topics.
- **Topic Extraction.** For each cluster $k \in \{1, \dots, K\}$, let $V_k \subset V$ denote the subset of vocabulary terms assigned to cluster k . We construct a document-term matrix using only the terms in V_k and apply Latent Semantic Analysis (LSA) via singular value decomposition (SVD). This yields topic loadings for each document, denoted $\theta_{i,k}$ for firm i , and word loadings for each topic, denoted f_k .

Before implementing these steps, we clean all business scope descriptions by removing boilerplate disclaimers and tokenize the text using Jieba. We then cluster the unique terms, producing an initial set of 8,460 clusters. To enhance computational efficiency, we estimate topics using a 5% random sample of firms established between 2015 and 2024, comprising 3,019,725 documents. Clusters with fewer than five terms are removed, leaving 4,967 economically interpretable topics. Finally, we compute topic loadings for all firms in the universe following Cong et al. (2025). For each firm i , we calculate a normalized word frequency vector and multiply it by the topic-specific term distribution, yielding the topic loading $\theta_{i,k}$, which represents the firm’s loading in topic k .

A.2.2 AI Relevance via Large Language Models

This subsection describes how we generate a structured and semantically consistent evaluation of each topic’s relationship to artificial intelligence using large language models (LLMs). Specifically, we employ GPT-4o (Hurst et al., 2024) to assess each topic along multiple dimensions of AI relevance, based on the topic’s most representative keywords and their associated weights derived from the topic model.

For each topic, we construct a concise text prompt containing a ranked list of top keywords and their corresponding importance weights. We then instruct the model to evaluate four conceptual dimensions, each scored on a scale from -100 to 100 :

1. **AI Upstream Score:** the extent to which the topic concerns upstream AI activities—such as model development, data infrastructure, computing resources, or foundational technologies.
2. **AI Downstream Score:** the extent to which the topic captures downstream AI activities—such as application development, product integration, or end-user services.
3. **Entrepreneurship Helpfulness Score:** the degree to which the topic is relevant or enabling for entrepreneurial or solo creative activities (e.g., content creation, personal branding, low-cost experimentation, or digital marketing), scored as:
 - -100 = obstructive or irrelevant to entrepreneurship,
 - 0 = neutral or unrelated,
 - 100 = highly enabling for entrepreneurial or small-scale creative work.

For each topic, the model also provides a brief justification and, when appropriate, examples of firms or industries where the topic is particularly relevant to AI-driven entrepreneurship. These responses are parsed and compiled into a structured dataset containing the three topic-level scores—*AI Upstream*, *AI Downstream*, and *Entrepreneurship Helpfulness*. To ensure robustness, the evaluation can be repeated multiple times to assess the stability of the scores.

Finally, we aggregate the topic-level evaluations into firm-level measures. For firm i , the AI-relevance score is computed as a weighted average of the topic-specific scores:

$$\text{AIScore}_i = \frac{\sum_k \theta_{i,k} s_k}{\sum_k \theta_{i,k}},$$

where $\theta_{i,k}$ denotes the topic loading for firm i on topic k , and s_k is the score assigned to topic k by GPT-4o. This approach yields a continuous, interpretable measure of each firm’s exposure to AI-related activities and its potential entrepreneurial complementarity with GenAI technologies.

A.2.3 Validation of the AI Relevance Score

This subsection validates the AI relevance score constructed above. Since LLM pipelines can suffer from hallucination and bias, we design a two-stage, human-in-the-loop procedure that leverages firms’ webpages with deterministic preprocessing, conservative fallbacks, and manual audits. Importantly, we do *not* reuse the “business scope” descriptions that are used to construct the AI relevance score so as to provide a clean validation. We

begin from a random sample of 27,239 firms and validate the scoring by comparing the keywords between the high- and low-score groups.

A Two-stage Procedure. We design a two-stage procedure for the validation: in Stage 1, we crawl firms’ webpages and extract raw keywords; and in Stage 2, we construct high-level keywords.

1. Stage 1: Websites collection and keywords extraction

We first obtain each firm’s official website and crawl its site for the initial random samples of 27,239 firms. During preprocessing, we remove malformed entries and uninformative URLs, such as using a generic e-commerce domain as the company homepage (e.g., `www.taobao.com`) while allowing store-specific URLs on such platforms. We then perform domain normalization and correction, retaining 24,528 firms.

We then identify and fetch the page most suitable for a business description by combining `sitemap.xml` cues, common “About/Business Introduction” paths (such as `/about`), homepage link analysis, and a weighted scoring scheme, with the homepage as a fallback. If the firm’s site is itself a path, as is common for e-commerce storefronts, we use it directly. We tried both simple HTTP requests and a headless browser (Playwright), both of which perform similarly. This step yields 8,729 useful websites.

To extract raw keywords, from each selected page, we convert HTML to clean text using `trafilatura` and then transform it into structured fields via batched LLM prompts that summarize the firm’s business description, primary industry, and finally extract raw keywords. We obtain 8,871 raw keywords from all the firms.

2. Stage 2: Human-audited semantic clustering and keyword taxonomy

In Stage 2, we consolidate the 8,871 raw keywords to construct a concise, high-level taxonomy, yielding 129 high-level keywords that are granular yet broad enough.

Specifically, we embed keywords with a multi-lingual LLM model `bge-m3` (Chen et al., 2024), reduce dimensionality using UMAP (McInnes et al., 2018), and form clusters with HDBSCAN (McInnes et al., 2017). To summarize the cluster, we involve human in the loop: we adopt a semi-automated, semi-manual workflow where AI proposes concise, cluster-level descriptors and human reviewers audit, refine, and, where necessary, conduct secondary clustering and re-summarization for mixed or oversized clusters. The result is a validated set of 129 high-level keywords that are granular yet broadly interpretable. For example, “data science,” “Internet-of-Things,” and “real estate.”

Results. To assess whether the AI relevance score is sensible, we compute for each keyword its relative frequency in the high- versus low-score groups; equivalently, we examine the ratio of keyword frequency in the high group to that in the low group. We highlight the top 20 with the largest ratios as the most differentiating terms by the three AI relevance scores respectively.

1. AI upstream score.

Figure A.1 shows the relative frequency of keywords in firms' website descriptions for the high- versus low-AI-upstream groups. In the high-upstream panel (Figure A.1a), the top keywords concentrate on the upstream industries of AI: data industry, communications networks, and the Internet of Things, followed by science- and hardware-intensive enablers such as biotechnology/drug R&D, intelligent manufacturing, consumer/electronic components, aviation/medical diagnostics, new energy, and semiconductors. These categories map naturally to data generation and curation, connectivity, sensing, and compute-adjacent hardware, which are precisely the upstream layers of the AI stack.

By contrast, the low-upstream panel (Figure A.1b) is dominated by traditional, service-oriented, or non-digital verticals: food industry, logistics, cultural industry, landscaping, customer service, leasing, interior design, apparel, and lodging, which are largely irrelevant to AI inputs or plausibly downstream users of technology rather than suppliers of core AI inputs.

Taken together, these keywords distributions based on the their websites contrast the upstream enablers in the high-scored group with the conventional brick-and-mortar activities in the low-scored group, providing evidence that the AI upstream score captures the intended construct.

2. AI downstream score.

Figure A.2 plots the relative frequency of keywords in firms' website descriptions for the high- versus low-AI downstream groups. In the high-downstream panel (Figure A.2a), top terms concentrate on end-user applications and distribution channels, including online games, office collaboration, media/entertainment, dating/relationships, image processing, information services, advertising/marketing—and on consumer-facing electronics. These are demand-side domains where AI can be directly embedded in products and user experiences.

By contrast, the low-downstream panel (Figure A.2b) is dominated by traditional, non-digital sectors, such as leasing, construction, food, interior design/renovation, cultural industry, landscaping, apparel, lodging, which have weaker direct ties to AI products. These patterns support the validity of the downstream score.

3. Entrepreneurship score.

Figure A.3 reports the analogous analysis for the entrepreneurship score. In the high-entrepreneur panel (Figure A.3a), top keywords cover internet markets and tools: computation power, consumer electronics, online games, media/entertainment, marketing, website/information services, translation services, artificial intelligence, “no-code” platforms (services that allows users to create applications, websites, and automations through visual interfaces and drag-and-drop tools instead of writing traditional programming code), information security, and general internet/O2O. These categories characterize rapid software iteration, low distribution costs, and scalable monetization, which are conducive to startup.

The low-entrepreneur panel (Figure A.3b) instead features asset-intensive industries with long capital cycles: metals, civil engineering, construction, chemical manufacturing, recycling, materials, logistics, supply chain, import/export, warehousing, project quoting, electric power, real estate, shipping, and mining, which are less conducive to early-stage, software-centric entrepreneurship. This contrast shows that the entrepreneur score also captures the intended construct.






Figure A.1: Relative frequency in the high- versus low- AI upstream groups.




(a) High AI downstream.




(b) Low AI downstream.

Figure A.2: Relative frequency in the high- versus low- AI downstream groups.



(a) High AI-entrepreneurship.




(b) Low AI-entrepreneurship.

Figure A.3: Relative frequency in the high- versus low- AI entrepreneurship groups.


B Appendix B

Figure B.1: Geographical Distribution of AI Patents in Major Cities


This figure plots the spatial distribution of AI patents across three representative cities—Beijing, Shanghai, and Shenzhen—using H3 grid cells at 2km resolution. Each point represents a hexagonal cell, with shading intensity power-normalized to reflect local AI patent density during the pre-ChatGPT period (2010–2019). The maps highlight substantial within-city heterogeneity in AI innovation activity, even within major metropolitan areas.



(a) Panel A: Beijing



(b) Panel B: Shanghai



(c) Panel C: Shenzhen

Table B.1: Restricting to AI-Active Grids

This table reports estimates of changes in the number of new firms following the release of ChatGPT. For robustness, we restrict the sample to grids with at least one prior AI patent. *Post* is an indicator equal to one for quarters after 2022Q4, and *HighAI* is a dummy equal to one for grids with more than eight AI patents (the 75th-percentile threshold) filed during the period 2010–2019. Columns (1)–(3) use the following dependent variables: (1) number of new firms; (2) number of small firms; and (3) number of large firms. All regressions include city-by-quarter fixed effects and grid-by-calendar-quarter fixed effects to absorb seasonal patterns. Standard errors are clustered at the city level. Variable definitions are provided in Section A.1. *Notes*: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1) Num new firms	(2) Num small firms	(3) Num large firms
Post × HighAI	6.568*** (1.962)	11.62*** (2.815)	-5.859*** (1.942)
Constant	40.81*** (0.235)	16.04*** (0.338)	22.69*** (0.233)
Observations	162,512	162,512	162,512
R-squared	0.842	0.709	0.796
City × Quarter FE	Yes	Yes	Yes
Grid × Cal QTR FE	Yes	Yes	Yes

Table B.2: Matched Comparison with Nearby Non-AI Grids

This table reports estimates of changes in the number of new firms following the release of ChatGPT. As a robustness check, we construct the sample by matching each grid with at least one AI patent to the five geographically closest grids without AI patents (with replacement). *Post* is an indicator equal to one for quarters after 2022Q4, and *HighAI* is a dummy equal to one for grids with at least one AI patent filed between 2010 and 2019. Columns (1)–(3) use the following dependent variables: (1) number of new firms; (2) number of small firms; and (3) number of large firms. All regressions include city-by-quarter fixed effects and grid-by-calendar-quarter fixed effects to absorb seasonal patterns. Standard errors are clustered at the city level. Variable definitions are provided in Section A.1. *Notes:* *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	(1) Num new firms	(2) Num small firms	(3) Num large firms
Post × HighAI	4.345*** (1.354)	6.584*** (1.484)	-2.621*** (0.693)
Constant	18.17*** (0.113)	7.333*** (0.124)	9.889*** (0.0578)
Observations	977,408	977,408	977,408
R-squared	0.794	0.615	0.772
City × Quarter FE	Yes	Yes	Yes
Grid × Cal QTR FE	Yes	Yes	Yes

Table B.3: Changing Capital Threshold for Firm Size

This table reports estimates of changes in the number of new firms with different levels of registered capital following the release of ChatGPT. As a robustness check, we vary the cutoff used to define small firms. Specifically, in Columns (1)–(2), small firms are defined as firms with registered capital below 2 million RMB, with larger firms classified as large firms. In Columns (3)–(4), the cutoff is set at 3 million RMB, and in Columns (5)–(6), at 5 million RMB. *Post* is an indicator equal to one for quarters after 2022Q4, and *HighAI* is a dummy equal to one for grids with at least one AI patent. All specifications include city-by-quarter fixed effects and grid-by-calendar-quarter fixed effects. Variable definitions are provided in Section A.1. Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	Cutoff 2 Millions		Cutoff 3 Millions		Cutoff 5 Millions	
	(1) Num small firms	(2) Num large firms	(3) Num small firms	(4) Num large firms	(5) Num small firms	(6) Num large firms
Post × HighAI	7.680*** (1.368)	-3.096*** (0.473)	7.380*** (1.384)	-2.796*** (0.440)	7.157*** (1.386)	-2.573*** (0.421)
Constant	2.977*** (0.0419)	1.434*** (0.0145)	3.319*** (0.0424)	1.092*** (0.0135)	3.536*** (0.0425)	0.875*** (0.0129)
Observations	2,658,304	2,658,304	2,658,304	2,658,304	2,658,304	2,658,304
R-squared	0.752	0.756	0.762	0.751	0.769	0.731
City × Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Grid × Cal QTR FE	Yes	Yes	Yes	Yes	Yes	Yes