

WHEN AI MEETS ENTREPRENEURSHIP: EVIDENCE FROM THE COMMERCIALIZATION OF CHATGPT

EDOARDO MARCHESI[†] AND SIXUN TANG^{*}

September 4, 2026

[Click here for the latest version of this paper](#)

ABSTRACT. This paper examines how the commercialization of generative artificial intelligence influences new firm creation, leveraging the release of ChatGPT as a quasi-exogenous shock to AI accessibility. We develop a novel industry-level measure of AI compatibility that captures the extent to which AI tools can support early-stage entrepreneurial tasks. Using a difference-in-differences design, we find that industries more compatible with AI experienced an increase of over 10% in firm formation following ChatGPT's release. To investigate why, we examine three mechanisms: reduced experimentation costs, capital reallocation, and skill gap bridging. While we find no evidence supporting the latter two, we show that the increase in firm formation is concentrated in industries where AI can effectively assist and experimentation is critical—evidence consistent with a reduction in the cost of experimentation as the key driver. Moreover, new firms created after ChatGPT's release in AI-compatible industries are more likely to survive, grow faster, and attract more educated workers. Taken together, the mechanism and quality results suggest that AI lowers the cost of experimentation in a way that disproportionately benefits high-ability entrepreneurs, leading to both greater and higher-quality firm formation.

JEL CLASSIFICATION CODES: L26, O33

We thank Liudmila Alekseeva (discussant), Gregory Brown, Emanuele Colonnelli, Andrea Contigiani, Bo Cowgill (discussant), Julie Fu, Deivy Houeix, Steven Kaplan, Young Soo Jang, Mario Leccese, Ralf Meisenzahl, Pascal Noel; and participants at the 3rd Chicago Entrepreneurial Finance Workshop; Revelio Labs Seminar Series; Inter-Finance PhD Seminar; Chicago Booth Finance Brownbag Seminar; Chicago FED Brownbag Seminar; REER 2025; HEC Paris Entrepreneurship Workshop 2025; Fordham University; Loyola University; SFWA 2026 for helpful comments and suggestions. All errors remain our own. This research is supported by the Chicago Booth Fama Miller Research Center and the Chicago Booth Stevens Doctoral Program. IRB Approval was obtained from the University of Chicago IRB (IRB 25-1533).

[†] Quinlan School of Business, Loyola University Chicago. email: emarchesi@luc.edu.

^{*} Booth School of Business, University of Chicago. email: sixun-tang@chicagobooth.edu.

1. INTRODUCTION

Whether new technology favors the entrant or the incumbent is one of the oldest questions in economics. Schumpeter himself never settled on an answer, assigning the innovative function first to the upstart entrepreneur in *The Theory of Economic Development* (1934) and later to the large established firm in *Capitalism, Socialism and Democracy* (1942). The commercialization of generative artificial intelligence poses the question anew, and in unusually sharp form. With the release of ChatGPT in November 2022, capabilities that previously required specialized skills, hired expertise, or substantial capital—writing code, researching markets, drafting contracts, producing content—became available to anyone with an internet connection, essentially overnight and at near-zero marginal cost. In this paper, we show that this shock reshaped American entrepreneurship: it raised not only the number of new firms, but their quality.

Ex ante, the effect of AI on entry alone could go either way. Large language models allow an individual to perform tasks that once required a team, lowering the skill and resource requirements of getting started. But the same technology could just as plausibly entrench incumbents, who command the data, distribution, and technical talent that make AI most valuable; if AI commoditizes precisely the capabilities a newcomer would otherwise differentiate on, lower entry costs may be offset by lower expected returns to entry. By contrast, the prediction on quality is sharper. Canonical models of firm entry imply that if AI lowers barriers to entry, the marginal entrant it attracts should be of lower ability, dragging down the average quality of new ventures (Lucas Jr, 1978; Hopenhayn, 1992). On that view, a surge of AI-enabled founding is a surge of marginal firms.

How entrepreneurs respond to AI is no longer a hypothetical question. In an original survey of 1,036 U.S. founders who launched or began pursuing a venture after January 2023, more than half report using AI tools in the early stages of venture formation (Figure 1, Panel A), and nearly one in five report that, without these tools, they would have been unlikely to pursue their idea at all (Figure 1, Panel B). For a meaningful share of entrepreneurs, AI is already decisive not just in how they build their businesses, but in whether they start them.

Moving from founders’ self-reports to causal evidence on firm formation requires knowing where AI can actually assist someone starting a business, and no existing measure captures this. A central contribution of this paper is therefore the construction of two novel industry-level measures: the *AI Assistance Score* and the *Experimentation Intensity Score*. These measures are constructed using a bespoke multi-step pipeline that leverages large language models applied to a rich corpus of textual data, including business guides, entrepreneurship handbooks, and case studies. We identify the universe of tasks that entrepreneurs must undertake when starting a firm, and then evaluate each task in terms of how useful AI tools can be in supporting its completion (*AI Assistance Score*), as well as how reliant the

task is on trial-and-error experimentation (*Experimentation Intensity Score*). This approach draws inspiration from previous works such as Eloundou et al. (2023) and Eisfeldt et al. (2025), which assess task-level AI exposure in labor market settings using LLMs. However, to the best of our knowledge, our measures are the first to focus explicitly on entrepreneurial activity, providing a new framework for quantifying how AI supports the early-stage business formation process across industries.

Because these measures are generated by large language models rather than observed from entrepreneurs’ behavior, we validate them against the assessments of entrepreneurs themselves. Respondents to our survey who used AI rated its helpfulness across the same task categories that underlie our pipeline, and their ratings align closely with the LLM-generated scores, with various measures of correlation around 0.6 across task categories. Our scores thus capture how entrepreneurs themselves experience the usefulness of AI in starting a business, not merely a language model’s assessment of it.

Armed with these measures, we present novel evidence that the sudden democratization of AI tools stimulated entrepreneurial activity, particularly in industries where the tasks entrepreneurs must perform when founding and first operating a firm fall most squarely within the capabilities of large language models. Exploiting the release of ChatGPT as a quasi-exogenous shock to AI accessibility, we implement a difference-in-differences framework with varying treatment intensity. Our results show that, following the release of ChatGPT, industries with greater AI compatibility experienced a significant and persistent increase in new firm creation, averaging more than 10%.

To explore the mechanisms driving the observed increase in firm creation, we investigate three potential channels, each motivated by distinct theoretical rationales and accompanied by empirical tests. First, we examine whether AI reduces the cost of entrepreneurial experimentation. Starting a new firm often involves iterative trial-and-error, including idea validation, product testing, and business model refinement. AI tools may lower the cognitive and time costs associated with such activities, particularly in industries where experimentation plays a more central role. If this mechanism is operative, we expect the increase in firm creation to be concentrated in industries that are both highly AI-compatible and experimentation-intensive. We test this hypothesis using a triple-difference (DDD) design: specifically, we interact the post-ChatGPT indicator with both our AI Assistance measure and our Experimentation Intensity measure. The resulting triple interaction isolates whether the post-ChatGPT increase in firm creation among AI-compatible industries is disproportionately larger in industries where experimentation is most important.

Second, we assess whether the effect of AI on entrepreneurship operates through capital reallocation. In this view, AI-compatible industries may attract more funding post-ChatGPT, as investors anticipate future productivity gains. Increased investment could

in turn stimulate new firm formation, not because AI directly enables entrepreneurs, but because entrepreneurs follow capital into AI-compatible sectors. We test this hypothesis using mediation analysis, assessing whether investment flows into AI-compatible industries increased following the release of ChatGPT and whether these flows explain the observed rise in firm entry.

Third, we explore whether AI tools help individuals overcome human-capital barriers that would otherwise deter them from entering certain industries. Entrepreneurs often refrain from launching businesses in sectors where they lack the requisite skills or domain-specific knowledge. If AI-powered tools like ChatGPT can substitute for some of these missing capabilities, they may allow individuals to pursue self-employment despite skill mismatches. To test this hypothesis, we construct a novel, time-varying measure of individual-level skill gaps vis-à-vis any given industry. This measure compares each individual’s accumulated occupational skills—derived from employment histories linked to O*NET data—with the skill requirements of each industry, which is based on the skill profile displayed by individuals who successfully became entrepreneurs in the period prior to our study. We then examine whether individuals with larger skill gaps became relatively more likely to enter self-employment after the release of ChatGPT, compared to before.

Our findings point to the first mechanism. The increase in firm formation is concentrated in industries that are both highly AI-compatible and experimentation-intensive, consistent with the view that AI reduces the cost of entrepreneurial experimentation. Founders’ own accounts point in the same direction: when survey respondents who used AI describe how they employed it, the most frequently cited activities—generating and refining business ideas, researching customers and competitors, and testing marketing and content—are precisely the domains where trial-and-error and rapid iteration are central to venture formation. In contrast, we find little evidence that capital supply or skill-bridging channels are primary drivers of the observed effect.

Finally, we examine whether newly established firms differ in quality relative to their pre-AI counterparts. Against the canonical prediction that lower entry barriers attract lower-ability entrants, we find that firms founded after ChatGPT in highly AI-compatible industries exhibit higher short-term survival rates, expand faster, and attract more educated workers. This pattern suggests that the productivity gains from AI are tilted toward high-ability entrepreneurs, who benefit disproportionately from the technology, leading to both greater firm entry and higher average quality among new ventures.

To discipline the interpretation of these findings, we develop a stylized two-period model in which heterogeneous individuals choose between salaried employment and entrepreneurship. Individuals differ along two independent dimensions: entrepreneurial capability, which determines the fixed cost of starting a firm, and an idiosyncratic entry friction—orthogonal

to ability—that governs occupational sorting. When AI arrives, it asymmetrically reduces entry costs for high-capability individuals. The model delivers two key predictions: both the total mass of entrepreneurs and the share of high-capability individuals among them increase. The critical ingredient is the entry friction, which prevents the standard result that lower barriers attract only marginal, lower-quality entrants. Instead, AI unlocks entry for high-capability individuals who were previously “stuck” in salaried employment due to unfavorable friction draws, raising both the quantity and average quality of new ventures.¹

This paper contributes to two distinct strands of research. First, it advances the literature on the economic consequences of artificial intelligence. A growing body of work has examined how AI technologies affect worker productivity and firm performance. Field experiments and case studies show that access to AI tools can raise worker-level productivity in a variety of settings. These gains, however, are heterogeneous across worker types: improvements accrue primarily to lower-skilled workers (Dell’Acqua et al., 2023; Noy and Zhang, 2023; Brynjolfsson et al., 2025b). They are also heterogeneous across occupations: the threat of AI-driven replacement falls disproportionately on high-paying occupations (Eloundou et al., 2023; Hampole et al., 2025). A recent study further highlights this heterogeneity, showing that workers whose skills rely heavily on textbook and codified knowledge (typically younger workers) are especially vulnerable to job displacement (Brynjolfsson et al., 2025a).

At the firm level, both AI adoption and AI compatibility have been linked to changes in performance, employment, and valuations. Firms more engaged in AI development exhibit higher productivity (Yang, 2022), while firms investing more heavily in AI display higher sales growth and profitability (Babina et al., 2024). AI is also associated with aggregate increases in employment (Yang, 2022; Babina et al., 2024) and higher market valuations (Babina et al., 2024; Eisfeldt et al., 2025).

Much of this research has focused on how AI displaces or complements existing workers and how incumbent firms integrate AI into their operations.² By contrast, research on the economic consequences of AI for entrepreneurs (Otis et al., 2024; Bao et al., 2025; Gupta et al., 2025; Bena et al., 2026) has only recently started to flourish. We contribute to this literature by introducing a novel measure of how much AI can assist entrepreneurs in carrying out early-stage tasks (*AI Assistance*) and by documenting that industries where AI provides greater assistance experienced an increase of over 10% in firm formation. Our study differs from, yet complements, each of these papers in scope and design. Otis et al. (2024) examine AI’s impact on post-entry outcomes for a small group of entrepreneurs; instead, we focus on the decision to become one. Bao et al. (2025) study how AI may enable entrepreneurship in STEM sectors; our analysis extends to the full cross-section of

¹The theoretical framework is discussed in Subsection 6.3, with the full model presented in Appendix G.

²A more extensive review of the current findings on the economic impact of AI can be found in Subsection F.1 of the Online Appendix.

industries. In contemporaneous work, [Bena et al. \(2026\)](#) measure exposure from the task content of existing occupations; we construct an exposure measure tailored specifically to entrepreneurial tasks, capturing the nuances of early-stage firm creation rather than general occupational susceptibility to AI. Also in contemporaneous work, [Gupta et al. \(2025\)](#) study high-growth, VC-backed startups and trace increased entry to shifts in venture financing; we examine the universe of new firms and locate the entry margin in the entrepreneur’s own decision problem: AI lowers the cost of the experimentation inherent in early-stage tasks, making entry attractive for founders regardless of access to venture capital.

Second, this paper contributes to the literature on the determinants of entrepreneurship. One line of this work studies how entrepreneurial activity responds to exogenous shocks, including economic, environmental, and policy disruptions. Studies show that trade liberalization reshapes local firm entry ([Cui and Li, 2023](#)), natural disasters alter entrepreneurial dynamics ([Boudreaux et al., 2019](#)), input-price movements reallocate business formation ([Davis et al., 1996](#)), and commodity price shocks can lead to firm creation in areas highly dependent on local demand conditions ([Bernstein et al., 2022](#)). A subset of this literature focuses on technological shocks, showing that improvements in digital and communication infrastructure can lower entry costs and spur firm creation. Broadband expansion has been linked to increases in new firm formation, particularly in service-oriented sectors ([Hjort and Poulsen, 2019](#); [Hasbi, 2020](#); [Stephens et al., 2022](#)), while mobile technology diffusion has generated similar effects in developing economies ([Aker and Mbiti, 2010](#); [Hjort and Tian, 2021](#)). Other work examines how information frictions shape trade ([Steinwender, 2018](#)). We contribute to this literature by studying the entrepreneurial response to a frontier AI tool (ChatGPT) that directly augments or substitutes for entrepreneurial tasks, rather than to broad improvements in digital infrastructure. This approach allows us to explain heterogeneous effects across industries by linking AI compatibility to post-shock firm creation.

A separate line of inquiry on the determinants of entrepreneurship emphasizes the role of frictions such as access to capital ([Evans and Jovanovic, 1989](#); [Gompers et al., 2005](#); [Kerr and Nanda, 2011](#)), skill acquisition and human capital ([Lazear, 2005](#); [Levine and Rubinstein, 2017](#)), and the costs of experimentation ([Kerr et al., 2014](#); [Manso, 2016](#)). We contribute to this literature by providing causal evidence on the role of experimentation costs—an important but difficult-to-measure friction—using an exogenous technological shock that plausibly relaxes these constraints in a subset of industries. Our results align with experimental studies showing that AI is particularly effective at supporting idea generation, problem solving, and enhancing human creative work ([Boussioux et al., 2024](#); [Doshi and Hauser, 2024](#); [Meincke et al., 2024](#)).

The rest of the paper is structured as follows. Section 2 describes our data sources, and Section 3 the construction of the AI Assistance and Experimentation Intensity measures

along with other key variables. Section 4 presents the empirical strategy and the main findings on AI compatibility and firm creation. Section 5 explores three potential mechanisms and evaluates their validity. Section 6 examines firm-level outcomes, and Section 7 concludes.

2. DATA SOURCES

The objective of this paper is to study firm creation following the arrival of ChatGPT, a recent event. In principle, the ideal data would be administrative records, such as tax filings, which capture the full universe of new businesses. However, these sources are either not available at the frequency needed for our analysis or are released only with long delays, making them unsuitable for examining such a recent event. To circumvent this limitation, we use granular data from Revelio Labs, which provides near-real-time information on new firm creation and workforce histories.

2.1. Revelio Labs Data. To measure entrepreneurial activity and workforce dynamics, we draw on detailed workforce data from Revelio Labs (RL), which provides broad coverage of U.S. firm creation and employment histories from 2020Q4 to 2024Q1. RL aggregates publicly available professional profile information, primarily sourced from platforms such as LinkedIn. Firms appear in the dataset when they or their workers maintain a presence on the platforms RL aggregates. Crucially, a firm’s entry date is not the date its company page was created: we date entry using the founding date reported on the founder’s profile or, where unavailable, the earliest start date employees report for positions at the firm. Because entry timing is elicited from reported histories rather than platform activity, a firm that establishes its LinkedIn presence years after founding still enters our panel at its founding date. Platform adoption therefore affects which firms we observe, but not when we observe them entering. This dataset has been widely used in recent research to study human-capital flows, workforce composition, and organizational dynamics (Kwan et al., 2025; Labaschin et al., 2025; Wang et al., 2025).

We use RL’s firm-level panel to measure new firm creation. Firms in RL data are tagged with North American Industry Classification System (NAICS) codes, allowing us to define new firms by their entry date as elicited above, assigning each firm to the quarter of its founding. For the firms identified, we know which industry they operate in, their location, their date of registration and date of shutdown (if any).

At the individual level, we leverage RL’s employment history data. RL data allow us to construct comprehensive resumes of workers’ education and employment histories, including degree earned, jobs held, job titles, and employment durations. Moreover, each position is classified into occupational categories consistent with O*NET Standard Occupational Classification (SOC) codes. We extract data for all U.S. individuals who were active between 2020Q4 and 2024Q1.

It is important to note that RL data are limited to firms and individuals present on LinkedIn and similar platforms, introducing coverage bias. Consequently, our sample may overrepresent some portions of the economy or employment types. To address these concerns, we benchmark our data against official statistics from the U.S. Census Bureau. Specifically, we benchmark RL against the Business Dynamics Statistics (BDS), the gold-standard source for firm creation data in the United States. In Subsection C.2 of the Online Appendix, we document a series of validation exercises that assess how well RL replicates the key features of new firm formation captured in BDS. These tests compare RL and BDS along several dimensions, including the distribution of new firms across industries, the time series of firm creation activity, and the relationship between firm creation levels in the two datasets. We find that, to a large extent, RL closely mirrors BDS in terms of both the relative share and the rank-ordering of industries by new firm births. Furthermore, the number of new firms recorded in RL tracks that in BDS with high correlation across industries and years, despite RL’s known undercoverage of certain traditional or low-tech sectors. These findings suggest that, for the purpose of analyzing patterns of firm creation across industries and over time—as we do in this paper—RL provides a reliable and representative proxy for official data. While RL may not be representative of all economic measures or subpopulations, its coverage is sufficiently broad and consistent to support the main analyses in this paper. In Subsection C.1 of the Online Appendix we provide further information regarding the data-generating process underlying RL, while in Subsection C.3 we detail the data manipulation steps to create the variables from the various RL datasets.

2.2. Textual Data Sources. To construct our novel industry-level measures of AI compatibility and experimentation intensity, we leverage detailed textual data from entrepreneurship handbooks, business guides, and case studies. This subsection describes these textual sources and motivates their suitability for extracting the universe of tasks that entrepreneurs must perform when establishing new firms.

Since our objective is to establish to what extent—if any—LLMs can facilitate entrepreneurial entry and whether this is conditional on the type of tasks for which they provide assistance, we first generate the universe of tasks that entrepreneurs have to successfully carry out in order to get their business established and starting to operate. One major limitation we face is that, while there are repositories of tasks entailed by different jobs and occupations, to date, there is no source of data listing all the activities a new entrepreneur needs to carry out. To overcome this challenge, we take advantage of textual data from various sources such as (a) books focused on entrepreneurship or becoming an entrepreneur, (b) business guides from official sources,³ and (c) business school case studies. The unstructured sources

³Examples of these sources are the Small Business Administration and Chambers of Commerce.

of data are then used as raw data for an LLM to extract tasks entrepreneurs have to carry out.

In Subsection D.1 of the Online Appendix we offer additional information on the textual data sources we used.

2.3. Preqin Data. To capture shifts in private capital availability that could influence firm creation, we draw on investment data from Preqin, a dataset that collects detailed information on venture capital and private equity funding. We extract data on the investments made between 2020Q4 and 2024Q1 for which we know: (1) the industry of the company receiving the funding and (2) the amount of funding. We subset the dataset to include only funding coming from venture capital and private equity investors that is directed to ventures operating in the U.S. Finally, from this subset of investment we retain only the first instance of equity funding a given company received⁴.

Because Preqin primarily sources its data from investors with disclosure obligations, some transactions may be missing or reported incompletely, introducing potential measurement error. However, as long as investors with and without disclosure requirements do not substantially differ in their investment style—in terms of industry targets and timing—measurement error should be of little concern.

2.4. O*NET Data. To examine whether AI facilitates entrepreneurship by bridging human-capital gaps, we rely on occupational skill data from the Occupational Information Network (O*NET), the primary occupational information database maintained by the U.S. Department of Labor’s Employment and Training Administration (USDOL/ETA). The O*NET database provides standardized measures of skills required across nearly 1,000 U.S. occupations. These data are collected through structured surveys administered to job incumbents and occupational experts, ensuring relevance and accuracy in reflecting contemporary occupational requirements.

In the O*NET framework, each skill associated with an occupation is rated along two key dimensions: *importance* and *level*. For our analysis, we focus on the skill *level* ratings, which capture the degree of proficiency or complexity required for effective performance within an occupation. This measure can be interpreted as the intensity of a skill that an individual is expected to possess to perform successfully in a given occupation. Level ratings are determined through survey responses using an anchored numeric scale from 0 (no proficiency required) to 7 (highest proficiency). Behavioral examples provided within O*NET documentation anchor each numeric rating, offering practical interpretations of skill proficiency at different scale points. For example, lower-level anchors reflect basic capabilities

⁴The object is to create a comprehensive data set of early stage equity investments injected into U.S. entrepreneurial firms during our period of interest.

suitable for simple tasks, whereas higher-level anchors correspond to advanced capabilities necessary for more complex or specialized occupational tasks.

Each occupation, identified by its O*NET Standard SOC code, is linked systematically to its relevant skill-level ratings, creating a comprehensive occupation-skill dataset.

2.5. Additional Data Sources. We complement the four primary data sources with four auxiliary ones, which supply control variables, rival exposure measures, and identification checks. Their construction details are reported alongside the analyses that use them.

First, we use firm-level financial data from Compustat to construct time-varying industry controls that account for pre-existing economic conditions: Return on Assets (income before extraordinary items over total assets), R&D intensity (R&D expenses over revenue), CapEx intensity (capital expenditures over revenue), cash holdings over total assets as a proxy for liquidity, and market leverage (long-term and current debt over the market value of equity). Since the analysis is at the industry level, we compute median values at the three-digit NAICS level.

Second, we use the detailed fixed-asset estimates published by the Bureau of Economic Analysis (BEA) to measure each industry’s IT share of the capital stock, fixed at 2018 values. This measure enters the identification analysis of Subsection 4.3 as a rival exposure capturing pre-existing digitalization; its construction is detailed in Subsection A.1.6 of the Online Appendix.

Third, we use the occupation-level teleworkability classification of Dingel and Neiman (2020), aggregated to three-digit NAICS industries using OEWS employment weights. This measure enters the identification analysis of Subsection 4.3 as a rival exposure capturing remote-work capability; its construction is likewise detailed in Subsection A.1.6 of the Online Appendix.

Fourth, we use the Census Bureau’s Business Formation Statistics (BFS), which record weekly applications for an Employer Identification Number at the three-digit NAICS level. Because EIN applications exist independently of any professional platform, the BFS provide a measure of entry that is immune to platform coverage; we use them in Subsection 4.3 to measure each industry’s pandemic-era entry boom. Finally, Subsection 4.4 describes an original survey of new and aspiring U.S. entrepreneurs that we fielded to validate the AI Assistance measure.

3. CONSTRUCTION OF KEY MEASURES

Since our objective is to study how the expansion of access to AI tools influences firm creation, a central requirement of our analysis is a cross-sectional measure of exposure. To this end, we construct two novel industry-level indices: the *AI Assistance Score*, which captures the extent to which large language models can aid entrepreneurs in completing

early-stage tasks, and the *Experimentation Intensity Score*, which quantifies how reliant the completion of those tasks is on trial-and-error processes.

Beyond these task-based indices, this section also introduces a complementary measure designed to evaluate whether AI helps potential entrepreneurs overcome human-capital barriers. Specifically, we develop an individual-by-industry *Skill Gap* measure that compares an individual’s accumulated skills to the requirements of different industries, enabling us to test whether AI reduces frictions that would otherwise deter entry.

3.1. AI Assistance and Experimentation Intensity Scores. To measure cross-sectional exposure, we construct two novel indices: *AI Assistance* and *Experimentation Intensity*. A key challenge is that, while task repositories exist for occupations (e.g., O*NET), none exists for the activities entrepreneurs undertake when starting a firm. We address this gap by leveraging textual sources such as entrepreneurship handbooks, business guides, and case studies, which we process with large language models to extract the universe of entrepreneurial tasks. In what follows, we detail the workflow used to derive the two measures. In the interest of transparency, we report, in Subsection D.2 of the Online Appendix, the prompts, system messages, and model configurations used.

In a preliminary step, we take each source of textual data described in Subsection 2.2 and divide it into smaller chunks to ensure that the LLM can process the content effectively⁵. Each chunk is then submitted to ChatGPT-4o via the API for processing.

In the first layer of the workflow, we instruct the model to impersonate an experienced entrepreneur and focus on relevant sections of its knowledge base when processing each request. Each text chunk is presented individually to the LLM, which is then asked to identify any task related to establishing and first running a firm, and to produce a detailed description of the task. The resulting list of tasks includes activities such as “*Complete and submit necessary paperwork for firm registration*”, “*Design a minimum viable product (MVP)*”, “*Acquire early customer feedback on product/service*”, “*Establish and negotiate initial relations with suppliers*”.

Once the model in the first layer has produced the universe of identified tasks from the textual data, the extracted tasks contain duplicates both across and within sources, as the same task may appear multiple times within a source or recur across sources. We use a second LLM to identify duplicates based on semantic similarities and produce condensed versions of the task description. More specifically, we first embed task descriptions to identify groups of similar tasks. Text embedding converts textual content (e.g., words, sentences, or

⁵This pre-processing step is standard practice. Operationally, LLMs have a maximum context window (e.g., 4k, 32k, even 200k tokens depending on the model); if the text to be analyzed exceeds this limit, it cannot be processed in a single pass. Moreover, even when a text technically fits within the context window, very long inputs can reduce the model’s ability to attend to all parts equally. Chunking allows the model to focus on coherent segments, thereby improving accuracy and consistency.

documents) into numerical vectors for use in machine learning models. The key idea is to represent text in a way that captures its semantic meaning—so that similar texts are close to each other in vector space. Once we have grouped the tasks, we ask an LLM to come up with a task name for the group and produce a “collapsed” version of the task description which distills the main elements of the task descriptions collapsed. Through this process we are able to produce the universe of unique entrepreneurial tasks.

At this stage, we have the universe of tasks entrepreneurs must carry out; however, a given task may not apply to all industries, and even when it does, its description may not be specific enough to convey the nuances of performing it in a particular industry. We overcome this issue by bringing a third LLM layer into play. For this layer, we ask ChatGPT to impersonate an entrepreneur with extensive experience in running a firm in a specific industry. We then form all possible task–industry pairs at the three-digit NAICS level and present them to the model. For each pair, the model is given the task name and description and the industry name and description, and asked whether the task is relevant to that industry. If the task is deemed relevant, the model is then asked to use the industry description to tailor the task description to the industry. This process results in the list of relevant entrepreneurial tasks for each industry, with descriptions tailored to the sector’s specifics. Across the 81 three-digit NAICS industries in our sample, we extract 15,957 relevant tasks drawn from 847 unique tasks, with each industry containing 197 distinct tasks on average. Figure 2 presents the distribution of the number of tasks across industries.

At this stage, we have the list of tasks relevant to each three-digit NAICS industry. We introduce the fourth and last LLM layer to score these tasks across three dimensions. The first dimension represents how “important” the task is—where importance is defined as how recurrent, time-consuming or prevalent the task is for entrepreneurs starting a company in a specific industry. The second dimension is the degree of AI compatibility of the task, which we call the AI Assistance Score. The AI Assistance Score reports the degree to which an LLM can provide support in completing the task, with higher scores representing higher levels of support. The third dimension, which we term Experimentation Intensity, summarizes the degree of experimentation required by an entrepreneur in order to successfully carry out a task; in short, this dimension measures “linearity” versus trial-and-error based task execution.

We ask the fourth LLM to impersonate an entrepreneur with extensive industry knowledge—based on the industry at hand—and provide it with scoring rubrics to produce values for task Importance, AI Assistance and Experimentation Intensity. To score the tasks across the three dimensions, we implement Chain-of-thought (CoT), which means that instead of asking the model to provide a numerical score, we first ask the model to reason and produce

a rationale as to why a task may be “Important,” “AI Assistable,” or “Experimentation Intense,” and then, based on the rationale, produce a numerical score. Chain-of-thought helps models reason more effectively and accurately, especially on complex or multi-step problems. Instead of jumping directly to an answer, CoT encourages the model to generate intermediate reasoning steps, mimicking how humans solve problems by thinking out loud. Once the model has produced a rationale concerning the dimension at hand, the LLM is then asked to produce a numerical score based on it. We run the score estimation procedure three times and average the results in order to smooth out possible instances of hallucinations.

Since our analysis is conducted at the industry level, we collapse the task-level scores into industry-level measures using an importance-weighted average over the tasks relevant to each industry, so that marginal or rarely performed tasks do not receive disproportionate weight. For industry i with relevant-task set $\mathcal{T}_i$, the AI Assistance Score is

$$(1) \quad \text{AI Assistance}_i = \sum_{k \in \mathcal{T}_i} \frac{\text{Importance}_{ik}}{\sum_{l \in \mathcal{T}_i} \text{Importance}_{il}} \text{AI Assistance}_{ik},$$

the importance-weighted mean of the task-level scores. The Experimentation Intensity Score is constructed analogously, replacing $\text{AI Assistance}_{ik}$ with $\text{Experimentation Intensity}_{ik}$. Figure 3 presents the distributions of the AI Assistance (Panel (a)) and Experimentation Intensity (Panel (b)) scores.

In Table 1 we report the top 20 and bottom 20 industries by AI Assistance Score. Industries at the top of the AI Assistance distribution—Web Search Portals & Information Services (NAICS 519), Data-Processing and Hosting Services (518), and the broader family of Professional, Scientific & Technical Services (541)—all revolve around the creation, manipulation, or dissemination of digital information. Their core tasks (e.g. software engineering, data curation, text generation, legal drafting, and graphic design) are language- or pattern-based and can be executed, accelerated, or even autonomously handled by large language models and other machine-learning systems. Seemingly “old-economy” sectors such as Electronics Manufacturing (334) and several retail subsectors (Electronics & Appliance Stores, General Merchandise, Clothing & Accessories) also score highly, for a more revealing reason than their digital counterparts. Launching a firm in these industries requires standing up an operational backbone—production planning, quality and inventory control, supplier coordination, and demand forecasting—that a founder would otherwise have to build personally or staff with a dedicated operations team. Large language models can now scaffold much of this setup work, from designing stock-management and forecasting routines to drafting supplier specifications and contracts, lowering a fixed cost of entry that is especially steep in physical, inventory-intensive businesses. By contrast, industries at the bottom—Crop Production, Forestry and Logging, Oil & Gas Extraction, Primary Metal Manufacturing,

and Pipeline Transportation—are dominated by highly physical, location-specific production processes governed by engineering constraints, heavy regulation, and harsh operating environments. While narrow AI can aid predictive maintenance or drilling optimization, the share of tasks that can be materially substituted or augmented by foundation models remains small, keeping their AI Assistance Scores low.

Table 2 reports the top 20 and bottom 20 industries by Experimentation Intensity. Topping this distribution are creative and digital content sectors such as Motion Picture & Sound Recording (512), Broadcasting & Content Providers (516), and again Web Content & Information Services (519). These industries live on rapid A/B testing, audience feedback, and continual format tweaks, making experimentation intrinsic to value creation. Food Services & Drinking Places (722) also score highly because menu design, pricing, and service concepts are constantly refined in response to local consumer tastes. Even the “plumbing” of the digital economy, Data-Hosting & Web-Infrastructure (518), ranks near the top: cloud providers roll out new features, security patches, and pricing models continuously, learning from usage telemetry in real time. At the other extreme lie capital-intensive, heavily regulated sectors such as Pipeline Transportation (486), Utilities (221), Heavy & Civil Engineering Construction (237), and Primary Metal Manufacturing (331). Their production functions entail large irreversible investments and stringent safety or environmental standards that severely limit scope for mid-course experimentation. Similarly, Hospitals (622) and Nursing & Residential Care Facilities (623) face ethical and regulatory constraints that suppress rapid trial-and-error, anchoring them to the lower tail despite growing analytics adoption.⁶

3.2. Skill Gap Measure. To empirically evaluate the hypothesis that AI tools reduce entry barriers by bridging human-capital deficits, we construct a novel, time-varying measure of individual skill gaps relative to industry requirements. In what follows, we detail the construction of this measure and clarify how it enables rigorous tests of the human-capital bridging mechanism.

One of the potential mechanisms linking AI to firm creation we envisioned revolves around LLMs bridging potential entrepreneurs’ human-capital gaps. To test this mechanism we estimate the skill gap for individual j in quarter t vis-à-vis industry i by using the skills s connected to each occupation o reported in O*NET and their skill level ratings described in Subsection 2.4.

We first create a measure of skill profile, which is the collection of skills retained by an individual j and their respective level of proficiency. The skill profile is time varying, as

⁶To further clarify what the AI Assistance measure captures, we correlate it with various industry characteristics in Table OA1 (Online Appendix, Subsection A.1). AI Assistance does not exhibit meaningful correlations with these characteristics, underscoring that the measure reflects a distinct dimension of the entrepreneurial process—namely, the extent to which AI can support early-stage entrepreneurial tasks.

it depends on the occupations covered by the individual over time. Implicitly, we assume that skills do not depreciate. First, let O_j^t be the set of occupations held by individual j at or before time t . Let Skill Level_{so} be the level of proficiency required for skill s when an individual is employed in occupation o .⁷ Then the skill profile—the skill strength retained by the individual—for individual j at time t with respect to skill s is:

$$(2) \quad \text{Skill Profile}_{jts} = \max_{o \in O_j^t} \{\text{Skill Level}_{so}\}$$

Since we assume that skills do not depreciate over time, and multiple occupations may require the same skill, we assign each skill the highest level attained across all occupations the individual has ever held, hence the $\max_{o \in O_j^t}$.

Next, we create a measure of *industry requirements* using the skill profiles of individuals who successfully became entrepreneurs prior to the period of analysis to identify the skills historically associated with entry in each industry. The intuition is that these profiles represent the empirical combination of capabilities that allowed entrepreneurs to overcome the barriers to entry in a given sector before the widespread availability of generative AI.

Formally, let $\mathcal{E}_i$ denote the set of individuals who became entrepreneurs in industry i between 2010Q1 and 2022Q3,⁸ and let t_j^* be the quarter in which individual $j \in \mathcal{E}_i$ first appears as an entrepreneur. For each of these individuals, we extract their skill profile in the quarter immediately preceding entry, $\text{Skill Profile}_{j,t_j^*-1,s}$, where each element s corresponds to one of the standardized O*NET skills described above.

We then define the *industry requirement* for industry i and skill s as the mean pre-entry skill level across all historical entrepreneurs in that industry:

$$(3) \quad \text{Industry Requirement}_{is} = \frac{1}{N_i} \sum_{j \in \mathcal{E}_i} \text{Skill Profile}_{j,t_j^*-1,s},$$

where $N_i = |\mathcal{E}_i|$ is the number of entrepreneurs used to construct the benchmark for industry i .

The resulting vector $\text{Industry Requirement}_i = (\text{Industry Requirement}_{is})_{s=1}^S$ provides an empirically grounded representation of the skill bundle that historically characterized founders in each industry prior to the diffusion of AI. By anchoring the benchmark to pre-ChatGPT entrepreneurs, the measure captures the human-capital profile that historically characterized successful entrants in the pre-AI regime and remains unaffected by post-treatment changes in the composition of skills among entrepreneurs.

Next, we calculate the skill deficit of individual j for each skill s in quarter t with respect to industry i . The deficit captures only shortfalls relative to the industry benchmark: skills

⁷The same skill may be required by multiple occupations, but its required level may differ across them.

⁸Although our extraction covers individuals active between 2020Q4 and 2024Q1 (Subsection 2.1), each profile contains the individual’s full retrospective employment history. Entrepreneurial entries prior to the extraction window are therefore observable through the employment histories of sampled individuals.

in which the individual meets or exceeds the requirement are set to zero, so that surpluses in one skill do not offset deficiencies in another. Formally,

$$(4) \quad \text{Deficit}_{jits} = \max \{0, \text{Industry Requirement}_{is} - \text{Skill Profile}_{jts}\}$$

To calculate the skill gap for a given individual j in quarter t with respect to industry i we aggregate over the skill deficit values, using a Euclidean Norm to penalize large shortfalls more heavily:

$$(5) \quad \text{Skill Gap}_{jit} = \sqrt{\sum_s (\text{Deficit}_{jits})^2}$$

For the analysis, we standardize the skill gap measure using the pre-ChatGPT mean and standard deviation. Anchoring the transformation to the pre-treatment distribution ensures that the scale of the variable is fixed prior to the arrival of ChatGPT, avoiding the risk that post-treatment dynamics mechanically influence the normalization itself. This preserves the comparability of the measure across time and is essential for any attempt at a causal interpretation of pre/post differences. In addition, standardization enhances statistical stability in our regressions: it prevents coefficients from being sensitive to the arbitrary scale of the raw Euclidean norm and facilitates estimation in specifications with interaction terms (e.g., Skill Gap $\times$ Post), where unstandardized variables could otherwise generate unstable or difficult-to-interpret results.

4. THE EFFECTS OF AI ON FIRM CREATION

In this section, we turn to the central empirical question of the paper: does the commercialization and widespread accessibility of AI technologies influence entrepreneurial activity? To answer this, we exploit variation in industry-level AI compatibility around the release of ChatGPT in late 2022, which we treat as a quasi-exogenous shock to the availability of generative AI.

Our dataset is structured at the industry-by-quarter level from 2020Q4 through 2024Q1. It combines measures of firm creation, financial and economic controls from the sources outlined in Section 2, and our LLM-generated measures of AI Assistance and Experimentation Intensity. Industries are defined at the three-digit NAICS level, and we restrict the sample to sectors where private firm formation is economically meaningful and comparable. Table 3 reports descriptive statistics for the key variables in our empirical analysis. Panel (a) presents industry-by-quarter variables, including our primary outcome measure—the number of newly established firms—as well as investment flows and financial controls. Firm creation exhibits substantial variation across industries and over time, with a mean of 14,064 new firms per industry-quarter and a standard deviation of over 32,000. Investment flows are similarly heterogeneous, with a mean of \$1,730,000 and a maximum of nearly \$95 million.

Panel (b) summarizes industry-level variables that remain constant over time, including our standardized AI Assistance and Experimentation Intensity measures.

The remainder of this section introduces our empirical strategy, outlines the identification assumptions, and presents baseline estimates of the impact of AI Assistance on new firm creation. Robustness tests and validation exercises supporting the interpretation of the AI Assistance measure appear at the end of the section.

4.1. Empirical Strategy. Our goal is to estimate how the increased accessibility of generative AI tools, exemplified by the public release of ChatGPT, affects entrepreneurial activity. Identification exploits the fact that potential entrepreneurs are differentially exposed to AI depending on the industry in which they plan to operate.⁹

The release of ChatGPT in November of 2022 was an unprecedented and unanticipated shock to the availability of LLM technologies for the public. While LLMs had been available as early as 2018, their use was restricted due to the proprietary nature of the technology. The release of ChatGPT marked the first instance of true democratization of the technology. We treat the arrival of ChatGPT as an unexpected shock to the availability of the technology and use it to study its effects on firm formation. The choice of this event as a shock is guided by the fact that the timing of the release was driven by ethical and safety considerations and model readiness—factors unrelated to business dynamism. Figure 4 provides additional evidence of the absence of anticipation using Google Trends search data for the term “ChatGPT.” Search intensity is flat at or near zero throughout 2022 until the week beginning November 27, 2022 (which contains the November 30 release date), when it rises sharply. The lack of any prior trend in search activity confirms that the public debut of ChatGPT did not follow a period of increasing attention, further supporting our interpretation of the release as an unanticipated shock to AI tool accessibility.

In our DiD setting, cross-sectional exposure arises from the ex-ante benefit entrepreneurs derive from using AI virtual assistants to start a firm in a given industry.

We use a difference-in-differences design with varying treatment intensity to understand how the arrival of ChatGPT impacts business formation, conditional on a given industry’s compatibility with AI.

At the industry-by-quarter level, the main specification is:

$$(6) \quad \text{IHS}(\text{N. New Firms})_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

where the dependent variable $\text{IHS}(\text{N. New Firms})_{it}$ is the inverse hyperbolic sine of the total number of new firms. AI Assistance_i is the industry level proxy measuring LLMs’ ability to support entrepreneurs through the completion of various tasks, which we standardize for ease

⁹Different industries represent distinct collections of tasks, which are differentially exposed to the benefits of AI in the process of task completion.

of interpretation. Post_t is an indicator equal to 1 from 2022Q4 onward, the periods during which ChatGPT became publicly available. μ_i are industry fixed effects used to control for industries' intrinsically higher levels of firm creation, which may not be related to the arrival of ChatGPT. μ_t are quarter fixed effects, which allow us to net out common and aggregate time trends unrelated to the dissemination of the new technology, which may still shape business formation. We also include a series of industry level, time-varying controls which collectively help isolate the effect of AI on firm entry from broader economic and financial conditions.

In our setting, recovering the dose-response effect¹⁰ requires (conditional) comparability across industries prior to the shock in terms of outcomes. To lend credibility to this assumption we inspect the coefficients from the dynamic model

$$(7) \quad \text{IHS}(\text{N. New Firms})_{it} = \beta_0 + \sum_{t \neq 2022q3} \beta_t \text{AI Assistance}_i \times 1\{D = t\}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The coefficients on the interaction between the AI Assistance measure and a quarter dummy identify the differential change in outcome due to a change in the independent variable (holding within industry and quarter characteristics fixed). For industries to be comparable, the coefficients on the leads (β_t) need to be insignificant.

4.2. Results and Interpretation. We first ask whether firm formation is associated with the assistance LLMs provide to entrepreneurs.

The main analysis uses as level of industry aggregation three-digit NAICS industries, over the period 2020Q4 to 2024Q1. In this analysis, we exclude certain industries that are either related to public administration or sectors where firm formation by private entrepreneurs is not applicable¹¹. These sectors are excluded because they primarily consist of government

¹⁰With a binary treatment, β_1 reduces to the canonical difference-in-differences estimate of the average treatment effect on the treated (ATT). With a continuous treatment intensity, it instead identifies an average causal response—the average effect of a one-standard-deviation increase in AI Assistance on the change in firm creation. Treatment timing is common to all industries, so the specification is a standard two-way fixed-effects estimator and the concerns associated with staggered adoption do not arise; because treatment intensity is continuous, however, the causal interpretation of β_1 rests on the strong parallel-trends assumption of Callaway et al. (2024)—that, absent treatment, average entry paths would have evolved identically across industries at every level of AI Assistance. Three sets of results support this assumption. The leads are flat over an estimation window extended into the pre-pandemic period (Subsection 4.3). Binarizing treatment at the mean or median of AI Assistance—which requires only the standard parallel-trends assumption—delivers the same conclusion (Subsection A.1.3 of the Online Appendix). And the response is monotone across deciles of AI Assistance, indicating that the estimate does not rest on adverse weighting across the dose distribution (Subsection A.1.4 of the Online Appendix).

¹¹Specifically, we drop the following three-digit NAICS industries: 491 (Postal Service), 521 (Monetary Authorities-Central Bank), 525 (Funds, Trusts, and Other Financial Vehicles), 813 (Religious, Grantmaking, Civic, Professional, and Similar Organizations), 711 (Performing Arts, Spectator Sports, and Related Industries), 712 (Museums, Historical Sites, and Similar Institutions), 713 (Amusement, Gambling, and

functions or nonprofit organizations where traditional firm entry is either not possible or not relevant to entrepreneurial dynamics studied in this paper.

To understand industries' comparability across different AI Assistance levels, we begin by examining the coefficients from the dynamic specification presented in Equation 7. As shown by the insignificant coefficients on the leads, we do not observe meaningful pre-trends in firm creation. Moreover, the coefficients on the lags are positive and fairly similar across quarters, suggesting a persisting impact (Figure 5). These results remain robust after controlling for broader economic and financial conditions (Figure 5, Panel (b)).

Next, we test for existence of dose-response using the compact regression in Equation 6. Industries that are more compatible with the use of LLMs in entrepreneurial tasks—measured by AI's ability to assist entrepreneurs in the process of venture formation—experience a positive and significant increase in firm creation (Table 4, Column (1)). These effects do not appear to be driven by contemporaneous economic and financial conditions: the coefficient on the main variable hardly changes when controls are added (Table 4, Column (2)). When evaluated at the pre-treatment mean number of firms (13,673), this 0.105 coefficient corresponds to an increase of roughly 11 percent.¹² To gauge economic significance, a coefficient of 0.105 implies that a one-standard deviation increase in AI Assistance raises new firm creation by over 10 percent. Evaluated at the pre-treatment mean number of new firms (13,673), this corresponds to about 1,513 additional firms, not a new total. This calculation uses the mean of the outcome in levels, since interpreting magnitudes requires returning to the original scale rather than the IHS-transformed mean.

Taken together, these findings suggest that the increased accessibility of AI tools coincided with a notable rise in new firm creation in AI-compatible industries. The remaining question is whether this differential reflects the arrival of AI or the pre-existing dynamics of the industries most compatible with it; the next subsection addresses that identification challenge directly.

Recreation Industries), and the series 921 to 928 which cover various segments of Public Administration (921 – Executive, Legislative, and Other General Government Support; 922 – Justice, Public Order, and Safety Activities; 923 – Administration of Human Resource Programs; 924 – Administration of Environmental Quality Programs; 925 – Administration of Housing Programs, Urban Planning, and Community Development; 926 – Administration of Economic Programs; 927 – Space Research and Technology; 928 – National Security and International Affairs).

¹²With $\beta_1 = 0.105$ and pre-treatment mean $y_0 = 13,673$, the implied post-period change for a one-standard-deviation increase in AI Assistance is

$$\Delta y = y_0(\cosh \beta_1 - 1) + \sqrt{1 + y_0^2} \sinh \beta_1 \approx 1,513 \text{ firms,}$$

which corresponds to

$$\frac{\Delta y}{y_0} \approx e^{\beta_1} - 1 \approx 11.07\%.$$

4.3. Threats to Identification. Identification in our design rests on a counterfactual: absent ChatGPT, entry after 2022Q3 would not have shifted differentially across industries by AI Assistance. Any force that moves high-AI industries differently around that date can therefore masquerade as a treatment effect, and our setting faces two distinct layers of this risk. The first layer is specific to the pandemic and its timing: the baseline pre-period sits entirely inside the pandemic era. The eight quarters from 2020Q4 to 2022Q3 contain the decay of the largest business-formation surge in the modern record, and any residual pandemic dynamics in exposed industries—continued decay, recovery, normalization—would load onto the post-ChatGPT indicator and masquerade as a treatment effect. The same shock accelerated adoption of the professional-profile platforms from which our outcome data derive, so the measurement of entry could inherit pandemic dynamics as well. The second layer is not specific to the pandemic. The industries that rank high on AI Assistance are digital, knowledge-intensive industries that were remote-capable, IT-intensive, and in some cases fast-growing well before 2020; these characteristics could have produced a distinctive post-2022 trajectory on their own, and the pandemic may have accelerated them, but the confound would exist even had it not. The relevant threat, in either layer, is therefore not correlation per se: high-AI industries can be more digital, more remote-capable, or larger without invalidating the design. To confound the estimate, a characteristic must generate a differential change in entry timed to late 2022. Each exercise below confronts a candidate with that requirement.

We first re-estimate the specification of Subsection 4.1, unchanged, over an estimation window extended to 2018Q4. The extension adds eight quarters untouched by the pandemic, so the parallel-trends assumption can be judged against a clean baseline. Figure 6 reports the event study over 2018Q4–2024Q1. The leads are flat throughout—across the pre-pandemic quarters, through the onset and peak of the pandemic surge, and up to the reference quarter—and the coefficients break sharply upward at 2022Q4. The flat pre-pandemic leads show no evidence that high-AI-Assistance industries were already on differential entry trajectories before COVID; the similarly flat coefficients through the pandemic indicate that the pandemic entry surge itself did not differentially load on AI Assistance. The compact estimate over the extended window is 0.112, statistically significant at the 1% level and nearly identical to the baseline (Table 5).¹³ Entry paths did not diverge by AI compatibility when the pandemic arrived; they diverged when ChatGPT did. Because the results carry over to the extended window, we retain 2020Q4–2024Q1 as the main estimation window for ease of exposition: the shorter window balances pre-treatment and post-treatment quarters and keeps the effects of the shock in visual focus, rather than compressing them at the end of a long pre-period.

¹³The estimates with and without controls differ only at the fourth decimal and therefore coincide as printed in Table 5.

The extended window shows that the pandemic surge did not load on AI Assistance. A residual concern is dynamic: industries that boomed during the pandemic may follow distinctive post-boom paths—continued decay, recovery, normalization—and because pandemic exposure and AI Assistance are correlated across industries, those paths would load onto the post-ChatGPT indicator. Absorbing them requires measuring how hard the pandemic hit each industry’s entry and removing whatever dynamics accompany that exposure. For each industry we therefore compute a continuous dose of pandemic exposure: the entry boom over 2020Q3–2021Q4 relative to a 2018Q4–2019Q4 baseline, from two independent sources—our own Revelio panel and the Census Bureau’s Business Formation Statistics (BFS). Each dose enters as a full set of quarter interactions, allowing industries with greater pandemic exposure to follow a different entry path—one unrestricted over time in the dose dimension—before the AI Assistance coefficient is estimated.¹⁴ Absorbing the Revelio dose leaves the AI Assistance estimate at 0.100; the BFS dose, at 0.105; both jointly, at 0.107—against a baseline of 0.099 estimated on the same sample.¹⁵ If the post-ChatGPT coefficient were capturing recovery from the pandemic boom, conditioning flexibly on the magnitude of that boom should attenuate it. However, it does not: results on both estimation windows, reported in Tables OA7 and OA8, leave the estimate untrimmed. The event-study version tells the same story: the post-2022Q4 path of the AI Assistance coefficients is unchanged whether estimated unconditionally or conditional on either dose (Figure OA6).

A distinct concern is that the outcome itself could inherit pandemic dynamics. Revelio’s coverage depends on professional-platform presence, and platform adoption accelerated during the remote-work era, disproportionately in the knowledge-intensive industries that rank high on AI Assistance. Two features of our design limit this channel. First, as described in Subsection 2.1, a firm’s entry date is elicited from reported founding and employment start dates, not from the timing of its platform presence; late platform adoption shifts which firms we observe, not when we observe them entering. Second, the BFS dose in the preceding analysis is constructed from EIN applications, which exist independently of any platform—the stability of the estimate under that dose is evidence against a coverage artifact timed to the treatment. What remains is selection: firms never present on the platforms are absent from our panel, and Subsection C.2 of the Online Appendix bounds this by benchmarking Revelio against the Business Dynamics Statistics on industry shares, rank orderings, and

¹⁴For completeness, we note that the Census identifies some applications to a two-digit sector but not to any three-digit subsector, and our BFS dose is constructed from named subsectors only, discarding this residual. Column (5) of Table OA7 therefore recomputes the BFS dose allocating each sector’s residual across its subsectors in proportion to their own application counts; the estimate is unchanged at 0.105.

¹⁵The dose specifications run on 79 industries: NAICS 112 and 814 fall outside the BFS Business Applications universe, and we hold the sample constant across all columns, baseline included, so that movement across specifications reflects the dose rather than a sample change.

entry trends. We cannot rule out residual coverage growth correlated with treatment; we note the limitation and rely on the breadth of checks above.

The second layer of the concern remains: AI Assistance could proxy for a pre-existing industry characteristic—a confounder that would have produced a distinctive post-2022 trajectory on its own. We address each candidate confounder with a horse race. Each rival is not a control but a competing treatment: if AI Assistance predicts entry only as a proxy for the rival, the rival’s interaction should carry the post-2022 differential when both enter together. Each candidate is measured continuously at the industry level, standardized, and entered in the specification of Subsection 4.1 both alone and jointly with AI Assistance; Table 6 reports the results. If the post-ChatGPT differential belongs to the confounder, its interaction with Post should survive the joint specification and AI Assistance should collapse. Instead, the opposite occurs for every candidate.

Teleworkability is the first candidate: remote-capable industries were on a distinctive trajectory before and through the pandemic for reasons unrelated to AI, and a measure built from text-intensive entrepreneurial tasks could load on precisely that capability. We construct the Dingel and Neiman (2020) measure at the three-digit industry level and race it against AI Assistance.¹⁶ Alone, teleworkability predicts post-ChatGPT entry (0.069). Jointly, however, it collapses to an insignificant 0.018, while AI Assistance remains at 0.090 (Table 6, Columns (2)–(3)). The separation is not an artifact of weak overlap between the concepts: Dingel and Neiman’s construct captures physical constraints on where work can be performed, whereas AI Assistance captures the cognitive content of the tasks a founder must complete—related in the cross-section, but distinct in what they measure.

IT capital intensity meets a different objection: that AI-compatible industries are simply the already-digital ones, whose post-2022 trajectory owes nothing to ChatGPT. We measure each industry’s IT share of capital from the BEA detailed fixed-asset estimates, fixed at 2018 values. This is a demanding rival—it correlates with AI Assistance at 0.63. Alone, it predicts post-ChatGPT entry (0.068). Jointly, however, it collapses to an insignificant 0.010, while AI Assistance remains at 0.093 (Table 6, Columns (4)–(5)). Pre-existing digitalization predicts post-ChatGPT entry only insofar as it overlaps with AI Assistance; the reverse does not hold.

A steeper pre-existing trajectory is the remaining candidate: AI-exposed industries may simply have been growing faster all along. This test is distinct from the event-study leads of Figure 6: the leads ask whether entry diverged by AI Assistance within the observed pre-treatment window, whereas the race here asks whether AI Assistance proxies for an independently constructed measure of each industry’s secular growth, fitted entirely before

¹⁶We aggregate the occupation-level Dingel–Neiman classification to three-digit NAICS 2022 using OEWS employment weights; construction details are in Subsection A.1.6 of the Online Appendix.

the estimation sample opens. We fit each industry’s trend in firm formation over 2015Q1–2018Q3—a window closing before the estimation sample opens, so the measure and the sample share no observations. Alone, the pre-trend does not predict a post-ChatGPT jump (0.022, insignificant). Jointly, it collapses further to 0.006, while AI Assistance moves only from 0.100 to 0.099 (Table 6, Columns (6)–(7)). The conclusion is unchanged when the pre-trend interacts with every quarter rather than a single post indicator, and when the trend is fitted over 2006Q1–2018Q3, which yields a wrong-signed pre-trend coefficient (Table OA13, Columns (4) and (6)).

No single exercise proves that the post-2022 differential is the effect of AI; the identifying assumption of Subsection 4.1 is not directly testable. Yet taken together, these results sharply constrain what an alternative explanation must accomplish. A confounding characteristic would have to be correlated with AI Assistance yet distinct from teleworkability, IT capital intensity, and prior growth; generate no entry differential across 2018–2022, including through the largest business-formation shock on record; survive flexible conditioning on each industry’s measured pandemic boom and on sector-by-quarter shocks (Subsection A.1.2 of the Online Appendix); and begin moving entry precisely in 2022Q4. Across all of these specifications, the AI Assistance estimate ranges from 0.090 to 0.107. We therefore interpret the post-2022 break as the effect of the commercialization of generative AI.

4.4. Validation of AI Assistance Measure. A key concern with any new measure is whether it captures the underlying concept it is designed to reflect. To validate our AI Assistance Score, we compare it to two widely used measures in the literature. The first one is the O*NET-based measure of LLM exposure (AI Task Exposure) developed by Eloundou et al. (2023) and employed by several papers such as Gmyrek et al. (2023); Auer et al. (2024); Lindebaum and Fleming (2024); Chen et al. (2025); Eisfeldt et al. (2025); Ghosh et al. (2025) among others. This benchmark evaluates the extent to which the tasks performed in each occupation can be automated or augmented by large language models, and is widely accepted as a proxy for how much LLMs can support human work. We aggregate it to the industry level using employment shares, so that it is directly comparable to our AI Assistance Score. The second measure is inspired by Acemoglu et al. (2022) and captures the share of industry level job postings in the Revelio Labs’ data requiring AI-related skills using keywords in job descriptions; we term this measure AI Demand. This benchmark evaluates the industry-level demand for AI skills and, to the extent that skill demand is higher in industries where AI use is posited to be more beneficial, it can proxy for complementarity between an industry’s operations and AI tools.

For the analysis, we fix both measures at pre-ChatGPT levels (2022Q3) and standardize them. We begin with simple correlation analysis between each alternative measure and our AI Assistance measure. The Pearson correlation between AI Assistance and AI Task

Exposure is 0.53 ($p < 0.001$), and 0.32 ($p = 0.004$) with AI Demand. The Spearman rank correlation is 0.43 ($p < 0.001$) with AI Task Exposure and 0.27 ($p = 0.015$) with AI Demand. These results indicate that at the industry level, neither alternative measure perfectly correlates with the AI Assistance variable, highlighting that our variable conveys information not captured by broad industry-level AI compatibility measures.

To systematically determine whether the information conveyed by the AI Assistance measure can be superseded by other measures that broadly describe industry level AI compatibility, we use a horse race design. The idea of this approach is straightforward: we place AI Assistance and an alternative measure of AI–industry relevance side-by-side in the same difference-in-differences specification and evaluate which variable retains explanatory power when both compete to explain the same outcome. The specification is

$$\begin{aligned} \text{IHS(N. New Firms)}_{it} = & \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t \\ & + \beta_2 \text{Alternative Measure}_i \times \text{Post}_t + \mu_i + \mu_t + \epsilon_{it} \end{aligned}$$

where the benchmark is either AI Task Exposure or AI Demand. The specification omits the time-varying controls: the exercise here concerns what the treatment measure captures, not the identification of the treatment effect, and as Column (2) of Table 4 shows, the AI Assistance estimate is insensitive to their inclusion. However, the horse races of Subsection 4.3 include the controls throughout, because there industry conditions are themselves candidate confounds. The key test in this design is whether the coefficient on the competing benchmark measure remains significant once AI Assistance is included.

We report the horse-race results in Table 7. Column (1) reproduces the baseline estimate from the main analysis, while Columns (2) and (3) report the horse race against AI Task Exposure and AI Demand, respectively. Across all specifications, we find that the interaction between AI Assistance and the post-ChatGPT period remains large and statistically significant, whereas the coefficients on both AI Task Exposure (Table 7, Column (2)) and AI Demand (Table 7, Column (3)) are small and statistically insignificant. This pattern indicates that the variation in AI Assistance that predicts firm creation is not subsumed by existing measures of AI exposure or AI adoption.

The horse race design therefore provides evidence that AI Assistance captures an aspect of AI complementarity that is specific to entrepreneurship—namely, the extent to which AI helps founders perform the broad, cross-functional tasks involved in starting and running a new business—rather than characteristics of AI that are relevant primarily for incumbent firms.

While the horse race shows that our AI Assistance Score captures variation in AI complementarity not subsumed by existing measures of industry-level AI exposure, it remains

important to establish that the measure is grounded in how entrepreneurs themselves experience the usefulness of AI tools. To this end, we conducted an original survey of new entrepreneurs in the United States.

The survey was fielded between November 2025 and May 2026 and disseminated through Dynata’s commercial respondent panel¹⁷. We targeted individuals who either started a company or began actively working on a business idea from January 2023 onward, a period coinciding with the widespread availability of generative AI tools. Respondents who did not satisfy these criteria were screened out at the beginning of the survey¹⁸. The final sample comprises 1,036 respondents. The full text of the survey instrument as well as additional results from the survey are reported in Section E of the Online Appendix.

To ensure that respondents could engage meaningfully with the survey, we presented broad task categories that capture the spirit of underlying entrepreneurial activities, rather than asking about narrowly defined tasks that might not resonate with the diversity of businesses represented in our sample. The tasks extracted by our LLM pipeline are necessarily more granular, with a separate AI Assistance estimate for each specific task within each industry. To bring the two measures onto a common scale, we map each LLM-extracted task to the broader survey category it most closely corresponds to. We then aggregate the LLM-generated scores to the task-category level by computing an importance-weighted average of the AI Assistance values across all LLM-extracted tasks mapped to each category. The importance weights are drawn from the LLM pipeline’s task-level Importance Scores, which capture how central a given task is to entrepreneurial work. To construct task category-level survey scores, we restrict the sample to respondents who report having used AI for the task and then average their helpfulness ratings.

Figure 7 plots the LLM-generated AI Assistance Score against the survey-based measure across the 20 task categories. The figure also overlays the fitted OLS regression line and its 95% confidence band. The two measures are positively correlated: the Pearson correlation is 0.571 ($p = 0.008$) and the Spearman rank correlation is 0.614 ($p = 0.004$). These results suggest that the AI Assistance Score captures a meaningful dimension of how entrepreneurs perceive the usefulness of AI tools, lending credibility to its use as a proxy for AI complementarity in entrepreneurial work. Together with the horse race evidence against existing measures, this validation exercise strengthens confidence that our main results reflect genuine variation in how much AI can support entrepreneurial activity across industries.

¹⁷Dynata is a global research data company that maintains an opt-in panel of respondents who are compensated for completing surveys. Respondents who do not satisfy the qualifying criteria are screened out and do not complete the survey.

¹⁸We implement several quality checks to ensure the integrity of responses, including screening for straightlining (i.e., selecting the same response across all items in a battery) and verifying that open-ended responses were substantive rather than random or nonsensical text.

5. POTENTIAL MECHANISMS LINKING AI AND FIRM CREATION

Thus far, we have shown that AI-compatible industries experienced relatively higher levels of firm formation following the release of ChatGPT, suggesting that AI technology has the potential to stimulate entrepreneurial activity. In this section we present three potential mechanisms linking AI availability and firm formation.

5.1. Does AI Reduce the Cost of Experimentation for New Entrepreneurs? A

central feature of entrepreneurship is iterative trial-and-error, encompassing idea validation, product testing, and business model refinement. If AI tools lower the costs associated with these activities, then the positive effect of AI on firm creation should be especially pronounced in industries where experimentation plays a central role. To test this mechanism, we build directly on the dataset introduced in Section 4, which combines measures of firm creation with industry-level indicators of AI Assistance and Experimentation Intensity. This ensures that our analysis of the experimentation channel is based on the same empirical foundation as the main results.

5.1.1. *Entrepreneurship as an Experimental Process.* Scholars increasingly conceptualize entrepreneurship as a process of experimentation, emphasizing the iterative and exploratory nature of creating and developing new ventures. This perspective views entrepreneurship not merely as the pursuit of opportunities, but as a systematic approach to testing hypotheses, learning from outcomes, and adapting strategies accordingly. The lean startup methodology, championed by [Ries \(2011\)](#), advocates for building minimal viable products, conducting experiments to validate business models, and iterating rapidly in response to market reactions. This experimental approach aligns with the notion that entrepreneurial success often hinges on the ability to pivot and refine ideas through continuous testing and learning ([Blank, 2020](#)).

AI has the potential to increase entrepreneurs' productivity in completing experimental tasks related to setting up a firm. This, in turn, reduces the cost of experimentation, making it relatively more appealing to start a firm and leading to higher levels of firm formation.

Our hypothesis is that industries in which AI significantly reduces experimentation costs should experience a greater increase in business formation after ChatGPT's release. To proxy for the benefit accrued from reducing the cost of experimentation, we interact our measure of AI Assistance with the measure of Experimentation Intensity we devised.

5.1.2. *Empirical Strategy for Testing the Experimentation Channel.* We study the relation between the cost of experimentation and firm formation by first sorting industries into groups based on their levels of Experimentation Intensity. Specifically, we define a binary variable equal to one if an industry falls within the top quartile of the distribution of Experimentation Intensity Scores, and zero otherwise. We use a top-quartile indicator because our hypothesis

concerns a specific set of industries: those in which experimentation is so central to venture formation that its cost plausibly constitutes a first-order barrier to entry. A continuous interaction would impose linearity in experimentation intensity, assuming the mechanism operates uniformly across the distribution; instead, we isolate the industries where experimentation is most likely to be a binding constraint and contrast them with the rest. This choice is not driven by the particular cutoff: defining the indicator at the sample mean or median yields very similar results (Subsection A.2).

We start our mechanism analysis by revisiting Equation 6, now interacting AI Assistance with the Experimentation Intense indicator presented above in a DDD design. At the industry-by-quarter level, the main specification is:

(8)

$$\begin{aligned} \text{IHS(N. New Firms)}_{it} = & \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Experimentation Intense}_i \times \text{Post}_t \\ & + \beta_2 \text{AI Assistance}_i \times \text{Post}_t + \beta_3 \text{Experimentation Intense}_i \times \text{Post}_t \\ & + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it} \end{aligned}$$

where the dependent variable $\text{IHS(N. New Firms)}_{it}$ is the inverse hyperbolic sine of the total number of new firms. AI Assistance_i is the industry level proxy measuring LLMs' ability to support entrepreneurs through the completion of various tasks, which we standardize for ease of interpretation. $\text{Experimentation Intense}_i$ is a binary variable equal to one if an industry falls within the top quartile of the distribution of Experimentation Intensity Scores, and zero otherwise. Post_t is an indicator switching to 1 from 2022Q4 forward—periods during which ChatGPT is publicly available. μ_i are industry fixed effects used to control for industries' intrinsically higher levels of firm creation, which may not be related to the arrival of ChatGPT. In this specification we replace quarter fixed effects (μ_t) with sector-by-quarter fixed effects (μ_{It}), where sector is defined at the two-digit NAICS level. We introduce a richer fixed effects structure to ensure that our results are not driven by sector-wide trends that may correlate with Experimentation Intensity (e.g., sectoral adoption of AI tools, recovery dynamics post COVID or broad sector policy changes). This allows us to isolate the variation across three-digit industries within a given sector and quarter, thereby ensuring that our estimates are not confounded by broader sectoral dynamics and reflect genuine heterogeneity in how AI interacts with the cost of experimentation. In particular, we are concerned about sector-specific trends that could disproportionately affect experimentation-intensive industries. For example, some sectors—such as Information (NAICS 51) and Professional Services (NAICS 54)—may have seen a wave of AI-related innovation, investment, and entrepreneurial interest after the release of ChatGPT. These sectors also tend to include experimentation-heavy sub-industries (e.g., digital media, marketing, consulting), where trial-and-error is critical to new firm development. If these sector-wide surges in firm creation are not controlled for, we risk

mistakenly attributing them to the interaction between AI Assistance and Experimentation Intensity, when in fact they reflect a general sectoral response to the AI shock.

In our specification, the triple interaction AI Assistance $\times$ Experimentation Intense $\times$ Post implements a triple-difference design. Intuitively, the estimator compares the post-ChatGPT increase in firm creation for AI-compatible industries relative to less AI-compatible industries, and then asks whether this differential effect is itself larger in industries where experimentation plays a more central role. In other words, the DDD isolates the *additional* post-ChatGPT increase in entry that arises only when an industry is both (i) highly AI-assistable and (ii) highly experimentation-intensive. By netting out the DiD effects associated with AI compatibility alone and experimentation intensity alone, the triple interaction coefficient β_1 captures the incremental effect attributable specifically to a reduction in experimentation costs. This allows us to test directly whether industries that rely more heavily on trial-and-error benefit disproportionately from the arrival of AI tools.

We also revise the dynamic specification to tailor it to the mechanism we are considering so to visually inspect parallel trends. The event study specification is

$$\begin{aligned}
 (9) \quad \text{IHS(N. New Firms)}_{it} = & \beta_0 + \sum_{t \neq 2022q3} \beta_t \text{AI Assistance}_i \times \text{Experimentation Intense}_i \times 1\{D = t\}_t \\
 & + \sum_{t \neq 2022q3} \gamma_t \text{AI Assistance}_i \times 1\{D = t\}_t \\
 & + \sum_{t \neq 2022q3} \delta_t \text{Experimentation Intense}_i \times 1\{D = t\}_t \\
 & + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it}
 \end{aligned}$$

Here, we are interested in studying the sign and significance of the coefficients on the triple interaction (β_t), as they inform us as to whether different industries differ in the joint effect of AI Assistance and Experimentation Intensity prior to the shock.

5.1.3. Results: AI Assistance in Experimentation Intense Industries. As a first step in investigating the mechanism, we examine the coefficients from the dynamic regression in Equation 9. Analyzing the significance of the event study coefficients allows us to assess whether industries systematically differed in business formation based on their AI compatibility and reliance on experimentation prior to the arrival of ChatGPT. As shown in the event study plots (Figure 8), the coefficients for the lead periods are statistically insignificant and relatively stable in magnitude, suggesting that pre-existing cross-industry differences are unlikely to drive the results, and that there is no evidence of anticipatory effects. The effects on the lags are positive and incremental in size over quarters, likely reflecting differential rates of

take-up over time, as is expected for complex technologies (Figure 8). The findings are unaffected when including the vector controls to account for varying financial and economic conditions (Figure 8, Panel (b)).

Next, we test for existence of a dose-response effect using the compact regression in Equation 8. Industries where the benefits from reducing the cost of experimentation are higher—industries where experimentation is prominent and AI can strongly assist with task completion—experience a statistically significant increase in firm creation (Table 8, Column (1)). Moreover, neither the coefficient on $\text{AI Assistance} \times \text{Post}$ nor that on $\text{Experimentation Intense} \times \text{Post}$ is individually significant, indicating that the post-ChatGPT increase in firm creation is concentrated in industries that are jointly AI-assistable and experimentation-intensive, rather than operating through either dimension alone. It is important to acknowledge that AI may still influence firm creation through mechanisms that operate uniformly within broad sectors. However, because the specification includes sector-by-quarter fixed effects, such common sector-wide effects are absorbed and cannot be separately identified in our analysis.

The results in the analysis display economically significant effects: one standard deviation increase in the AI Assistance measure leads to about 3,500 new firms being created in high Experimentation Intensity industries, which equates to roughly a 25%¹⁹ increase from the pre-treatment mean.

To the extent to which our battery of industry-by-quarter varying financial and economic controls can proxy for time varying within industry economic and financial conditions, the results are not driven by that. On the contrary, the inclusion of the control set slightly increases the magnitude of the effect of the mechanism (Table 8, Column (2)).

In Section B.2 of the Online Appendix we show that our results are unchanged when using $\text{Log}(\text{N. New Firms})$ in place of IHS, assuaging concerns regarding the results being an artifact of the IHS specification. We also test the robustness of our results by providing different binarized version of the Experimentation Intense indicator. In Subsection A.2 of the Online Appendix we categorize industries as being Experimentation Intense using as cutoff either the sample mean or median. The results remain qualitatively unchanged. In the same subsection of the Online Appendix, we further report the fully continuous triple-interaction specification: the triple-interaction coefficient is positive and consistent in sign with the threshold result, though imprecisely estimated given the limited independent variation in the product of two continuous measures across 81 industries.

Survey evidence provides further support for this interpretation. In addition to asking entrepreneurs whether they use AI tools, our survey asked AI users to describe, in their own

¹⁹This does not mean the mechanism regression contradicts the baseline results of 11%; rather, it shows that the baseline average of 1,500 firms aggregates both low- and high-experimentation industries, and once we isolate the channel where AI is most relevant, the impact is amplified. In other words, the mechanism identifies the subset of industries where AI’s ability to lower experimentation costs is most potent.

words, how they use AI in their business. We classify these open-ended responses using a large language model, allowing each response to be assigned to one or more substantive use categories. The resulting patterns are consistent with the experimentation-cost mechanism. The most frequently mentioned use of AI is ideation and brainstorming, reported by nearly half of AI users, followed by operations and administration, market research, marketing and content, and information retrieval (Figure 9). These activities closely correspond to the types of tasks through which entrepreneurs experiment: generating and refining business ideas, learning about customers and competitors, testing messaging and content, gathering information, and iterating over operational choices. Thus, the survey evidence indicates that entrepreneurs are not merely using AI for generic productivity gains; rather, they report using it precisely in domains where trial-and-error, feedback, and rapid iteration are central to the venture formation process. This reinforces the interpretation of the triple-difference results as evidence that AI stimulates firm creation by lowering the cost of entrepreneurial experimentation.

Overall, the evidence presented in this section supports the interpretation that the rise in firm creation following the release of ChatGPT is concentrated in industries where AI can meaningfully reduce the cost of experimentation. These findings are not explained by pre-trends, sector-wide shocks, or time-varying financial and economic conditions, reinforcing the role of reduced experimentation costs as a key mechanism.

5.2. AI as a Catalyst of Private Financing. Beyond reducing the costs of experimentation, AI may also influence firm creation by shaping how capital is allocated. If investors perceive AI as enhancing founders’ productivity or mitigating early-stage risks, industries more compatible with AI should attract relatively more private financing. Such reallocation of capital could further amplify entry in AI-compatible sectors. To test this mechanism, we rely on the same dataset described in Section 4, which integrates firm creation and capital allocation data with industry-level measures of AI Assistance.

5.2.1. Do Entrepreneurs Cater to Investors? Industries inherently more compatible with AI should experience productivity gains after the deployment of the technology, consequently attracting more capital, due to higher expected returns. The heightened flow of capital into highly AI-exposed industries entices entrepreneurs, as it increases their perceived probability of success. More specifically, entrepreneurs cater to investors by starting companies in industries where financing is abundant. Therefore, an alternative explanation as to why we see increased firm creation in AI-compatible industries—following the arrival of ChatGPT—is that industries where AI is highly beneficial for entrepreneurs experience an increase in funding by investors, which encourages prospective entrepreneurs to start companies in that same industry.

5.2.2. *Mediation Strategy to Test the Capital Allocation Channel.* We investigate the validity of this potential mechanism using mediation analysis. First, we estimate the Total Impact of AI-compatibility on firm creation:

$$(10) \quad \text{IHS}(\text{N. New Firms})_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \mu_i + \mu_{It} + \epsilon_{it}$$

where the dependent variable $\text{IHS}(\text{N. New Firms})_{it}$ is the inverse hyperbolic sine of the total number of new firms. AI Assistance_i is the industry level proxy measuring LLMs' ability to support entrepreneurs through the completion of various tasks, which we standardize for ease of interpretation. Post_t is an indicator switching to 1 from 2022Q4 forward—periods during which ChatGPT is publicly available. μ_i are industry fixed effects used to control for industries' intrinsically higher levels of firm creation, which may not be related to the arrival of ChatGPT. μ_{It} are sector-by-quarter²⁰ fixed effects, which allow us to net out time trends that are common to macro industry groups and which may shape business formation. If β_1 is significant, it supports the idea that firm creation increases more in AI-compatible industries after the release of ChatGPT.

Next, we check whether AI compatibility results in higher levels of capital allocated post-ChatGPT:

$$(11) \quad \text{IHS}(\text{Investments})_{it} = \gamma_0 + \gamma_1 \text{AI Assistance}_i \times \text{Post}_t + \mu_i + \mu_{It} + \epsilon_{it}$$

where the dependent variable $\text{IHS}(\text{Investments})_{it}$ is the inverse hyperbolic sine of the total amount of VCPE investments in millions of dollars²¹. AI Assistance_i is the industry level proxy measuring LLMs' ability to support entrepreneurs through the completion of various tasks, which we standardize for ease of interpretation. Post_t is an indicator switching to 1 from 2022Q4 forward—periods during which ChatGPT is publicly available. μ_i are industry fixed effects used to control for industries' intrinsically different levels of allocated capital, which may not be related to the arrival of ChatGPT. μ_{It} are sector-by-quarter fixed effects, which allow us to net out time trends that are common to macro industry groups and which may shape business formation. The sign and significance of γ_1 will determine whether funding availability acts as a conduit for the effect of the popularization of AI technologies on firm formation.

²⁰This specification differs from Equation 6 in Section 4, as it is more stringent in terms of fixed effects structure. We include sector-by-quarter fixed effects to keep the regression specification consistent with the test presented in Subsection 5.1.

²¹It is important to notice that while the total capital allocated to a given industry is indeed an equilibrium outcome, which could be driven by both demand and supply factors, the tests are not aimed at understanding which side is responding, rather their objective is to elicit the presence of an overall capital shift to the arrival of AI that might coincide with changes in the level of firm creation and potentially act as the main driver of increased firm creation.

Finally, we estimate the Mediated Effect of the treatment after controlling for investment flows. We include Investment in the firm creation regression as a mediator:

(12)

$$\text{IHS}(\text{N. New Firms})_{it} = \delta_0 + \delta_1 \text{AI Assistance}_i \times \text{Post}_t + \delta_2 \text{IHS}(\text{Investments})_{it} + \mu_i + \mu_{It} + \epsilon_{it}$$

where all the variables are measured as usual. δ_1 measures how much AI compatibility directly affects firm creation, controlling for investment (Direct Effect). δ_2 is the residual explanatory power of capital supply after conditioning on the treatment (AI Assistance $\times$ Post).

5.2.3. Results: Investment Flows and Firm Creation. As a preliminary step for the mediation analysis, we examine the event study for the total amount of investments (Figure 10). The coefficients on the leads are all insignificant with no clear pattern arising, suggesting two things. First, more AI compatible industries were not experiencing differential levels of investments prior to the arrival of ChatGPT. Second, investors did not anticipate the productivity-boosting effect of the technology and therefore did not rush to fund AI-compatible industries in advance of its arrival.

We present the results of the mediation analysis in Table 9. The treatment displays a positive and significant total impact: industries that are more compatible with LLMs display relatively higher firm creation following the arrival of ChatGPT (Column (1)). On the other hand, when we look at the association between the treatment and the total amount invested by VCPE funds—the mediating variable—we find that the coefficient is negative and significant, suggesting that the arrival of AI coincides with lower levels of capital deployed in AI compatible industries (Table 9, Column (2)). This finding is not consistent with the hypothesis that the effect of an increase in access to AI impacts firm creation via a capital allocation channel, as we actually observe reallocation away from AI compatible industries. In Table 9, Column (3), we examine the direct impact of the treatment, now mediating for investment levels. We find that the effect of the release of ChatGPT remains unchanged after including the mediator, which itself is statistically insignificant, suggesting that the effect of the treatment is not flowing through a change in capital allocation. These findings also help address the concern that the main effects (Table 4) might be driven by changes in the investment opportunity set rather than by AI lowering the cost of entrepreneurial experimentation. If the release of ChatGPT had primarily expanded profitable investment opportunities—either directly, by creating new AI-related markets, or indirectly, by increasing incumbents’ productivity and stimulating downstream or upstream demand—we would expect to observe a corresponding rise in aggregate investment following the diffusion of AI. The results in this section show no evidence that investment rose more in high-AI-Assistance

industries relative to less AI-compatible industries in the post-ChatGPT period. The absence of an investment surge suggests that the overall pool of available capital remains stable, and that the observed increase in firm formation arises from entrepreneurs’ enhanced ability to operate and experiment with existing resources rather than from a broad expansion in profitable opportunities. In this sense, the evidence is more consistent with AI reducing the fixed costs of entry and facilitating experimentation, rather than shifting the investment opportunity set.

All in all, given the current evidence at hand, it is unlikely that firm formation is driven by entrepreneurs catering to investors, as we do not find evidence that firm creation increases in AI compatible industries following the arrival of ChatGPT because investors anticipate productivity growth and flood the industries with funds. It is important to highlight that our findings do not suggest that relatively less capital was allocated to AI industries following the arrival of ChatGPT. Rather, our treatment group is defined as the set of industries where entrepreneurs would benefit the most from using LLMs to carry out the tasks involved in establishing and running a firm. Indeed, private investment in AI-related companies has increased substantially in recent years (Maslej et al., 2025). Hence, these findings further confirm that our measure of AI compatibility does not perfectly overlap with the traditional group of AI and AI-related industries. This distinction ensures that our design, which exploits the arrival of ChatGPT as a quasi-exogenous shock, does not merely capture a technological shock confined to technology firms.

While the use of IHS is innocuous for firm creation, since the dependent variable is strictly positive, the critique of Chen and Roth (2024) has sharper relevance in the context of investment flows. In some industry–quarter cells, total VCPE investment is zero, which means that conventional log specifications are undefined and IHS implicitly combines extensive and intensive margins. To address this concern, in Subsection B.3 we implement a battery of robustness tests. In particular, we (i) replicate the specifications using logs whenever feasible, (ii) re-estimate the investment regressions using Poisson pseudo-maximum likelihood (PPML), which is transformation-free and well defined in the presence of zeros, and (iii) explicitly separate investments into extensive and intensive margins in a two-part specification. Across all approaches, the coefficient on *AI Assistance* $\times$ *Post* remains stable and the mediator continues to be insignificant. This demonstrates that our conclusion—that capital reallocation does not explain the rise in firm creation in AI-compatible industries—is not an artifact of the IHS functional form.

Lastly, we acknowledge that the relation between capital allocation and AI compatibility is not ex-ante obvious. On the one hand, higher expected productivity may attract greater investment into AI-compatible industries. On the other hand, firms in these industries may require less external financing if AI reduces the cost of key inputs, in which case average

deal sizes would decline rather than increase. In Subsection A.3 of the Online Appendix we test this alternative channel by replacing total capital with the average size of investment rounds as a mediator. We find little evidence in support of this mechanism: while average deal sizes are somewhat smaller in AI-compatible industries after ChatGPT, these changes do not predict entry, and the treatment effect on firm creation remains unchanged. Overall, capital does not appear to be reallocated in AI-compatible industries, and it is therefore unlikely that financing dynamics are the primary driver of the increase in firm creation.

5.3. Does AI Bridge Human-Capital Gaps? A further channel through which AI may affect entrepreneurship is by alleviating human-capital constraints. Many potential founders possess incomplete skill sets relative to the demands of operating a business in their chosen industry. If AI tools complement individual abilities and mitigate these deficiencies, then industries where potential founders have larger skill gaps should see a disproportionate boost in entry following the arrival of ChatGPT. To test this hypothesis, we construct a novel individual-level dataset from Revelio Labs, which tracks employment histories and job titles of U.S. workers. By linking these histories to standardized task and skill data from O*NET, we build dynamic skill profiles for each worker and compare them against industry-specific requirements. This allows us to quantify skill gaps and examine whether AI reduces the barriers to firm creation posed by human-capital limitations.

5.3.1. Missing Skills as a Barrier to Entrepreneurial Entry. A lack of skills is a legitimate barrier that prevents potential entrepreneurs from entering certain industries. AI—particularly LLM-powered virtual assistants—can supplement the skill sets of entrepreneurs by allowing them to complete tasks they previously could not perform. Individuals whose skill set is better complemented by AI see a reduction in barriers to entry, and should be more likely to sort into self-employment. Therefore, a potential channel linking AI availability to firm creation is LLMs’ ability to bridge human-capital gaps. This channel is distinct from a reduction in the cost of experimentation: the latter applies to tasks that individuals can already perform but can now complete more efficiently thanks to the presence of LLMs. Here, AI supplements the entrepreneur’s abilities, enabling them to perform tasks for which they previously lacked the skills. AI, in this sense, fills skill gaps, letting entrepreneurs operate businesses they otherwise wouldn’t be able to run.

An implication of this mechanism is that individuals whose skill sets are better complemented by AI should sort themselves into unfamiliar industries at higher rates following the arrival of ChatGPT.

If following the arrival of ChatGPT, entrepreneurs become more likely to start firms in industries for which they do not have the appropriate skill set, but these are also industries where ChatGPT can help bridge the skill gap to start a venture, then firm formation is likely fueled by LLMs’ ability to bridge skill gaps.

5.3.2. *Empirical Design for Measuring Skill Gaps and Entry.* We estimate the regression at the Individual-by-industry-by-quarter level (jit)

$$(13) \quad P(\text{Entrepreneur})_{jit} = \beta_0 + \beta_1 \text{Skill Gap}_{jit} \times \text{Post}_t + \beta_2 \text{Skill Gap}_{jit} + \mu_{ji} + \mu_{jt} + \mu_{it} + \epsilon_{jit}$$

$P(\text{Entrepreneur})_{jit}$ is an indicator for whether individual j started a firm in industry i in quarter t , multiplied by 10,000 for ease of exposition. Skill Gap_{jit} is the skill gap value for individual j with respect to industry i at quarter t . Post_t is an indicator for the post-ChatGPT period (2022Q4 onward). μ_{ji} are industry by individual fixed effects, which control for any time-invariant preference or sorting between an individual and an industry; for example, if an individual is more likely to start a business in a given industry due to prior experience or interest, this FE absorbs that. μ_{jt} are individual by quarter fixed effects, which control for individual-level shocks or trends over time, such as job loss, career shifts, or general changes in propensity to become an entrepreneur. Finally, μ_{it} are industry by quarter fixed effects, which control for all industry-specific shocks over time, such as business cycle exposure, regulatory changes, or industry-specific AI trends.

The coefficient β_2 reflects the importance of a skill gap for an individual’s ability to penetrate a given industry, while β_1 represents the differential effect of the existence of a skill gap in penetrating a given industry before and after the arrival of ChatGPT. When β_1 is positive, entrepreneurs sort themselves into harder to penetrate industries at higher rates after the arrival of ChatGPT. Moreover, when β_2 is negative, we know that a given industry is unconditionally harder to penetrate given the individual’s skill set. Therefore, if β_1 is positive and β_2 is negative, we can conclude that entrepreneurs sort themselves into industries where ChatGPT is able to provide support to bridge the skill gap and said industries are generally harder to break into for entrepreneurs without the right set of skills, implying that ChatGPT eases the barrier to entry due to skill gaps.

5.3.3. *Results: Skill Gaps and Entry into Entrepreneurship.* In Table 10 we present the results of the skill gap mechanism test. The sample used includes individuals present on RL who were actively working between 2020q4 and 2024q1; for computational ease we draw a random sample of 200,000 individuals to use for the analysis. First, we assess the reliability of our skill gap measure (Table 10, Column (1)) and find that our proxy for skill mismatch predicts whether an individual will become an entrepreneur: the coefficient is negative and significant: a one standard deviation increase in the skill gap measure—i.e., moving away from the optimal skill benchmark—leads to a decline of 0.02265 percentage points²² in the likelihood of starting a firm in a specific industry-quarter, which equates to a 3% reduction relative to the pre-treatment entrepreneurial flow (0.755%).

²²The binary outcome is multiplied by 10,000 for readability, so the reported coefficient of -2.265 corresponds to a change of 0.02265 percentage points in the entry probability.

Next, we estimate Equation 13 to test the mechanism (Table 10, Column (2)). The skill-gap coefficient is negative and significant, whereas the interaction with the post-ChatGPT indicator is positive but statistically insignificant. Although the signs align with the hypothesized mechanism, only the main effect (β_2) is statistically significant, reiterating that entrepreneurs tend to enter industries that better match their skills. On the other hand, the coefficient on the interaction between the skill gap and the post variable is insignificant, although it is positive, consistent with the mechanism’s prediction.

We next restrict the sample to individuals who eventually become entrepreneurs during the sample period and re-estimate Equation 13. This exercise isolates variation in where entrepreneurs choose to enter—holding fixed the decision to become a founder—and allows us to test whether the availability of AI tools shifts entrepreneurs toward industries that are farther from their prior skill set. First, the skill mismatch variable accurately predicts the industry in which the individual will seek to start a firm: a higher level of skill misalignment with respect to a given industry’s requirement correlates negatively with the likelihood of becoming an entrepreneur in the same industry (Table 11, Column (1)). On the other hand, the interaction captures whether post-ChatGPT entrants are relatively more likely to start firms in industries in which they initially lacked traditional skills. The results, reported in Table 11, Column (2), show that the interaction coefficient remains positive but statistically insignificant. This suggests that, conditional on deciding to become an entrepreneur, the introduction of generative AI has not yet led to a meaningful reallocation of founders across industries according to their prior skill gaps.

Overall, the combined evidence in Table 10 and 11 suggests that while AI might modestly ease entry into industries requiring unfamiliar skills, the effect is not strong enough to be statistically detected. At the same time, the strong negative relationship between the skill-gap measure and entry prior to the AI shock supports the validity of the measure as a proxy for human-capital mismatch across industries. Hence, AI’s ability to bridge such mismatches appears limited during our sample period and is unlikely to be a primary driver of the post-ChatGPT rise in firm creation.

6. WHO ARE THE NEW ENTREPRENEURS? EVIDENCE ON POST-AI FIRM QUALITY

As a final step in our analysis, we examine whether firms created after the release of ChatGPT differ intrinsically from comparable firms established before its arrival.

A central implication of our mechanism is that AI lowers the cost of experimentation, thereby reducing barriers to entry for potential entrepreneurs. Canonical economic models predict that a reduction in entry barriers, while raising entry, dilutes average firm quality, as the influx of new entrants draws in marginal, lower-ability founders (Lucas Jr, 1978; Hopenhayn, 1992). If lower entry costs enable the formation of ventures that previously

would not have been viable, the post-AI increase in firm formation might be accompanied by a composition shift toward lower-quality firms: a “lemon market” for new ventures. However, this need not always be the case. In an alternative framework, the effect of lower entry barriers on average firm quality depends on how these barriers are reduced. If AI primarily lowers fixed entry costs—a uniform reduction that benefits all potential founders equally—then canonical predictions hold and average firm quality declines as marginal, low-ability entrepreneurs enter. By contrast, if AI disproportionately reduces the convex costs of experimentation and scaling, the gains from AI are *ability-complementary*: higher-ability individuals benefit more because they engage in more experimentation and operate at larger scale. In this case, the post-AI payoff gain is increasing and convex in entrepreneurial ability, inducing both higher entry rates and a shift in the composition of entrants toward higher-ability types. Consequently, firm formation rises while average firm quality improves.

Guided by this distinction, we turn to the data to test whether the new firms spawned in the wake of ChatGPT differ in quality from their predecessors. To do so, we construct a firm-level dataset from Revelio Labs that tracks incorporation dates, industry affiliations, survival, employment growth, and average worker education for firms established between 2020Q4 and 2024Q1, with outcomes observed through 2025Q1²³. By linking each firm to its industry’s AI Assistance Score, we can directly compare proxies of firm quality for post-AI entrants with those of earlier cohorts.

6.1. Empirical Strategy. We assess whether AI-driven firm creation results in lower-quality entrants or in the formation of higher-quality ventures by examining three firm-level post-entry outcomes that capture complementary dimensions of venture quality.

First, *firm survival* reflects basic viability: higher-quality firms are more likely to endure beyond the initial quarters of operation, whereas lower-quality entrants tend to exit quickly. Second, *employment growth* proxies for a firm’s ability to expand and scale, which depends on managerial capability, resource acquisition, and market success—key manifestations of operational quality. Third, the *average education level of employees* captures the quality of the workforce that new firms are able to attract. Because firm and worker quality tend to be positively assortatively matched, the education level of employees provides an informative proxy for underlying firm quality: if newly founded firms in highly AI-compatible industries attract more educated workers, they are likely to be of higher quality.

Because firm quality is not directly observable, we examine multiple complementary dimensions—viability, growth capacity, and human-capital composition—to provide a broader picture of how the post-AI surge in firm creation relates to the quality of new entrants.

²³We refresh the workforce data up to 2025Q1, so that we are able to calculate proxies of quality for firms registered in most recent years.

Our empirical strategy uses a difference-in-differences framework that leverages cross-sectional variation in industry AI Assistance and temporal variation around the release of ChatGPT in 2022Q4. Specifically, we estimate the following cross-sectional firm-level specification:

$$(14) \quad Y_f^{(\tau)} = \beta_0 + \beta_1 \text{AI Assistance}_{i(f)} \times \text{Post}_{t(f)} + \mu_{i(f)} + \mu_{t(f)} + \mu_{s(f)} + \epsilon_f$$

$Y_f^{(\tau)}$ denotes one of three outcomes observed τ quarters after registration: an indicator for survival ($P(\text{Survived})_f^{(\tau)}$), the proportional change in employment since establishment ($\text{Emp. Growth}_f^{(\tau)}$), or the average years of education among the firm’s active workforce ($\text{Avg. Education}_f^{(\tau)}$), with $\tau \in \{4, 6, 8\}$. $\text{AI Assistance}_{i(f)}$ is the standardized AI Assistance Score of firm f ’s industry, and $\text{Post}_{t(f)}$ is an indicator for registration after the release of ChatGPT. We include industry ($\mu_{i(f)}$), registration quarter ($\mu_{t(f)}$), and state ($\mu_{s(f)}$) fixed effects to account for unobserved heterogeneity.

6.2. Results and Interpretation. Table 12 shows that firms in highly AI-compatible industries founded after the release of ChatGPT exhibit significantly higher survival probabilities across all three time horizons. The interaction coefficient β_1 is positive and statistically significant in all specifications, with point estimates ranging from 0.003 to 0.004. The estimated effects on firm survival are modest in absolute terms but economically meaningful. At one year, a one-standard deviation increase in AI Assistance raises survival by 0.3 percentage points, equivalent to a 0.36 percent relative increase over the pre-ChatGPT baseline survival rate of roughly 84%. At 18 months, the effect rises to 0.4 percentage points (0.52% relative to a 77% baseline), and at two years it remains at 0.3 percentage points (0.41% relative to a 74% baseline).

While the magnitudes appear smaller at first sight, they are well within the range of effects documented in the literature. For instance, Myers et al. (2025) finds that receipt of an SBIR grant raises annual survival probabilities by about 1.4 percentage points on a baseline survival rate of 80%, whereas Schwartz (2011) shows that business incubators often have no effect (and in some cases a negative effect) on long-term survival. However, it is important to note that the SBIR study examines a highly idiosyncratic sample, as the treated firms are those able to secure a competitive R&D grant, while our estimates reflect the broad effect of AI tools across entire industries.

Turning to employment dynamics, Table 13 reports the effect of AI compatibility on post-entry size. We find that firms founded after the release of ChatGPT in industries with higher AI Assistance expand more rapidly in their early years. At one year, a one-standard deviation increase in AI Assistance is associated with 0.7 percentage points higher employment growth, which corresponds to a 4.7% increase relative to the pre-ChatGPT mean growth rate of 15%.

At 18 months, the effect is 0.9 percentage points, or 5.6% relative to the baseline mean of 16%. At two years the estimated coefficient is smaller: 0.6 percentage points, which implies a 3% increase relative to the baseline mean of 20%. This pattern suggests that the positive impact of AI on employment growth is mostly concentrated in the early stages of firm development.

In addition, we examine whether post-AI firms differ in the level of human capital they attract, conditional on the level of AI compatibility (Table 14). The analysis shows that firms founded after the arrival of ChatGPT in industries with higher AI Assistance tend to have a more educated workforce. At one year, a one-standard-deviation increase in AI Assistance is associated with an additional 0.796 years of educational attainment among the active workforce, a 5% increase relative to the pre-ChatGPT baseline mean of 15.914. At 18 months, the effect is 0.812 years (5.1% relative to the baseline mean of 15.954), and at two years the estimated coefficient remains positive though somewhat smaller at 0.719 years (a 4.5% increase relative to the baseline mean of 15.975). This pattern suggests that the positive impact of AI on firm creation is concentrated among entrepreneurs and ventures able to attract higher-quality human capital.

Taken together, these results are consistent with the alternative framework in which AI reduces the convex costs of experimentation and early-stage scaling, rather than merely lowering fixed entry costs. In such a setting, AI acts as an ability-complementary technology: it enhances the effectiveness of skilled founders in refining products, reaching customers, and organizing production, thereby improving the survival and early growth prospects of new firms. The finding that post-ChatGPT entrants in AI-compatible industries both survive longer and expand faster suggests that AI tools do not simply enable more entry, but also help new ventures execute more effectively once they are formed. In other words, generative AI appears to lower the cost of experimentation in a way that raises—not dilutes—the average quality of entrepreneurial outcomes.

6.3. Rationalizing the Quality Results: A Theoretical Framework. The empirical findings presented in the preceding subsections reveal a pattern that stands in tension with canonical models of industry dynamics. In the frameworks of Lucas Jr (1978) and Hopenhayn (1992), a reduction in entry barriers leads to an influx of marginal entrepreneurs, diluting the average quality of the entrant pool. Yet we observe the opposite: following the release of ChatGPT, industries with greater AI compatibility exhibit not only higher entry rates but also higher survival, faster growth, and a more educated workforce among new entrants. In this subsection, we sketch the key ingredients of a theoretical framework that rationalizes these findings.

The resolution hinges on two features that canonical models typically abstract from. First, the cost reduction brought by AI is *asymmetric*: it disproportionately benefits high-capability entrepreneurs. Second, the population of potential entrepreneurs is subject to

idiosyncratic entry frictions—capturing factors such as heterogeneous access to early-stage financing, variation in the degree of uncertainty surrounding entrepreneurial outcomes, and differences in the availability of experimentation and testing strategies—that are orthogonal to entrepreneurial capability. Together, these two features allow the model to generate both higher entry and higher average quality simultaneously.

To see the intuition, consider an economy populated by two types of individuals who differ in entrepreneurial capability, which we define as the ability to start a firm at lower fixed cost. In each period, individuals choose between salaried employment and entrepreneurship. In addition to the capability dimension, each individual faces an idiosyncratic friction that governs their occupational choice. Crucially, this friction is independent of capability: some high-capability individuals remain in salaried positions because they face unfavorable draws of the entry friction, while some low-capability individuals become entrepreneurs because their friction draws are favorable. This separation between capability and the entry decision is central to the model’s predictions.

When AI arrives, it reduces the fixed cost of starting a firm, but only for high-capability individuals.²⁴ This asymmetric cost reduction has two consequences. First, the gap between the entrepreneurial payoff and the salaried wage widens more for high-capability individuals than for low-capability ones, drawing a disproportionate share of high-capability individuals into the entrepreneur pool. Second, the resulting increase in labor demand and reduction in labor supply push up the equilibrium wage. The higher wage, in turn, raises the opportunity cost of entrepreneurship for all individuals, but since low-capability individuals did not benefit from the cost reduction, some of them are pushed out of the entrepreneur pool. The net effect is that both the total mass of entrepreneurs and the share of high-capability individuals among them increase.

Two formal propositions emerge from the model. The first establishes that the share of high-capability entrepreneurs in the total entrepreneur pool rises upon the arrival of AI (*higher entrepreneur quality*). The second establishes that the total mass of entrepreneurs increases (*higher entry rate*). The proof of the latter proceeds by contradiction: if the mass of entrepreneurs were to decline, labor demand would fall and entry costs would decrease (due to fewer, but more efficient, firms), generating excess labor supply that would push wages down and output up. But lower wages and higher output would increase entrepreneurial profits, strengthening the incentive to enter—contradicting the premise that entry declined.

²⁴This asymmetric fixed-cost reduction is a reduced-form representation of the ability-complementary channel. The model’s comparative statics require only that AI raise the entrepreneurial payoff more for high-ability than for low-ability individuals; whether this asymmetry stems from a lower fixed entry cost that only capable founders can exploit, or from lower convex costs of experimentation and scaling—which higher-ability founders incur more of, and therefore save more on—is immaterial for the predictions that follow. We adopt the fixed-cost formulation for tractability.

The role of the idiosyncratic entry friction deserves emphasis, as it is the feature that distinguishes the model’s predictions from canonical results. In the absence of this friction, high-capability individuals would *always* enter the entrepreneur pool before low-capability ones. In such a world, any reduction in entry costs would operate exclusively at the margin, drawing in lower-capability individuals and diluting average quality—precisely the prediction of Lucas Jr (1978) and Hopenhayn (1992). By contrast, with the friction in place, some high-capability individuals are “stuck” in salaried employment despite having the ability to run successful ventures. The arrival of AI unlocks entry for these individuals by widening the gap between entrepreneurial and salaried payoffs sufficiently to overcome their unfavorable friction draws. It is this *inframarginal* reallocation of high-capability individuals—rather than marginal entry of low-capability ones—that drives the improvement in average quality.

This mechanism maps directly onto our empirical findings. The model’s prediction that AI raises both entry and quality is consistent with the higher survival rates, faster employment growth, and more educated workforces documented in Tables 12–14. Moreover, the model’s emphasis on the cost of experimentation as the channel through which AI operates aligns with the evidence from Section 5 that the increase in firm creation is concentrated in industries where experimentation is most salient. The framework thus provides a unified theoretical account of the paper’s main results: AI lowers the cost of experimentation in a way that disproportionately benefits high-ability entrepreneurs, leading to both greater and higher-quality firm formation.

The full model, including the formal environment, equilibrium definition, and proofs of both propositions, is presented in Section G of the Online Appendix. While the current setup is deliberately parsimonious, it can be naturally extended in future work to explicitly incorporate the cost of experimentation as a modeled primitive, or to embed elements of Bayesian entrepreneurship in which individuals draw signals to learn about the quality of their ventures prior to entry.

7. CONCLUSION

This paper examines how the commercialization and widespread accessibility of artificial intelligence—exemplified by the release of ChatGPT in late 2022—have influenced entrepreneurial activity in the United States. While prior research has explored how AI reshapes labor markets and firm-level productivity, we focus on a previously underexplored margin: new firm creation. Specifically, we ask whether AI tools can lower barriers to entry for prospective entrepreneurs by supporting early-stage business formation tasks.

To answer this question, we construct novel industry-level measures of AI compatibility and experimentation intensity using a multi-step pipeline built on large language models. These measures are designed to capture (i) how useful AI is in assisting entrepreneurs with

task completion in a given industry (AI Assistance), and (ii) the extent to which early-stage firm formation relies on trial-and-error (Experimentation Intensity). We link these measures to newly established firms across industries and quarters between 2020 and 2024.

Using a difference-in-differences framework with continuous treatment intensity, we exploit the arrival of ChatGPT as a quasi-exogenous shock to AI availability. Our results show that industries more compatible with AI experienced a significant increase in new firm creation following the release of ChatGPT. A one-standard deviation increase in AI Assistance leads to over a 10% increase in firm entry, holding other factors constant. This effect is not driven by pre-existing trends and remains robust to controls for funding conditions and industry fixed effects.

The mechanism linking AI availability and firm formation is one entailing reductions in the cost of experimentation: industries that both rely on experimentation and benefit strongly from AI assistance exhibit the largest post-ChatGPT increase in firm creation. In contrast, we find little empirical support for alternative explanations based on capital supply or AI bridging human-capital gaps. While these mechanisms are theoretically plausible, their contributions appear limited in practice during the period we study.

Moreover, our analysis of firm-level outcomes shows that the rise in entrepreneurship in AI-compatible industries is not accompanied by a decline in firm quality. On the contrary, new firms created after the release of ChatGPT are more likely to survive, expand faster, and attract more educated workers, indicating that AI lowers frictional barriers to entry without encouraging the formation of lower-quality ventures.

Our findings contribute to a growing literature on the economic implications of AI by offering new evidence that generative AI tools can foster entrepreneurial activity—not merely by enhancing productivity, but by lowering cognitive and operational barriers to firm formation. More broadly, these results highlight how democratized access to advanced technologies can unlock new entry margins, especially in sectors where innovation and experimentation are central.

Several questions remain open; we highlight two. First, the shock is still recent, so the long-run trajectory of post-ChatGPT entrants—whether their early survival and growth advantages persist as these firms mature and as incumbents adopt AI—is not yet observable. Second, beyond the rate and quality of entry, AI may also reshape its *direction*—whether cheaper experimentation broadens the set of ideas entrepreneurs pursue or instead channels them toward similar, AI-amenable ventures. We leave both to future work.

REFERENCES


- ACEMOGLU, D., D. AUTOR, J. HAZELL, AND P. RESTREPO (2022): “Artificial intelligence and jobs: Evidence from online vacancies,” *Journal of Labor Economics*, 40, S293–S340.
- AKER, J. C. AND I. M. MBITI (2010): “Mobile phones and economic development in Africa,” *Journal of economic Perspectives*, 24, 207–232.
- AUER, R., D. KÖPFER, AND J. ŠVÉDA (2024): “The rise of generative AI: modelling exposure, substitution, and inequality effects on the US labour market,” Tech. rep., CESifo Working Paper.
- BABINA, T., A. FEDYK, A. HE, AND J. HODSON (2024): “Artificial intelligence, firm growth, and product innovation,” *Journal of Financial Economics*, 151, 103745.
- BAO, J., B. LOU, AND H. SUN (2025): “Generative AI and Entrepreneurial Entry,” *Available at SSRN 5167018*.
- BENA, J., B. BIAN, AND M. GIANNETTI (2026): “Prompted to Start: How Generative AI is Transforming Entrepreneurship,” *HKU Jockey Club Enterprise Sustainability Global Research Institute Paper*.
- BERNSTEIN, S., E. COLONNELLI, D. MALACRINO, AND T. MCQUADE (2022): “Who creates new firms when local opportunities arise?” *Journal of Financial Economics*, 143, 107–130.
- BICK, A., A. BLANDIN, AND D. J. DEMING (2024): “The rapid adoption of generative ai,” Tech. rep., National Bureau of Economic Research.
- BLANK, S. (2020): *The four steps to the epiphany: successful strategies for products that win*, John Wiley & Sons.
- BONNEY, K., C. BREAU, C. BUFFINGTON, E. DINLERSOZ, L. FOSTER, N. GOLDSCHLAG, J. HALTIWANGER, Z. KROFF, AND K. SAVAGE (2024): “The impact of AI on the workforce: Tasks versus jobs?” *Economics Letters*, 244, 111971.
- BOUDREAUX, C. J., M. P. ESCALERAS, AND M. SKIDMORE (2019): “Natural disasters and entrepreneurship activity,” *Economics Letters*, 182, 82–85.
- BOUSSIOUX, L., J. N. LANE, M. ZHANG, V. JACIMOVIC, AND K. R. LAKHANI (2024): “The crowdless future? Generative AI and creative problem-solving,” *Organization Science*, 35, 1589–1607.
- BRYNJOLFSSON, E., B. CHANDAR, AND R. CHEN (2025a): “Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence,” .
- BRYNJOLFSSON, E., D. LI, AND L. RAYMOND (2025b): “Generative AI at work,” *The Quarterly Journal of Economics*, qjae044.
- CALLAWAY, B., A. GOODMAN-BACON, AND P. H. SANT’ANNA (2024): “Difference-in-differences with a continuous treatment,” Tech. rep., National Bureau of Economic Research.
- CHEN, J. AND J. ROTH (2024): “Logs with zeros? Some problems and solutions,” *The Quarterly Journal of Economics*, 139, 891–936.
- CHEN, Q., J. GE, H. XIE, X. XU, AND Y. YANG (2025): “Large language models at work in China’s labor market,” *China Economic Review*, 102413.
- CONTIGIANI, A. (2023): “Experimentation and appropriability in early-stage ventures: Evidence from the US software industry,” *Strategic Management Journal*, 44, 2128–2174.
- CONTIGIANI, A. AND T. YOUNG-HYMAN (2022): “Experimentation, planning, and structure in early-stage ventures: Evidence from pitch decks,” *Strategic Entrepreneurship Journal*, 16, 425–459.
- CUI, C. AND L. S.-Z. LI (2023): “Trade policy uncertainty and new firm entry: Evidence from China,” *Journal of Development Economics*, 163, 103093.

- DAVIS, S. J., J. HALTIWANGER, AND S. SCHUH (1996): “Small business and job creation: Dissecting the myth and reassessing the facts,” *Small business economics*, 8, 297–315.
- DELL’ACQUA, F., E. MCFOWLAND III, E. R. MOLLICK, H. LIFSHITZ-ASSAF, K. KELLOGG, S. RAJENDRAN, L. KRAYER, F. CANDELON, AND K. R. LAKHANI (2023): “Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality,” *Harvard Business School Technology & Operations Mgt. Unit Working Paper*.
- DINGEL, J. I. AND B. NEIMAN (2020): “How many jobs can be done at home?” *Journal of public economics*, 189, 104235.
- DOSHI, A. R. AND O. P. HAUSER (2024): “Generative AI enhances individual creativity but reduces the collective diversity of novel content,” *Science advances*, 10, eadn5290.
- EISFELDT, A. L., G. SCHUBERT, B. TASKA, AND M. B. ZHANG (2025): “Generative AI and Firm Values,” *Available at SSRN 4436627*.
- ELOUNDOU, T., S. MANNING, P. MISHKIN, AND D. ROCK (2023): “Gpts are gpts: An early look at the labor market impact potential of large language models,” *arXiv preprint arXiv:2303.10130*.
- ERICSON, R. AND A. PAKES (1995): “Markov-perfect industry dynamics: A framework for empirical work,” *The Review of economic studies*, 62, 53–82.
- EVANS, D. S. AND B. JOVANOVIC (1989): “An Estimated Model of Entrepreneurial Choice under Liquidity Constraints,” *Journal of Political Economy*, 97, 808–827.
- EWENS, M., R. NANDA, AND M. RHODES-KROPF (2018): “Cost of experimentation and the evolution of venture capital,” *Journal of Financial Economics*, 128, 422–442.
- GHOSH, D., R. GHOSH, S. ROY CHOWDHURY, AND B. GANGULY (2025): “AI-exposure and labour market: a systematic literature review on estimations, validations, and perceptions,” *Management Review Quarterly*, 75, 677–704.
- GMYREK, P., J. BERG, AND D. BESCOND (2023): “Generative AI and jobs: A global analysis of potential effects on job quantity and quality,” *ILO Working paper*, 96.
- GOMPERS, P., J. LERNER, AND D. SCHARFSTEIN (2005): “Entrepreneurial spawning: Public corporations and the genesis of new ventures, 1986 to 1999,” *The journal of Finance*, 60, 577–614.
- GUPTA, A., F. QIAN, E. SIMINTZI, AND Y. SUN (2025): “Generative AI and Entrepreneurship,” *Available at SSRN 5971474*.
- HALTIWANGER, J. C. (2022): “Entrepreneurship during the COVID-19 pandemic: Evidence from the business formation statistics,” *Entrepreneurship and Innovation Policy and the Economy*, 1, 9–42.
- HAMPOLE, M., D. PAPANIKOLAOU, L. D. SCHMIDT, AND B. SEEGMILLER (2025): “Artificial intelligence and the labor market,” Tech. rep., National Bureau of Economic Research.
- HASBI, M. (2020): “Impact of very high-speed broadband on company creation and entrepreneurship: Empirical Evidence,” *Telecommunications Policy*, 44, 101873.
- HELLER, D. AND D. ASAM (2024): “Generative AI and Firm-level Productivity: Evidence from Startup Funding Dynamics,” *Available at SSRN 4877505*.
- HELLMANN, T. AND M. PURI (2000): “The interaction between product market and financing strategy: The role of venture capital,” *The review of financial studies*, 13, 959–984.
- (2002): “Venture capital and the professionalization of start-up firms: Empirical evidence,” *The journal of finance*, 57, 169–197.
- HJORT, J. AND J. POULSEN (2019): “The arrival of fast internet and employment in Africa,” *American Economic Review*, 109, 1032–1079.


- HJORT, J. AND L. TIAN (2021): “The economic impact of internet connectivity in developing countries,” *Annual Review of Economics*, 17.
- HOPENHAYN, H. A. (1992): “Entry, Exit, and Firm Dynamics in Long Run Equilibrium,” *Econometrica*, 60, 1127–1150.
- HUI, X., O. RESHEF, AND L. ZHOU (2024): “The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market,” *Organization Science*, 35, 1977–1989.
- HUMLUM, A. AND E. VESTERGAARD (2025): “The unequal adoption of ChatGPT exacerbates existing inequalities among workers,” *Proceedings of the National Academy of Sciences*, 122, e2414972121.
- JOVANOVIĆ, B. (1982): “Selection and the evolution of industry,” *Econometrica*, 50, 649–670.
- KERR, W. R. AND R. NANDA (2011): “Financing constraints and entrepreneurship,” *Handbook of Research on Innovation and Entrepreneurship*. Cheltenham: Elgar, 88–103.
- KERR, W. R., R. NANDA, AND M. RHODES-KROPF (2014): “Entrepreneurship as experimentation,” *Journal of Economic Perspectives*, 28, 25–48.
- KNIGHT, F. H. (1921): *Risk, uncertainty and profit*, Houghton Mifflin Company.
- KONING, R., S. HASAN, AND A. CHATTERJI (2022): “Experimentation and start-up performance: Evidence from A/B testing,” *Management Science*, 68, 6434–6453.
- KWAN, A., B. MATTHIES, R. R. TOWNSEND, AND T. XU (2025): “Entrepreneurial spawning from remote work,” Tech. rep., National Bureau of Economic Research.
- LABASCHIN, B., T. ELOUNDOU, S. MANNING, P. MISHKIN, AND D. ROCK (2025): “Extending “GPTs Are GPTs” to Firms,” in *AEA Papers and Proceedings*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, vol. 115, 51–55.
- LAZEAR, E. P. (2005): “Entrepreneurship,” *Journal of Labor Economics*, 23, 649–680.
- LEVINE, R. AND Y. RUBINSTEIN (2017): “Smart and illicit: who becomes an entrepreneur and do they earn more?” *The Quarterly journal of economics*, 132, 963–1018.
- LINDEBAUM, D. AND P. FLEMING (2024): “ChatGPT undermines human reflexivity, scientific responsibility and responsible management research,” *British Journal of Management*, 35, 566–575.
- LUCAS JR, R. E. (1978): “On the size distribution of business firms,” *The Bell Journal of Economics*, 508–523.
- MANSO, G. (2011): “Motivating innovation,” *The journal of finance*, 66, 1823–1860.
- (2016): “Experimentation and the Returns to Entrepreneurship,” *Review of Financial Studies*, 29, 2319–2340.
- MASLEJ, N., L. FATTORINI, R. PERRAULT, Y. GIL, V. PARLI, N. KARIUKI, E. CAPSTICK, A. REUEL, E. BRYNJOLFSSON, J. ETCHEMENDY, ET AL. (2025): “Artificial intelligence index report 2025,” *arXiv preprint arXiv:2504.07139*.
- MEINCKE, L., E. R. MOLICK, AND C. TERWIESCH (2024): “Prompting diverse ideas: Increasing AI idea variance,” *arXiv preprint arXiv:2402.01727*.
- MYERS, K. R., L. LANAHAN, AND E. JOHNSON (2025): “Small Business Innovation Applied to National Needs,” Tech. rep., National Bureau of Economic Research.
- NOY, S. AND W. ZHANG (2023): “Experimental evidence on the productivity effects of generative artificial intelligence,” *Science*, 381, 187–192.
- OTIS, N., R. CLARKE, S. DELECOURT, D. HOLTZ, AND R. KONING (2024): “The uneven impact of generative AI on entrepreneurial performance,” *Available at SSRN 4671369*.
- RIES, E. (2011): *The lean startup: How today’s entrepreneurs use continuous innovation to create radically successful businesses*, Crown Currency.

- SCHUMPETER, J. A. (1934): *The Theory of Economic Development: An Inquiry into Profits, Capital, Credit, Interest, and the Business Cycle*, Cambridge, MA: Harvard University Press.
- (1942): *Capitalism, Socialism and Democracy*, New York: Harper & Brothers.
- SCHWARTZ, M. (2011): “Incubating an Illusion? Long-term incubator firm performance after graduation,” *Growth and change*, 42, 491–516.
- STEINWENDER, C. (2018): “Real effects of information frictions: When the states and the kingdom became united,” *American Economic Review*, 108, 657–696.
- STEPHENS, H. M., E. A. MACK, AND J. MANN (2022): “Broadband and entrepreneurship: An empirical assessment of the connection between broadband availability and new business activity across the United States,” *Telematics and Informatics*, 74, 101873.
- WANG, X., C. FENG, AND T. SUN (2025): “AI Spillover is Different: Flat and Lean Firms as Engines of AI Diffusion and Productivity Gain,” *arXiv preprint arXiv:2511.02099*.
- YANG, C.-H. (2022): “How artificial intelligence technology affects productivity and employment: Firm-level evidence from Taiwan,” *Research Policy*, 51, 104536.


FIGURES AND TABLES

Figure 1. AI Usage and Reliance Among Surveyed Entrepreneurs


Notes: This figure summarizes survey-reported AI usage and reliance among U.S. entrepreneurs. The survey was fielded between November 2025 and May 2026 through Dynata’s respondent panel. Qualified respondents are new or aspiring U.S. entrepreneurs who started a company or began pursuing a business idea on or after January 2023. The final analytic sample after quality filtering consists of 1,036 respondents. Panel A reports the distribution of self-reported AI usage across the full sample, in response to the question “How much have you used AI tools in connection with your business?” Panel B reports the distribution of self-reported AI reliance among the 818 respondents who use AI at all, in response to the question “If AI tools were not available, how likely is it that you would have pursued the same business idea?”

Figure 2. Distribution of the Number of Tasks per Industry


Notes: The figure presents the distribution of the number of tasks across all 81 three-digit NAICS industries in our sample. Tasks are extracted from the corpus of several texts focusing on entrepreneurship. After all the tasks have been extracted and deduplicated, an LLM assesses the relevance of any given task with respect to a particular industry. Once the universe of industry-specific tasks is created, we calculate the number of tasks per industry and plot it. The horizontal axis reports the number of tasks assigned to an industry; the vertical axis shows the share of industries in each bin. For reference, we overlay a normal density curve fitted to the sample mean and standard deviation.

Figure 3. AI Assistance and Experimentation Intensity Scores Distributions

Notes: This panel of plots displays the distributions of the two measures generated through the GenAI workflow across all 81 three-digit NAICS industries in our sample. Tasks are extracted from the corpus of several texts focusing on entrepreneurship. After all the tasks have been extracted and deduplicated, an LLM assesses the relevance of any given task with respect to a particular industry. Once the universe of industry-specific tasks is created, each task is scored on Importance Score, AI Assistance and Experimentation Intensity. We then take the industry level average of the two measures weighting by the Importance Score. Panel (a) shows the distribution of the AI Assistance Score, while panel (b) shows the distribution of the Experimentation Intensity Score. Both variables are presented at the industry level. For each of the variables, we also overlay a normal distribution curve for comparison.

Figure 4. Google Search Intensity for ChatGPT around Model Release


Notes: This figure plots weekly Google Trends search intensity for the term “ChatGPT.” Google Trends normalizes the number of searches within a specific period so to range on a 0–100 scale, where 100 represents the peak search volume during the sample period. The vertical red line marks the week in which ChatGPT was released. ChatGPT launched on November 30, 2022, which falls within the week beginning November 27, 2022 (as shown on the horizontal axis).

Figure 5. Event Study of AI Assistance on Firm Creation Around the Release of ChatGPT

Notes: The figure presents the event study for the sample of 81 three-digit NAICS industries in our analysis, over the time period 2020Q4–2024Q1. We plot the estimates from the dynamic difference-in-differences model for the interaction of the quarter-year dummies with the AI Assistance_{*i*} from the specification

$$\text{IHS(N, New Firms)}_{it} = \beta_0 + \sum_{t \neq 2022Q3} \beta_t \text{AI Assistance}_i \times 1\{D = t\}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. The specifications include both industry and quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level. The specification in Panel (a) does not include controls. The specification in Panel (b) includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage.


Figure 6. Event Study of AI Assistance on Firm Creation over the Extended Window

Notes: The figure presents the event study for the sample of 81 three-digit NAICS industries in our analysis, over the extended time period 2018Q4–2024Q1. We plot the estimates from the dynamic difference-in-differences model for the interaction of the quarter-year dummies with the $AI Assistance_i$ from the specification


$$IHS(N. New Firms)_{it} = \beta_0 + \sum_{t \neq 2022q3} \beta_t AI Assistance_i \times 1\{D = t\}_t + Controls_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. The extended window adds eight quarters untouched by the pandemic, so that the leads span both the pre-pandemic period and the pandemic-era entry surge. The specifications include both industry and quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level. The specification in Panel (a) does not include controls. The specification in Panel (b) includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage.

Figure 7. Survey Validation of the AI Assistance Score




Notes: This figure validates the LLM-generated AI Assistance Score against a survey-based measure of AI usefulness reported by new and aspiring U.S. entrepreneurs. The survey was fielded between November 2025 and May 2026 through Dynata’s respondent panel. Qualified respondents are new or aspiring U.S. entrepreneurs who started a company or began pursuing a business idea on or after January 2023. The analytic sample consists of 1,036 respondents. Each respondent was presented with 20 substantive task groups spanning the spectrum of entrepreneurial activity, and asked to rate AI’s usefulness for each. The y-axis plots the mean survey rating per task group, restricted to respondents who reported using AI for that task (Q5 = “Yes”). The x-axis plots the LLM-generated AI Assistance aggregated to the same task-group level. To construct this aggregate, each task extracted by our LLM pipeline is mapped to one of the 20 task groups, and the task-group AI Assistance is computed as the importance-weighted average of the LLM-generated AI Assistance values across its constituent tasks.

Figure 8. Event Study of Cost of Experimentation on Firm Creation Around the Release of ChatGPT

Notes: The figure presents the event study for the sample of 81 three-digit NAICS industries in our analysis, over the time period 2020Q4–2024Q1. We plot the estimates from the dynamic triple differences model for the interaction of the quarter-year dummies with the AI Assistance_{*i*} × Experimentation Intense_{*i*} from the specification


$$\begin{aligned}
 \text{IHS}(\text{N. New Firms})_{it} = & \beta_0 + \sum_{t \neq 2022q3} \beta_t \text{AI Assistance}_i \times \text{Experimentation Intense}_i \times 1\{D = t\}_t \\
 & + \sum_{t \neq 2022q3} \gamma_t \text{AI Assistance}_i \times 1\{D = t\}_t + \sum_{t \neq 2022q3} \delta_t \text{Experimentation Intense}_i \times 1\{D = t\}_t \\
 & + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it}
 \end{aligned}$$

The dependent variable is the Inverse Hyperbolic Sine of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment heterogeneity comes from the indicator variable Experimentation Intense, which equals 1 if the industry reports a value of Experimentation Intensity within the top quartile and 0 otherwise. The specifications include both industry and sector-by-quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level. The specification in Panel (a) does not include controls. The specification in Panel (b) includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage.

Figure 9. Most Mentioned AI Usage by Entrepreneurs

Notes: This figure summarizes the most frequently mentioned uses of AI tools among the 818 respondents in our survey who reported using AI in their business. To produce the underlying data, we asked respondents to describe in their own words how they use AI tools in their business. We then used a large language model to read each response and assign one or more content codes summarizing the types of AI use described. Codes were not pre-specified but emerged from the model’s reading of the responses; the model identified 22 distinct codes across the full set of responses, the five most frequently mentioned of which are shown here. Because a single response can describe multiple types of AI use, individual responses may receive multiple codes. Each bar shows the percentage of AI users whose response was tagged with the corresponding code. Bars highlighted in maroon correspond to ideation and brainstorming and marketing and content creation, two categories we identify as experimentation-intensive, in that they involve the iterative generation, evaluation, and refinement of new ideas. The full list of 22 codes and their frequencies is reported in Section E of the Online Appendix.

Figure 10. Event Study of AI Assistance on Allocated Capital Around the Release of ChatGPT



Notes: The figure presents the event study for the sample of 81 three-digit NAICS industries in our analysis, over the time period 2020Q4–2024Q1. We plot the estimates from the dynamic difference-in-differences model for the interaction of the quarter-year dummies with the AI Assistance_{*i*} from the specification

$$\text{IHS(Investments)}_{it} = \beta_0 + \sum_{t \neq 2022q3} \beta_t \text{AI Assistance}_i \times 1\{D = t\}_t + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it}$$

The dependent variable is the Inverse Hyperbolic Sine of the dollar amount of VCPE investments, in thousands. Treatment intensity is proxied by AI Assistance in its standardized form. The specification includes both industry and sector-by-quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level.

Table 1. Top and Bottom 20 Industries by AI Assistance Score

Spot	Industry Name	NAICS Code	Value
1	Web Search Portals, Libraries, Archives, and Other Information Services	519	47.226
2	Computing Infrastructure Providers, Data Processing, Web Hosting, and Related Services	518	47.021
3	Professional, Scientific, and Technical Services	541	46.140
4	Computer and Electronic Product Manufacturing	334	45.608
5	Broadcasting and Content Providers	516	45.234
6	Furniture, Home Furnishings, Electronics, and Appliance Retailers	449	44.940
7	Telecommunications	517	44.756
8	Administrative and Support Services	561	44.342
9	Merchant Wholesalers, Nondurable Goods	424	44.236
10	Sporting Goods, Hobby, Musical Instrument, Book, and Miscellaneous Retailers	459	44.197
11	Clothing, Clothing Accessories, Shoe, and Jewelry Retailers	458	44.132
12	Merchant Wholesalers, Durable Goods	423	44.070
13	Management of Companies and Enterprises	551	43.771
14	Publishing Industries	513	43.733
15	Health and Personal Care Retailers	456	43.432
16	Educational Services	611	43.248
17	Miscellaneous Manufacturing	339	43.160
18	General Merchandise Retailers	455	43.153
19	Rental and Leasing Services	532	43.113
20	Wholesale Trade Agents and Brokers	425	43.021
⋮	⋮	⋮	⋮
62	Support Activities for Agriculture and Forestry	115	40.001
63	Textile Mills	313	39.914
64	Gasoline Stations and Fuel Dealers	457	39.652
65	Support Activities for Mining	213	39.598
66	Private Households	814	39.546
67	Construction of Buildings	236	39.532
68	Animal Production and Aquaculture	112	39.461
69	Mining (except Oil and Gas)	212	39.406
70	Nursing and Residential Care Facilities	623	39.395
71	Rail Transportation	482	39.301
72	Hospitals	622	39.134
73	Water Transportation	483	38.970
74	Petroleum and Coal Products Manufacturing	324	38.784
75	Pipeline Transportation	486	38.354
76	Heavy and Civil Engineering Construction	237	38.143
77	Fishing, Hunting and Trapping	114	37.928
78	Forestry and Logging	113	37.907
79	Oil and Gas Extraction	211	37.815
80	Crop Production	111	37.642
81	Primary Metal Manufacturing	331	37.417

Notes: This table lists the 20 industries with the highest and the 20 industries with the lowest AI Assistance Scores among 81 three-digit NAICS industries. Tasks are extracted from a corpus of entrepreneurship-focused texts and, after extraction and deduplication, a large language model (LLM) assesses each task’s relevance to each industry and tailors the description accordingly. For each industry, the AI Assistance Score is the Importance Score-weighted average of task-level AI Assistance Scores; higher values indicate greater potential for AI to assist industry tasks. A detailed explanation of the workflow to generate the variables is presented in Subsection D.2 of the Online Appendix.

Table 2. Top and Bottom 20 Industries by Experimentation Intensity

Spot	Industry Name	NAICS Code	Value
1	Motion Picture and Sound Recording Industries	512	66.626
2	Broadcasting and Content Providers	516	66.194
3	Computing Infrastructure Providers, Data Processing, Web Hosting, and Related Services	518	66.165
4	Food Services and Drinking Places	722	65.983
5	Web Search Portals, Libraries, Archives, and Other Information Services	519	65.930
6	Clothing, Clothing Accessories, Shoe, and Jewelry Retailers	458	65.870
7	Sporting Goods, Hobby, Musical Instrument, Book, and Miscellaneous Retailers	459	65.820
8	Food and Beverage Retailers	445	65.700
9	Computer and Electronic Product Manufacturing	334	65.603
10	Professional, Scientific, and Technical Services	541	65.583
11	Educational Services	611	65.572
12	Accommodation	721	65.546
13	General Merchandise Retailers	455	65.544
14	Furniture, Home Furnishings, Electronics, and Appliance Retailers	449	65.498
15	Miscellaneous Manufacturing	339	65.497
16	Telecommunications	517	65.476
17	Publishing Industries	513	65.404
18	Apparel Manufacturing	315	65.263
19	Health and Personal Care Retailers	456	65.245
20	Rental and Leasing Services	532	64.932
⋮	⋮	⋮	⋮
62	Motor Vehicle and Parts Dealers	441	62.860
63	Lessors of Nonfinancial Intangible Assets (except Copyrighted Works)	533	62.839
64	Warehousing and Storage	493	62.804
65	Mining (except Oil and Gas)	212	62.791
66	Water Transportation	483	62.655
67	Rail Transportation	482	62.344
68	Construction of Buildings	236	62.241
69	Hospitals	622	61.873
70	Wholesale Trade Agents and Brokers	425	61.686
71	Chemical Manufacturing	325	61.605
72	Heavy and Civil Engineering Construction	237	60.863
73	Utilities	221	60.768
74	Paper Manufacturing	322	60.708
75	Primary Metal Manufacturing	331	60.696
76	Forestry and Logging	113	60.630
77	Nursing and Residential Care Facilities	623	60.216
78	Gasoline Stations and Fuel Dealers	457	59.166
79	Petroleum and Coal Products Manufacturing	324	59.143
80	Oil and Gas Extraction	211	58.639
81	Pipeline Transportation	486	57.516

Notes: This table lists the 20 industries with the highest and the 20 industries with the lowest Experimentation Intensity Scores among 81 three-digit NAICS industries. Tasks are extracted from a corpus of entrepreneurship-focused texts and, after extraction and deduplication, a large language model (LLM) assesses each task’s relevance to each industry and tailors the description accordingly. For each industry, the Experimentation Intensity Score is the Importance Score-weighted average of task-level Experimentation Intensity Scores; higher values indicate greater industry reliance on trial-and-error, feedback-driven refinement and iterative learning. A detailed explanation of the workflow to generate the variables is presented in the Subsection D.2 of the Online Appendix.

Table 3. Summary Statistics

	(1) N. Obs	(2) Min	(3) p50	(4) Mean	(5) Max	(6) SD
Panel (a): Industry $\times$ Quarter						
N. New Firms	1,134	26.000	5,401.000	14,064.000	322,847.000	32,689.507
Investments (\$ 000)	1,134	0.000	97.000	1,730.000	94,564.000	7,090.847
ROA	1,134	-5.844	0.005	-0.003	0.316	0.177
R&D Intensity	1,134	0.000	0.018	0.077	11.235	0.397
Capex Intensity	1,134	-36.034	0.013	0.003	17.816	1.346
Liquidity Ratio	1,134	0.004	0.096	0.127	0.726	0.103
Market Leverage Ratio	1,134	0.000	0.281	0.408	6.545	0.465
Panel (b): Industry						
AI Assistance Score	81	37.417	41.365	41.587	47.226	2.172
Experimentation Intensity	81	57.516	63.810	63.652	66.626	1.859

Notes: This table reports summary statistics for the variables used in the analysis over 2020Q4–2024Q1. Columns display the following: (1) number of observations, (2) minimum, (3) median, (4) mean, (5) maximum, and (6) standard deviation. Panel (a) presents summary statistics for variables varying at the industry-by-quarter level, while Panel (b) presents summary statistics for variables varying at the industry level. N. New Firms is the count of new firms created. Investments is the total amount (in thousands of dollars) of VCPE investments from Preqin. The following Compustat variables are industry medians. Return on Assets (ROA) is defined as income before extraordinary items divided by total assets. R&D Intensity is defined as R&D expenses over revenue. CapEx Intensity is defined as capital expenditures over revenue. Liquidity Ratio is cash holdings over total assets. Market Leverage Ratio is calculated as the sum of long-term and current debt divided by the market value of equity. AI Assistance Score is the industry-level average proxying for how much AI can help with task completion. Experimentation Intensity is the industry-level average proxying for how prominent experimentation is.

Table 4. Effect of AI Assistance on Firm Creation After the Release of ChatGPT

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms
AI Assistance $\times$ Post	0.105*** (0.025)	0.100*** (0.025)
Observations	1,134	1,134
Industry 3D FE	Yes	Yes
Quarter FE	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No
Controls	No	Yes
Adj. R. sq.	0.982	0.982

Notes: This table presents results testing the effect of AI's arrival on firm creation. The design is a difference-in-differences model at the industry-by-quarter level with continuous treatment. The sample includes 81 three-digit NAICS industries over 2020Q4–2024Q1. The empirical specification is

$$\text{IHS(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. All specifications include industry and quarter fixed effects. Column (2) additionally includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Table 5. Effect of AI Assistance on Firm Creation over the Extended Window

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms
AI Assistance $\times$ Post	0.112*** (0.027)	0.112*** (0.027)
Observations	1,782	1,782
Industry 3D FE	Yes	Yes
Quarter FE	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No
Controls	No	Yes
Adj. R. sq.	0.984	0.984

Notes: This table presents results testing the effect of AI's arrival on firm creation over the extended estimation window. The design is a difference-in-differences model at the industry-by-quarter level with continuous treatment. The sample includes 81 three-digit NAICS industries over 2018Q4–2024Q1. The empirical specification is

$$\text{IHS(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. All specifications include industry and quarter fixed effects. Column (2) additionally includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Table 6. Horse Race Analysis of Confounding Factors

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms	(3) IHS N. New Firms	(4) IHS N. New Firms	(5) IHS N. New Firms	(6) IHS N. New Firms	(7) IHS N. New Firms
AI Assistance $\times$ Post	0.100*** (0.025)		0.090*** (0.032)		0.093*** (0.032)		0.099*** (0.026)
Teleworkability $\times$ Post		0.069** (0.027)	0.018 (0.030)				
IT Capital Intensity $\times$ Post				0.068*** (0.024)	0.010 (0.031)		
Pre-Trend $\times$ Post						0.022 (0.025)	0.006 (0.023)
Observations	1,134	1,134	1,134	1,134	1,134	1,134	1,134
Industry 3D FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No	No	No	No	No	No
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R. sq.	0.982	0.982	0.982	0.981	0.982	0.981	0.982

Notes: This table presents results racing AI Assistance against three rival exposure measures. The design is a horse race difference-in-differences model at the industry-by-quarter level with continuous treatment. The sample includes 81 three-digit NAICS industries over 2020Q4–2024Q1. Each horse race regression uses the empirical specification

$$\text{IHS(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \beta_2 \text{Rival}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. In Column (1) we present the baseline results from the specification of Subsection 4.1. In Columns (2)–(3), the rival measure is Teleworkability, constructed from the occupation-level classification of [Dingel and Neiman \(2020\)](#) aggregated to three-digit NAICS using OEWS employment weights: Column (2) enters Teleworkability alone and Column (3) enters both measures jointly. In Columns (4)–(5), the same process is repeated using IT Capital Intensity, the IT share of the industry capital stock from the BEA detailed fixed-asset estimates, fixed at 2018 values. In Columns (6)–(7), the same process is repeated using the Pre-Trend, each industry’s OLS slope of IHS firm formation on time fitted over 2015Q1–2018Q3. All rival measures are standardized over the estimation sample. All specifications include industry and quarter fixed effects and the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Fully specified horse-race tables, including both estimation windows and residualized variants, are reported in Subsection A.1.6 of the Online Appendix. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Table 7. Horse Race Analysis of Competing AI Metrics

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms	(3) IHS N. New Firms
AI Assistance $\times$ Post	0.105*** (0.025)	0.087*** (0.029)	0.110*** (0.027)
AI Task Exposure $\times$ Post		0.036 (0.033)	
AI Demand $\times$ Post			-0.014 (0.019)
Observations	1,134	1,134	1,134
Industry 3D FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No	No
Controls	No	No	No
Adj. R. sq.	0.982	0.982	0.982

Notes: This table presents results probing the interpretation of the AI Assistance measure. The design is a horse race difference-in-differences model at the industry-by-quarter. The sample includes 81 three-digit NAICS industries over 2020Q4–2024Q1. Each horse race regression uses the empirical specification

$$\text{IHS}(\text{N. New Firms})_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \beta_2 \text{Alternative Measure}_i \times \text{Post}_t + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. In Column (1) we present the baseline results from Equation 6. In Column (2) the AI Assistance variable is juxtaposed to the AI Task Exposure measure, both in standardized form. In Column (3) the same process is repeated, using the AI Demand measure instead. All specifications include industry and quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Table 8. Effects of the Cost of Experimentation on Firm Creation

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms
AI Assistance $\times$ Experimentation Intense $\times$ Post	0.261** (0.111)	0.264** (0.112)
AI Assistance $\times$ Post	0.037 (0.054)	0.032 (0.054)
Experimentation Intense $\times$ Post	-0.190 (0.119)	-0.193 (0.118)
Observations	1,078	1,078
Industry 3D FE	Yes	Yes
Quarter FE	No	No
Industry 2D $\times$ Quarter FE	Yes	Yes
Controls	No	Yes
Adj. R sq.	0.981	0.981

Notes: This table presents results testing whether the reduced cost-of-experimentation mechanism operates. The design is a triple differences model at the industry-by-quarter level. The sample includes 81 three-digit NAICS industries over 2020Q4–2024Q1. The empirical specification is

$$\begin{aligned} \text{IHS(N. New Firms)}_{it} = & \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Experimentation Intense}_i \times \text{Post}_t \\ & + \beta_2 \text{AI Assistance}_i \times \text{Post}_t + \beta_3 \text{Experimentation Intense}_i \times \text{Post}_t \\ & + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it} \end{aligned}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. Treatment heterogeneity comes from the indicator Experimentation Intense, which equals 1 if the industry's Experimentation Intensity Score is in the top quartile and 0 otherwise. The specifications include industry and sector-by-quarter fixed effects. Column (2) additionally includes the following controls: lagged VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Table 9. Mediation Analysis: Capital Allocation, AI Assistance and Firm Creation

VARIABLES	(1) IHS N. New Firms	(2) IHS Investments	(3) IHS N. New Firms
AI Assistance $\times$ Post	0.102** (0.046)	-0.586** (0.239)	0.102** (0.046)
IHS Investments			0.000 (0.004)
Observations	1,078	1,078	1,078
Industry 3D FE	Yes	Yes	Yes
Quarter FE	No	No	No
Industry 2D $\times$ Quarter FE	Yes	Yes	Yes
Controls	No	No	No
Adj. R sq.	0.980	0.687	0.980

Notes: This table reports a mediation-style decomposition to study capital reallocation as a potential mechanism linking AI's arrival and firm creation. The analysis is at the industry-by-quarter level for 81 three-digit NAICS industries over 2020Q4–2024Q1. The mediator is total capital allocated, measured as the inverse hyperbolic sine (IHS) of total VCPE investments from Preqin (in \$000).

Column (1) reports the results for the estimation of the total impact that the treatment has on firm formation. The specification is

$$\text{IHS(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \mu_i + \mu_{It} + \epsilon_{it}$$

Column (2) assesses the linkage between the treatment and the total capital allocation, according to the specification

$$\text{IHS(Investments)}_{it} = \gamma_0 + \gamma_1 \text{AI Assistance}_i \times \text{Post}_t + \mu_i + \mu_{It} + \epsilon_{it}$$

Column (3) reports the results for the estimation of the mediated effect of the treatment once we control for total allocated capital. The specification is

$$\text{IHS(N. New Firms)}_{it} = \delta_0 + \delta_1 \text{AI Assistance}_i \times \text{Post}_t + \delta_2 \text{IHS(Investments)}_{it} + \mu_i + \mu_{It} + \epsilon_{it}$$

All specifications include both industry and sector-by-quarter fixed effects. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which is equal to 1 for all quarters after 2022Q3, and 0 otherwise. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Table 10. AI Availability, Missing Skills and Entrepreneurial Entry - All Individuals

VARIABLES	(1) P(Entrepreneur)	(2) P(Entrepreneur)
Skill Gap $\times$ Post		0.715 (0.907)
Skill Gap	-2.265*** (1.042)	-2.427*** (1.054)
Observations	173,968,981	173,968,981
Sample	All Individuals	All Individuals
User $\times$ Quarter FE	Yes	Yes
User $\times$ Industry 3D FE	Yes	Yes
Quarter $\times$ Industry 3D FE	Yes	Yes
Adj. R sq.	0.119	0.119

Notes: The table reports the results for the test assessing whether AI promotes self-employment by bridging human-capital gaps for potential entrepreneurs. The analysis is at the individual-by-industry-by-quarter level. The sample is composed of 200,000 individuals randomly sampled from the profiles on Revelio Labs. Industries are reported at the three-digit NAICS level. The dataset includes 81 industries over the period 2020Q4–2024Q1. The model used is

$$P(\text{Entrepreneur})_{jit} = \beta_0 + \beta_1 \text{Skill Gap}_{jit} \times \text{Post}_t + \beta_2 \text{Skill Gap}_{jit} + \mu_{ji} + \mu_{jt} + \mu_{it} + \epsilon_{jit}$$

The dependent variable $P(\text{Entrepreneur})_{jit}$ is an indicator for whether the individual becomes an entrepreneur in a specific industry and quarter, multiplied by 10,000 for ease of exposition. Skill Gap is the measure of human-capital gap calculated for each individual, in a given quarter, with respect to a particular industry. We standardize Skill Gap using the pre-shock mean and standard deviation. The binary indicator Post, which is equal to 1 for all quarters after 2022Q3, and 0 otherwise, proxies for the availability of ChatGPT. The specifications include individual-by-industry, individual-by-quarter and industry-by-quarter fixed effects. Column (1) reports results assessing the predictive power for the measure of skill mismatch, while Column (2) presents the results for the mechanism test. Standard errors are clustered at the individual level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Table 11. AI Availability, Missing Skills and Entrepreneurial Entry - Entrepreneurs Only

VARIABLES	(1) P(Entrepreneur)	(2) P(Entrepreneur)
Skill Gap $\times$ Post		1.939 (3.081)
Skill Gap	-24.051*** (6.214)	-24.784*** (6.288)
Observations	67,289,988	67,289,988
Sample	All Entrepreneurs	All Entrepreneurs
User $\times$ Quarter FE	Yes	Yes
User $\times$ Industry 3D FE	Yes	Yes
Quarter $\times$ Industry 3D FE	Yes	Yes
Adj. R sq.	0.096	0.096

Notes: The table reports the results for the test assessing whether AI determines the industry choice of entrepreneurs by bridging human-capital gaps. The sample includes all individuals who eventually became entrepreneurs between 2020Q4 and 2024Q1. The analysis is at the individual-by-industry-by-quarter level. Industries are reported at the three-digit NAICS level. The dataset includes 81 industries over the period 2020Q4–2024Q1. The model used is

$$P(\text{Entrepreneur})_{jit} = \beta_0 + \beta_1 \text{Skill Gap}_{jit} \times \text{Post}_t + \beta_2 \text{Skill Gap}_{jit} + \mu_{ji} + \mu_{jt} + \mu_{it} + \epsilon_{jit}$$

The dependent variable $P(\text{Entrepreneur})_{jit}$ is an indicator for whether the individual becomes an entrepreneur in a specific industry and quarter, multiplied by 10,000 for ease of exposition. Skill Gap is the measure of human-capital gap calculated for each individual, in a given quarter, with respect to a particular industry. We standardize Skill Gap using the pre-shock mean and standard deviation. The binary indicator Post, which is equal to 1 for all quarters after 2022Q3, and 0 otherwise, proxies for the availability of ChatGPT. The specifications include individual-by-industry, individual-by-quarter and industry-by-quarter fixed effects. Column (1) reports results assessing the predictive power for the measure of skill mismatch, while Column (2) presents the results for the mechanism test. Standard errors are clustered at the individual level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Table 12. Effect of AI Compatibility on New Firm Survival After ChatGPT

VARIABLES	(1) P(Survived) ⁽⁴⁾	(2) P(Survived) ⁽⁶⁾	(3) P(Survived) ⁽⁸⁾
AI Assistance $\times$ Post	0.003*** (0.000)	0.004*** (0.000)	0.003*** (0.001)
Observations	1,021,651	945,087	826,018
Industry 3D	Yes	Yes	Yes
Registration Quarter FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Controls	No	No	No
Adj. R sq.	0.014	0.018	0.018

Notes: The table reports the results for the test assessing whether firms established after the arrival of ChatGPT have better survival odds at different horizons (4, 6 and 8 quarters after registration). The analysis is at the firm level and we include all new firms created between 2020Q4 and 2024Q1 in any of the 81 three-digit NAICS industries of interest. The cross-sectional model specification is

$$P(\text{Survived})_f^{(\tau)} = \beta_0 + \beta_1 \text{AI Assistance}_{i(f)} \times \text{Post}_{t(f)} + \mu_{i(f)} + \mu_{t(f)} + \mu_{s(f)} + \epsilon_f$$

The dependent variable $P(\text{Survived})_f^{(\tau)}$ is an indicator equal to 1 if the firm is active τ quarters after establishment. The outcome variable is measured 4, 6, and 8 quarters from establishment. AI Assistance is measured at the three-digit NAICS industry level. The binary indicator Post is equal to 1 for all firms created after 2022Q3, and 0 otherwise. The specification includes industry, quarter of registration and state fixed effects. Robust standard errors are reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Table 13. Effect of AI Compatibility on New Firms' Employment Growth After ChatGPT

VARIABLES	(1) Emp. Growth ⁽⁴⁾	(2) Emp. Growth ⁽⁶⁾	(3) Emp. Growth ⁽⁸⁾
AI Assistance $\times$ Post	0.007*** (0.001)	0.009*** (0.002)	0.006** (0.003)
Observations	1,021,651	945,087	826,018
Industry 3D	Yes	Yes	Yes
Registration Quarter FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Controls	No	No	No
Adj. R sq.	0.013	0.014	0.012

Notes: The table reports the results for the test assessing whether firms established after the arrival of ChatGPT grow faster at different horizons (4, 6 and 8 quarters after registration). The analysis is at the firm level and we include all new firms created between 2020Q4 and 2024Q1 in any of the 81 three-digit NAICS industries of interest. The cross-sectional model specification is

$$\text{Emp. Growth}_f^{(\tau)} = \beta_0 + \beta_1 \text{AI Assistance}_{i(f)} \times \text{Post}_{t(f)} + \mu_{i(f)} + \mu_{t(f)} + \mu_{s(f)} + \epsilon_f$$

The dependent variable $\text{Emp. Growth}_f^{(\tau)}$ is the level change in employment between the quarter of establishment and τ quarters after establishment over the employment level in the quarter of establishment. The outcome variable is measured 4, 6, and 8 quarters from establishment. Employment growth for firms not surviving at a given horizon is coded as -1 (100%). AI Assistance is measured at the three-digit NAICS industry level. The binary indicator Post is equal to 1 for all firms created after 2022Q3, and 0 otherwise. The specification includes industry, quarter of registration and state fixed effects. Robust standard errors are reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Table 14. Effect of AI Compatibility on New Firms' Average Education Level After ChatGPT

VARIABLES	(1) Avg. Education ⁽⁴⁾	(2) Avg. Education ⁽⁶⁾	(3) Avg. Education ⁽⁸⁾
AI Assistance $\times$ Post	0.796*** (0.205)	0.812*** (0.263)	0.719** (0.219)
Observations	556,128	488,318	413,562
Industry 3D	Yes	Yes	Yes
Registration Quarter FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Controls	No	No	No
Adj. R sq.	0.067	0.069	0.072

Notes: The table reports the results for the test assessing whether firms established after the arrival of ChatGPT are able to attract more qualified individuals at different horizons (4, 6 and 8 quarters after registration). The analysis is at the firm level and we include all new firms created between 2020Q4 and 2024Q1 in any of the 81 three-digit NAICS industries of interest. The cross-sectional model specification is

$$\text{Avg. Education}_f^{(\tau)} = \beta_0 + \beta_1 \text{AI Assistance}_{i(f)} \times \text{Post}_{t(f)} + \mu_{i(f)} + \mu_{t(f)} + \mu_{s(f)} + \epsilon_f$$

The dependent variable $\text{Avg. Education}_f^{(\tau)}$ is the average number of years in education attained by workers active at the firm τ quarters after establishment. The outcome variable is measured 4, 6, and 8 quarters from establishment. AI Assistance is measured at the three-digit NAICS industry level. The binary indicator Post is equal to 1 for all firms created after 2022Q3, and 0 otherwise. The specification includes industry, quarter of registration and state fixed effects. Robust standard errors are reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Online Appendix for:

When AI Meets Entrepreneurship: Evidence from the Commercialization of ChatGPT

Table of Contents

- A. Additional Results: Summary Statistics and Analysis
 - A.1. The Effects of AI on Firm Creation
 - * A.1.1. AI Assistance and Industry Characteristics
 - * A.1.2. Richer Fixed Effects Structure
 - * A.1.3. Alternative Treatment Definitions
 - * A.1.4. Heterogeneity and Non-linearity Analysis
 - * A.1.5. Pandemic Exposure and Long-Term Industry Trajectories
 - * A.1.6. Rival Exposure Measures: Construction and Horse Races
 - A.2. Alternative Specifications for Cost of Experimentation Mechanism
 - A.3. Alternative Interpretation of the Capital Allocation Mechanism
- B. Robustness: Chen and Roth (2024)’s Critique
 - B.1. Replicating Effects of AI on Firm Creation Results
 - B.2. Replicating of Cost of Experimentation Mechanism Results
 - B.3. Testing Capital Reallocation Mechanism Results
 - * B.3.1. Log Specification Estimation
 - * B.3.2. Estimating Investments with PPML
 - * B.3.3. Two-Part Hurdle Specification: Extensive Versus Intensive Margin
- C. Additional Information on Revelio Labs
 - C.1. Sourcing of Revelio Labs Data
 - C.2. Representativeness of Revelio Labs Data
 - * C.2.1. Industry Distribution of New Firms in RL vs. BDS
 - * C.2.2. New Firm Counts and Trends: RL vs. BDS
 - * C.2.3. Summary and Caveats
 - C.3. Measures Creation Using Revelio Labs
- D. Generation of Measures Using LLMs
 - D.1. Textual Data Sources
 - D.2. Workflow Infrastructure
 - * D.2.1. LLM Layer 1: Task Extraction
 - * D.2.2. LLM Layer 2: Task Deduplication

- * D.2.3. LLM Layer 3: Assessing Industry Relevance and Adapting Task Description to Industry
 - * D.2.4. LLM Layer 4: Task Scoring
- E. Survey of New and Aspiring Entrepreneurs
 - E.1. Survey Instrument
 - E.2. Constructing the Analysis Sample
 - E.3. Sample Composition
- F. In Depth Literature Review
 - F.1. A Closer Look at the Impact of AI
 - F.2. A Closer Look at Entrepreneurship as a Process of Experimentation
- G. Theoretical Framework: AI and Entrepreneurial Entry
 - G.1. Overview
 - G.2. Model Setup
 - G.3. Heterogeneous Entrepreneurial Capabilities
 - G.4. Entrepreneurship Decision
 - G.5. Equilibrium
 - G.6. Main Results
 - G.7. Discussion

APPENDIX A. ADDITIONAL RESULTS: SUMMARY STATISTICS AND ANALYSIS

A.1. The Effects of AI on Firm Creation. The credibility of our main findings hinges on the validity of the identifying assumptions and on the robustness of our results to alternative specifications and interpretations. We therefore conduct a series of complementary analyses.

First, we verify that the AI Assistance measure is not simply capturing pre-existing industry characteristics by examining its relationship with a wide range of pre-period fundamentals.

Next, we assess the robustness of our estimates to alternative fixed-effects structures and treatment definitions, and we explore whether the effects display any meaningful nonlinearities or heterogeneity across industries.

We then show that our results are not driven by pandemic-era entry dynamics—absorbed through continuous measures of each industry’s entry boom—or by secular long-term growth in value added, which could otherwise confound the interpretation of post-ChatGPT dynamics.

Finally, we detail the construction of the rival exposure measures—teleworkability, IT capital intensity, and the pre-treatment growth trend—raced against AI Assistance in Subsection 4.3, and report the fully specified horse races on both estimation windows.

A.1.1. AI Assistance and Industry Characteristics. A potential concern is that our measure of AI Assistance may simply capture pre-existing industry fundamentals—such as profitability, R&D intensity, or investment activity—rather than a distinct dimension of entrepreneurial relevance. If that were the case, the observed increase in firm creation in high-AI Assistance industries could reflect differences in baseline industry strength rather than the extent to which AI tools facilitate entrepreneurial tasks.

To address this concern, we examine how AI Assistance correlates with standard industry characteristics in the pre-ChatGPT period. Table OA1 reports results using three alternative definitions of the pre-shock window: Column (1) uses only 2022Q3 values, Column (2) averages over the last four quarters (2021Q4–2022Q3), and Column (3) averages over the last eight quarters (2020Q4–2022Q3). Across all definitions, the correlations remain modest, ranging from about -0.18 to 0.35 , and are highly stable across panels. The strongest positive associations are with the number of investments and total investment, as well as with liquidity ratios. In contrast, AI Assistance is essentially uncorrelated with profitability, R&D intensity, and capex intensity, and is negatively but weakly related to market leverage.

Table OA1. Correlation of AI Assistance Score with Industry Characteristics

	(1) AI Assistance 2022Q3	(2) AI Assistance 2021Q4–2022Q3 Average	(3) AI Assistance 2020Q4–2022Q3 Average
Number of Investments	0.310	0.312	0.313
Total Investments (\$000)	0.334	0.319	0.339
Average Investment (\$000)	-0.025	0.021	0.128
ROA	0.028	0.025	0.098
R&D Intensity	0.024	-0.098	-0.093
Capex Intensity	0.073	0.092	0.075
Liquidity Ratio	0.260	0.317	0.357
Market Leverage Ratio	-0.138	-0.143	-0.178

Notes: This table presents the correlation between the AI Assistance measure and various industry characteristics, calculated on different time windows. The industry characteristics are the number of early stage equity investments, the total invested amount (in thousands of dollars), the average invested amount (in thousands of dollars), the Return on Assets (ROA, income before extraordinary items divided by total assets), R&D Intensity (R&D expenses over revenue), CapEx Intensity (capital expenditures over revenue), Liquidity Ratio (cash holdings over total assets), and Market Leverage Ratio (the sum of long-term and current debt divided by the market value of equity). Column (1) establishes industry characteristics using 2022Q3 values. Column (2) establishes the characteristics using the average industry value over the period 2021Q4 to 2022Q3. Column (3) establishes the characteristics using the average industry value over the period 2020Q4 to 2022Q3.

As an additional placebo-outcome check, we test whether the increase in firm creation in high-AI-Assistance industries coincides with differential changes in other industry fundamentals after ChatGPT’s release. To do so, we regress pre-ChatGPT industry characteristics—ROA, R&D intensity, capex intensity, liquidity ratios, leverage ratios, and investment variables—on the interaction between AI Assistance and Post, controlling for industry and time fixed effects. This specification examines whether high-AI industries experienced disproportionate shifts in these outcomes in the post-ChatGPT period.

As seen in Table OA2, the estimated coefficients are small and statistically insignificant across all outcomes except for the liquidity ratio, where we find a negative but economically modest association (-0.012). This coefficient is an order of magnitude smaller than the effect of AI Assistance $\times$ Post on new firm creation (0.105 in Table 4), implying that any relationship between AI Assistance and liquidity is negligible in economic terms. Moreover, the negative sign suggests that industries more compatible with AI exhibited slightly lower liquidity ratios, contrary to what one would expect if AI Assistance merely captured financially stronger sectors.

Table OA2. Effect of AI Assistance on Other Outcomes After the Release of ChatGPT

VARIABLES	(1) ROA	(2) R&D Intensity	(3) Capex Intensity	(4) Liquidity Ratio	(5) Market Leverage Ratio
AI Assistance $\times$ Post	-0.011 (0.008)	0.037 (0.032)	-0.015 (0.022)	-0.012*** (0.003)	0.003 (0.019)
Observations	1,134	1,134	1,134	1,134	1,134
Industry 3D FE	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No	No	No	No
Controls	No	No	No	No	No
Adj. R. sq.	0.003	0.253	-0.077	0.919	0.800

Notes: This table presents the results for the test assessing whether the treatment used in the main analysis affects other industry level outcomes that are not entrepreneurship specific. The design is a difference-in-differences model at the industry-by-quarter level with continuous treatment. The sample includes 81 three-digit NAICS industries over 2020Q4–2024Q1. The empirical specification is

$$Y_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is either the Return on Assets (ROA, income before extraordinary items divided by total assets) in Column (1), R&D Intensity (R&D expenses over revenue) in Column (2), CapEx Intensity (capital expenditures over revenue) in Column (3), Liquidity Ratio (cash holdings over total assets) in Column (4), and Market Leverage Ratio (the sum of long-term and current debt divided by the market value of equity) in Column (5). Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. All specifications include industry and quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Overall, these results indicate that AI Assistance is not simply a proxy for standard industry fundamentals such as size, profitability, or investment intensity, but rather captures a distinct dimension of the entrepreneurial process—namely, the extent to which AI can provide support for early-stage entrepreneurial tasks.

A.1.2. Richer Fixed Effects Structure. A primary identification concern in our analysis is the presence of unobserved shocks at the industry-by-quarter level, which could simultaneously influence firm creation and correlate with AI Assistance—omitted variable bias. Ideally, we would include industry-by-quarter fixed effects to completely absorb any shocks specific to individual industries each quarter. However, this approach would exhaust degrees of

freedom and prevent identification of the treatment effect. Therefore, we adopt sector-by-quarter fixed effects as a robustness test, where sectors are defined as industry groups at the two-digit NAICS level. While not fully achieving the ideal scenario, these fixed effects significantly mitigate concerns by removing broader sector-level shocks each quarter, ensuring that the variation we exploit for identification arises predominantly from differences across finer-grained (three-digit) industries within the same sector and quarter. We present the results in Table OA3 and Figure OA1.

The robustness of our estimated coefficient, remaining positive and significant, enhances confidence that the observed effects are driven by industry-specific treatment intensity rather than broader sectoral shocks (Table OA3).

Table OA3. Effect of AI Assistance on Firm Creation After the Release of ChatGPT Using Richer Fixed Effects Structure

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms
AI Assistance $\times$ Post	0.102** (0.046)	0.096** (0.046)
Observations	1,078	1,078
Industry 3D FE	Yes	Yes
Quarter FE	No	No
Industry 2D $\times$ Quarter FE	Yes	Yes
Controls	No	Yes
Adj. R. sq.	0.980	0.980

Notes: This table presents results testing the effect of AI’s arrival on firm creation using a richer fixed effects structure. The design is a difference-in-differences model at the industry-by-quarter level with continuous treatment. The sample includes 81 three-digit NAICS industries over 2020Q4–2024Q1. The empirical specification is


$$\text{IHS(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. All specifications include industry and sector-by-quarter fixed effects. Column (2) additionally includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

The event study analysis with the richer fixed effects structure (sector-by-quarter fixed effects) provides important insights (Figure OA1). Specifically, the coefficients on the lead terms—those representing periods prior to ChatGPT’s release—are statistically insignificant,

consistent with the validity of the parallel trends assumption. This indicates that industries with higher AI compatibility were not following differential pre-existing trends before the actual treatment event, strengthening the causal interpretation of our estimates.

However, coefficients on the lag terms—representing post-treatment periods—are generally positive but statistically insignificant. It is important to note that this statistical insignificance does not necessarily imply the absence of a meaningful economic effect. Instead, it arises primarily from the demanding structure imposed by the highly granular fixed effects. By including sector-by-quarter fixed effects, we remove substantial variation from the data at the broader sectoral level. Consequently, identification of treatment effects relies solely on variation across fine-grained (three-digit NAICS) industries within each two-digit sector and each specific quarter. Given that some two-digit sectors contain very few three-digit industries—occasionally only two or three—the remaining variation available for estimating coefficients is severely limited. This sparsity of observations within each cell inflates standard errors and results in less stable, statistically noisier estimates.

Figure OA1. Event Study of AI Assistance on Firm Creation with Richer Fixed Effects Structure

Notes: The figure presents the event study for the sample of 81 three-digit NAICS industries in our analysis, over the time period 2020Q4–2024Q1. We plot the estimates from the dynamic difference-in-differences model for the interaction of the quarter-year dummies with the AI Assistance_{*i*} from the specification

$$\text{IHS}(\text{N. New Firms})_{it} = \beta_0 + \sum_{t \neq 2022q3} \beta_t \text{AI Assistance}_i \times 1\{D = t\}_t + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. The specifications include both industry and sector-by-quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level. The specification in Panel (a) does not include controls. The specification in Panel (b) includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage.

Despite the statistical challenges entailed using a very rich fixed effects structure, the positive and significant coefficients in Table OA3, paired with the lack of pre-trends displayed in the plots of Figure OA1 provide reassuring qualitative support which align with our primary findings.

A.1.3. Alternative Treatment Definitions. The validity of our causal estimates relies critically on accurately measuring and defining treatment intensity, specifically industries' compatibility with generative AI technologies. Our main analysis in Section 4 employs a continuous measure (the AI Assistance Score), but this approach inherently risks being affected by measurement errors, sensitivity to extreme observations, and potential biases due to subtle inaccuracies or misclassification.

To address these potential concerns and to validate the robustness of our primary findings, we additionally test binary versions of our treatment definition—classifying industries simply as high or low compatibility based on whether their AI Assistance exceeds the mean

or median. By using these binary definitions, we effectively reduce the sensitivity of our estimates to minor measurement errors, mitigate the influence of extreme values that might disproportionately affect our continuous measure, and clearly distinguish treated from control groups.

This simplified binary approach also helps alleviate attenuation bias that can arise from measurement noise in continuous variables, while simultaneously serving as an important check against assumptions of strict linearity or monotonicity in the treatment effect.

We re-run the specification in Equation 6, now using the two binary version of treatment. Table OA4 and Table OA5 presents the results. The positive and significant coefficients confirm that our findings are not dependent on a continuous scoring system, indicating robustness of our qualitative conclusions.

The larger estimates from binary definitions potentially suggest the presence of nonlinearities. Specifically, the effect of AI compatibility may be more pronounced once a certain threshold of compatibility is reached. In other words, industries that surpass a certain critical level of AI compatibility might experience disproportionately stronger entrepreneurial responses after ChatGPT's release compared to incremental increases at lower levels of compatibility.

Table OA4. Effect of Above-Mean AI Assistance on Firm Creation After the Release of ChatGPT

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms
AI Assistance Above Mean $\times$ Post	0.128** (0.055)	0.112** (0.053)
Observations	1,134	1,134
Industry 3D FE	Yes	Yes
Quarter FE	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No
Controls	No	Yes
Adj. R. sq.	0.981	0.981

Notes: This table presents results testing the effect of AI's arrival on firm creation using an alternative definition of treatment status. The design is a difference-in-differences model at the industry-by-quarter level with binary treatment. The sample includes 81 three-digit NAICS industries over 2020Q4–2024Q1. The empirical specification is

$$\text{IHS(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance Above Mean}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment status is proxied by AI Assistance Above Mean, an indicator equal to 1 if the AI Assistance Score of the industry is above the sample mean, and 0 otherwise. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. All specifications include industry and quarter fixed effects. Column (2) additionally includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Table OA5. Effect of Above-Median AI Assistance on Firm Creation After the Release of ChatGPT


VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms
AI Assistance Above Median $\times$ Post	0.143*** (0.054)	0.131** (0.052)
Observations	1,134	1,134
Industry 3D FE	Yes	Yes
Quarter FE	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No
Controls	No	Yes
Adj. R. sq.	0.981	0.981

Notes: This table presents results testing the effect of AI's arrival on firm creation using an alternative definition of treatment status. The design is a difference-in-differences model at the industry-by-quarter level with binary treatment. The sample includes 81 three-digit NAICS industries over 2020Q4–2024Q1. The empirical specification is

$$\text{IHS(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance Above Median}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment status is proxied by AI Assistance Above Median, an indicator equal to 1 if the AI Assistance Score of an industry is above the sample median, and 0 otherwise. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. All specifications include industry and quarter fixed effects. Column (2) additionally includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.


In Figures OA2 and OA3 we present the event plot study. The results are largely unchanged from what we find in the main analysis.

Figure OA2. Event Study of AI Assistance Above Mean Group on Firm Creation

Notes: The figure presents the event study for the sample of 81 three-digit NAICS industries in our analysis, over the time period 2020Q4 to 2024Q1. We plot the estimates from the dynamic difference-in-differences model for the interaction of the quarter-year dummies with the AI Assistance Above Mean_{*i*} from the specification

$$\text{IHS}(\text{N. New Firms})_{it} = \beta_0 + \sum_{t \neq 2022Q3} \beta_t \text{AI Assistance Above Mean}_i \times 1\{D = t\}_t + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its binarized form (above versus below the sample mean). The specifications include both industry and quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level. The specification in Panel (a) does not include controls. The specification in Panel (b) includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage.

Figure OA3. Event Study of AI Assistance Above Median Group on Firm Creation

Notes: The figure presents the event study for the sample of 81 three-digit NAICS industries in our analysis, over the time period 2020Q4 to 2024Q1. We plot the estimates from the dynamic difference-in-differences model for the interaction of the quarter-year dummies with the AI Assistance Above Median_i from the specification

$$\text{IHS}(\text{N. New Firms})_{it} = \beta_0 + \sum_{t \neq 2022q3} \beta_t \text{AI Assistance Above Median}_i \times 1\{D = t\}_t + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its binarized form (above versus below the sample median). The specifications include both industry and quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level. The specification in Panel (a) does not include controls. The specification in Panel (b) includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage.

The consistent and significant positive results obtained from these binary treatment robustness checks reinforce the credibility and stability of our main results, by showing that they are not an artifact stemming from the choice of using a continuous variable as treatment definition.

A.1.4. Heterogeneity and Non-linearity Analysis. To further validate our findings and to explore possible nonlinearities, we perform a heterogeneity analysis. The motivation behind this analysis is twofold. First, our main analysis, employing a continuous measure of AI compatibility (AI Assistance Score), implicitly assumes a smooth, linear relationship between treatment intensity and firm creation. However, economically, the benefits of generative AI tools might differ significantly at higher levels of compatibility. Second, by explicitly examining subsets of industries—such as those in the top quartile or decile groups—we directly address concerns related to potential measurement error or misclassification bias in our continuous measure. Clearly defined groups based on explicit thresholds reduce sensitivity

to subtle measurement inaccuracies, thus helping to clarify the robustness of our estimated effects.

In the first set of heterogeneity tests we redefine treatment by grouping industries based on whether their AI Assistance Score belongs to the top quartile of the distribution. Table OA6 strongly supports the existence of meaningful heterogeneity and nonlinearities in the treatment effect. Specifically, we find that industries in the top quartile of AI compatibility exhibit significantly larger effects compared to the overall continuous measure in Table 4. This result indicates that the entrepreneurial response to generative AI is notably stronger once a critical threshold of compatibility is surpassed.

Table OA6. Effect of Top Quartile AI Assistance on Firm Creation After the Release of ChatGPT

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms
AI Assistance Top Quartile $\times$ Post	0.240*** (0.057)	0.228*** (0.058)
Observations	1,134	1,134
Industry 3D FE	Yes	Yes
Quarter FE	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No
Controls	No	Yes
Adj. R. sq.	0.982	0.982

Notes: The table presents results testing the functional shape of the effect of AI's arrival on firm creation. The design is a difference-in-differences model at the industry-quarter level with a binary treatment. The sample includes 81 three-digit NAICS industries over the period 2020Q4–2024Q1. The empirical specification is


$$\text{IHS(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance Top Quartile}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment status is proxied by AI Assistance Top Quartile, an indicator equal to 1 if the AI Assistance Score of an industry belongs to the top quartile of the sample distribution, and 0 otherwise. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. All specifications include industry and quarter fixed effects. Column (2) additionally includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

In Figure OA4 we plot the event study for the specification where we use as definition of treatment belonging to the top quartile of the AI Assistance Score distribution and not only we do not observe diverging pre-trends, much like in the main analysis, but the significance

of the lag coefficients is strong and the effects are incremental over time, consistent with idea of progressive take up.

Figure OA4. Event Study of Top Quartile AI Assistance on Firm Creation



Notes: The figure presents the event study for the sample of 81 three-digit NAICS industries in our analysis, over the time period 2020Q4 to 2024Q1. We plot the estimates from the dynamic difference-in-differences model for the interaction of the quarter-year dummies with the AI Assistance Top Quartile_{*i*} from the specification

$$\text{IHS(N. New Firms)}_{it} = \beta_0 + \sum_{t \neq 2022q3} \beta_t \text{AI Assistance Top Quartile}_i \times 1\{D = t\} + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$


The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. The specifications include both industry and quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level. The specification in Panel (a) does not include controls. The specification in Panel (b) includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&DIntensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage.

Additionally, we plot the estimated treatment effects by decile to investigate how the impact of AI compatibility on entrepreneurial activity varies across the entire distribution of AI Assistance Scores. By dividing industries into ten equally-sized groups based on their AI compatibility, we can carefully assess whether the entrepreneurial responses increase smoothly and monotonically with higher compatibility levels or if distinct thresholds exist beyond which these effects intensify significantly. This analysis allows us to identify potential nonlinearities and provides valuable insights into whether the benefits of generative AI accrue gradually or primarily above specific compatibility thresholds.

Econometrically speaking, we run our difference-in-differences regression again, but this time we replace continuous AI Assistance with a set of dummy variables for each decile. For instance, the 1st decile serves as the baseline reference group, and we estimate separate coefficients for the 2nd decile through the 10th decile, each interacted with the post-treatment

indicator. Thus, each estimated coefficient captures how industries in a particular decile of AI compatibility differ in their entrepreneurial responses after ChatGPT’s release relative to the lowest decile.

The decile analysis in Figure OA5 reveals a clear and systematic relationship between AI compatibility and entrepreneurial activity. Specifically, we observe that industries situated in lower deciles, representing those with comparatively low AI compatibility, exhibit modest or negligible increases in firm creation following ChatGPT’s release. In contrast, as we move progressively towards higher deciles, the estimated effects become increasingly pronounced, demonstrating a clear and consistent monotonic relationship. Industries in the highest deciles, characterized by substantially greater AI compatibility, show considerably larger and statistically significant impacts on new firm creation. These findings underscore the presence of meaningful nonlinearities, indicating that entrepreneurial benefits associated with generative AI intensify substantially beyond certain critical compatibility thresholds. Thus, this decile analysis not only corroborates our main results but also enriches our understanding by highlighting nuanced dynamics that inform strategic decision-making and targeted policy interventions.

Figure OA5. Effect of the Release of ChatGPT by AI Assistance Deciles

Notes: The figure presents the heterogeneity study for the sample of 81 three-digit NAICS industries in our analysis, over the time period 2020Q4–2024Q1. We plot the estimates from a compact difference-in-differences for the vector of decile based categorical groups of AI Assistance

$$\text{IHS(N. New Firms)}_{it} = \beta_0 + \sum_{d=2}^{10} \beta_d \text{AI Assistance Decile}_i^d \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity groups are created by binning industries into ten equally-sized deciles according to their AI Assistance Score, with the first decile (lowest AI compatibility) serving as the baseline category. The specification includes both industry and quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level. The specification in Panel (a) does not include controls. The specification in Panel (b) includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median Capex Intensity, median Liquidity Ratio and median Market Leverage.

Taken together, these heterogeneity analyses significantly enhance the interpretability and economic relevance of our primary findings by explicitly revealing that the entrepreneurial impact of generative AI compatibility does not merely increase uniformly across all industries. Instead, the results clearly indicate the presence of threshold-like or nonlinear effects, where the benefits of generative AI tools become especially pronounced beyond certain critical levels of compatibility. This finding has crucial implications: it suggests that policy interventions, resource allocations, or strategic industry investments aimed at fostering entrepreneurial activity through generative AI may be particularly effective when targeted at industries that surpass specific thresholds of AI compatibility, rather than applied broadly across all sectors. Furthermore, the monotonic increase in estimated effects across deciles provides a transparent and intuitive demonstration of a dose-response relationship, meaning industries progressively better suited to leveraging generative AI experience correspondingly larger entrepreneurial gains. These insights are economically meaningful as they clearly delineate where the potential for generative AI-driven entrepreneurial growth is highest,

thereby offering practical guidance for policymakers and industry stakeholders seeking to maximize the economic benefits of AI technology adoption.

It is important to emphasize that the presence of nonlinearities revealed by our heterogeneity analyses does not contradict or invalidate our primary empirical strategy, which employs a continuous exposure variable (the AI Assistance Score). Rather, the identification of nonlinear effects enriches our understanding by explicitly highlighting thresholds and nuanced dynamics that are inherent to many economic relationships. The continuous measure remains a valid and informative way to capture the general and overall relationship between generative AI compatibility and entrepreneurial outcomes. The robustness and heterogeneity analyses complement this approach by demonstrating that entrepreneurial responses to AI compatibility intensify disproportionately at higher compatibility levels. Far from invalidating our continuous specification, these findings clarify and deepen its economic interpretation, offering valuable insights for policy targeting and strategic investment decisions.

A.1.5. Pandemic Exposure and Long-Term Industry Trajectories. A natural concern is that the post-ChatGPT increase in firm creation we document could be driven by lingering dynamics from the COVID-19 pandemic. [Haltiwanger \(2022\)](#) shows that the surge in firm formation during COVID was highly concentrated in a small set of three-digit NAICS industries; if the industries that expanded rapidly during the pandemic also score highly on our AI Assistance measure, our baseline results could conflate pandemic-era entry dynamics with the effects of AI. Subsection 4.3 addresses this concern by measuring each industry’s pandemic exposure as a continuous dose—the entry boom over 2020Q3–2021Q4 relative to a 2018Q4–2019Q4 baseline, computed independently from Revelio and from the Census Business Formation Statistics (BFS)—and absorbing a full set of interactions between each dose and the quarter indicators. This subsection reports the complete results.

Tables [OA7](#) and [OA8](#) present the estimates on the baseline and extended estimation windows. Across every specification, the AI Assistance coefficient is at least as large as the baseline estimated on the same sample: absorbing the Revelio dose, the BFS dose, or both jointly leaves the estimate between 0.100 and 0.107 on the main window, against a baseline of 0.099, and between 0.102 and 0.118 on the extended window, against a baseline of 0.109.

Table OA7. Robustness to Continuous Measures of Pandemic Exposure

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms	(3) IHS N. New Firms	(4) IHS N. New Firms	(5) IHS N. New Firms
AI Assistance $\times$ Post	0.099*** (0.025)	0.100*** (0.025)	0.105*** (0.025)	0.107*** (0.026)	0.105*** (0.025)
COVID Dose (Revelio) $\times$ Quarter	No	Yes	No	Yes	No
COVID Dose (BFS) $\times$ Quarter	No	No	Yes	Yes	No
COVID Dose (BFS, NFD alloc.) $\times$ Quarter	No	No	No	No	Yes
BFS-matched sample (79 industries)	Yes	Yes	Yes	Yes	Yes
Observations	1,106	1,106	1,106	1,106	1,106
Industry 3D FE	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No	No	No	No
Controls	Yes	Yes	Yes	Yes	Yes
Adj. R. sq.	0.982	0.983	0.982	0.983	0.982

Notes: This table presents results absorbing pandemic exposure measured as a continuous dose. The design is a difference-in-differences model at the industry-by-quarter level with continuous treatment, over the baseline estimation window 2020Q4–2024Q1. The empirical specification is

$$\begin{aligned} \text{IHS(N. New Firms)}_{it} = & \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \sum_{t \neq 2022q3} \delta_t \text{COVID Dose}_i \times 1\{D = t\}_t \\ & + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it} \end{aligned}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. The COVID dose measures each industry’s entry boom over 2020Q3–2021Q4 relative to a 2018Q4–2019Q4 baseline. Column (1) presents the baseline results without a dose. Column (2) enters the dose computed from Revelio Labs. Column (3) enters the dose computed from Census Business Formation Statistics (BFS) applications. Column (4) enters both doses jointly, each as a full set of quarter interactions. For completeness, Column (5) recomputes the BFS dose allocating each two-digit sector’s residual of applications not identified to any three-digit subsector across its named subsectors in proportion to their own application counts; the dose in Columns (3)–(4) is constructed from named subsectors only. The sample comprises 79 of the 81 industries in the main analysis: NAICS 112 and 814 are excluded because the Business Formation Statistics omit them from the Business Applications universe, so no BFS dose can be constructed for them. The restriction applies uniformly to all columns, including the baseline in Column (1), so that estimates are comparable across columns within a constant sample. All specifications include industry and quarter fixed effects and the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Table OA8. Robustness to Continuous Measures of Pandemic Exposure: Extended Window


VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms	(3) IHS N. New Firms	(4) IHS N. New Firms	(5) IHS N. New Firms
AI Assistance $\times$ Post	0.109*** (0.027)	0.102*** (0.028)	0.117*** (0.027)	0.110*** (0.028)	0.118*** (0.028)
COVID Dose (Revelio) $\times$ Quarter	No	Yes	No	Yes	No
COVID Dose (BFS) $\times$ Quarter	No	No	Yes	Yes	No
COVID Dose (BFS, NFD alloc.) $\times$ Quarter	No	No	No	No	Yes
BFS-matched sample (79 industries)	Yes	Yes	Yes	Yes	Yes
Observations	1,738	1,738	1,738	1,738	1,738
Industry 3D FE	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No	No	No	No
Controls	Yes	Yes	Yes	Yes	Yes
Adj. R. sq.	0.984	0.985	0.984	0.985	0.984

Notes: This table presents results absorbing pandemic exposure measured as a continuous dose. The design is a difference-in-differences model at the industry-by-quarter level with continuous treatment, over the extended estimation window 2018Q4–2024Q1. The empirical specification is

$$\begin{aligned} \text{IHS(N. New Firms)}_{it} = & \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \sum_{t \neq 2022q3} \delta_t \text{COVID Dose}_i \times 1\{D = t\}_t \\ & + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it} \end{aligned}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. The COVID dose measures each industry's entry boom over 2020Q3–2021Q4 relative to a 2018Q4–2019Q4 baseline. Column (1) presents the baseline results without a dose. Column (2) enters the dose computed from Revelio Labs. Column (3) enters the dose computed from Census Business Formation Statistics (BFS) applications. Column (4) enters both doses jointly, each as a full set of quarter interactions. For completeness, Column (5) recomputes the BFS dose allocating each two-digit sector's residual of applications not identified to any three-digit subsector across its named subsectors in proportion to their own application counts; the dose in Columns (3)–(4) is constructed from named subsectors only. The sample comprises 79 of the 81 industries in the main analysis: NAICS 112 and 814 are excluded because the Business Formation Statistics omit them from the Business Applications universe, so no BFS dose can be constructed for them. The restriction applies uniformly to all columns, including the baseline in Column (1), so that estimates are comparable across columns within a constant sample. All specifications include industry and quarter fixed effects and the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Figure OA6 presents the event-study version of the same exercise. The path of the AI Assistance coefficients is unchanged whether estimated unconditionally or conditional on either dose, and in particular the post-2022Q4 break does not flatten.

Figure OA6. Event Study Conditional on Continuous Pandemic Exposure

Notes: The figure presents the event study for the sample of 79 three-digit NAICS industries for which a COVID dose can be constructed, over the extended time period 2018Q4–2024Q1. We plot the estimates from the dynamic difference-in-differences model for the interaction of the quarter-year dummies with the AI Assistance_{*i*} from the specification

$$\begin{aligned} \text{IHS(N. New Firms)}_{it} = & \beta_0 + \sum_{t \neq 2022q3} \beta_t \text{AI Assistance}_i \times 1\{D = t\}_t \\ & + \sum_{t \neq 2022q3} \delta_t \text{COVID Dose}_i \times 1\{D = t\}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it} \end{aligned}$$


The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. The series marked Unconditional omits the COVID dose interactions; the series marked Cond. on Revelio dose enters the dose computed from Revelio Labs; the series marked Cond. on BFS dose enters the dose computed from Census Business Formation Statistics applications. The specifications include both industry and quarter fixed effects and the controls listed in Table OA7. Standard errors are clustered at the three-digit NAICS industry level; bars represent 95% confidence intervals.

A second concern is that the set of high AI Assistance industries could coincide with industries that, independently of AI, experienced strong secular growth in value added between 1970 and 2010. In other words, heightened firm creation in these industries might mechanically reflect long-run industry trajectories rather than a distinct dimension of entrepreneurial relevance. The worry is that industries with persistent increases in value added over several decades are precisely those that we rank highly on AI Assistance, so that our

results could be interpreted as a continuation of long-run trends rather than a causal effect of the commercialization of AI.

To address this concern, we construct a measure of long-run industry growth based on the Bureau of Economic Analysis (BEA) data on Gross Value Added by industry. Specifically, for each three-digit NAICS industry we compute the average annual log growth rate in real value added and the cumulative growth between the 1970s and the 2010s. We then correlate this measure of long-run growth with our AI Assistance ranking. The correlation for both is modest, at 0.39 (0.387 and 0.392), implying that only about 15% of the variation in AI Assistance across industries can be linearly explained by differences in long-run value added growth. This finding suggests that while industries that have grown more over the past decades tend to be somewhat more AI-assistable, the majority of the cross-industry variation in AI Assistance reflects other dimensions that are orthogonal to long-run secular growth trends. Additionally, we scatter both measures of growth in value added against the AI Assistance variable in Figure OA7 to show that no striking pattern emerges.

Figure OA7. Correlation between AI Assistance and Long-Term Growth in Value Added



Notes: This set of plots displays the correlation between the AI Assistance measure and estimates of long-term growth in value added. Plot (a) scatters the AI Assistance measure against the cumulative growth in value added over the period 1970-2010, while Plot (b) scatters it against the annualized log growth over the same period.

Taken together, this evidence alleviates the concern that our AI Assistance measure is simply a relabeling of historically “winning” industries. Moreover, our empirical design—with industry fixed effects and pre-trend tests—ensures that we identify the discrete shift in entrepreneurial activity that follows the release of ChatGPT, rather than a continuation of long-run trajectories.

A.1.6. *Rival Exposure Measures: Construction and Horse Races.* This subsubsection details the construction of the three rival exposure measures raced against AI Assistance in Subsection 4.3 and reports the fully specified horse races. For teleworkability and IT capital intensity, each table presents four columns: the rival alone, AI Assistance alone, both jointly, and AI Assistance residualized on the rival in the industry cross-section—the residualized column matters because the joint column is estimated off correlated regressors, and the residualized coefficient is what survives once the rival’s variation is removed entirely. For the pre-treatment growth trend, the table adds the nonparametric, trend-absorption, and long-window variants. Each horse race is reported on both estimation windows—the baseline window 2020Q4–2024Q1 and the extended window 2018Q4–2024Q1—as separate tables, and all rival measures are standardized over the estimation sample.

Teleworkability is the occupation-level classification of [Dingel and Neiman \(2020\)](#), which codes whether each occupation’s work can be performed entirely at home based on its O*NET task content. We aggregate the classification to three-digit NAICS 2022 industries using OEWS May 2018 employment weights, so that the measure captures the employment-weighted share of each industry’s work that is remote-capable, fixed at pre-pandemic occupational structures. Tables [OA9](#) and [OA10](#) report the horse race on the baseline and extended estimation windows.

Table OA9. Horse Race Against Teleworkability

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms	(3) IHS N. New Firms	(4) IHS N. New Firms
AI Assistance $\times$ Post		0.100*** (0.025)	0.090*** (0.032)	
Teleworkability $\times$ Post	0.069** (0.027)		0.018 (0.030)	
AI Assistance (resid.) $\times$ Post				0.088*** (0.033)
Observations	1,134	1,134	1,134	1,134
Industry 3D FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No	No	No
Controls	Yes	Yes	Yes	Yes
Adj. R. sq.	0.982	0.982	0.982	0.982

Notes: This table presents results racing AI Assistance against Teleworkability. The design is a horse race difference-in-differences model at the industry-by-quarter level with continuous treatment. The sample includes 81 three-digit NAICS industries over the baseline estimation window 2020Q4–2024Q1. Each horse race regression uses the empirical specification

$$\text{IHS(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \beta_2 \text{Teleworkability}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. Teleworkability is constructed from the occupation-level classification of [Dingel and Neiman \(2020\)](#), aggregated to three-digit NAICS 2022 using OEWS employment weights, and standardized over the estimation sample. Column (1) enters Teleworkability alone. Column (2) enters AI Assistance alone. Column (3) enters both measures jointly. In Column (4), AI Assistance is residualized on Teleworkability in the industry cross-section and the standardized residual replaces AI Assistance. All specifications include industry and quarter fixed effects and the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Table OA10. Horse Race Against Teleworkability: Extended Window

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms	(3) IHS N. New Firms	(4) IHS N. New Firms
AI Assistance $\times$ Post		0.112*** (0.027)	0.087** (0.034)	
Teleworkability $\times$ Post	0.095*** (0.028)		0.044 (0.031)	
AI Assistance (resid.) $\times$ Post				0.086** (0.036)
Observations	1,782	1,782	1,782	1,782
Industry 3D FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No	No	No
Controls	Yes	Yes	Yes	Yes
Adj. R. sq.	0.983	0.984	0.984	0.983

Notes: This table presents results racing AI Assistance against Teleworkability. The design is a horse race difference-in-differences model at the industry-by-quarter level with continuous treatment. The sample includes 81 three-digit NAICS industries over the extended estimation window 2018Q4–2024Q1. Each horse race regression uses the empirical specification

$$\text{IHS(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \beta_2 \text{Teleworkability}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. Teleworkability is constructed from the occupation-level classification of [Dingel and Neiman \(2020\)](#), aggregated to three-digit NAICS 2022 using OEWS employment weights, and standardized over the estimation sample. Column (1) enters Teleworkability alone. Column (2) enters AI Assistance alone. Column (3) enters both measures jointly. In Column (4), AI Assistance is residualized on Teleworkability in the industry cross-section and the standardized residual replaces AI Assistance. All specifications include industry and quarter fixed effects and the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

IT capital intensity is each industry's IT share of the net capital stock, computed from the Bureau of Economic Analysis detailed nonresidential fixed-asset estimates and fixed at 2018 values, before the estimation window opens. The IT stock sums eleven BEA asset series: eight computers-and-peripherals equipment categories (mainframes, personal computers, direct access storage devices, printers, terminals, tape drives, storage devices, and system integrators; BEA codes EP1A–EP1H) and three software categories from intellectual property products (prepackaged, custom, and own-account; ENS1–ENS3). The denominator is the industry's total private nonresidential fixed assets—equipment, structures, and intellectual property products (EQ00, ST00, IP00)—so the measure is IT's share of the industry's

entire capital stock. The measure correlates with AI Assistance at 0.63 in the industry cross-section, making it the most demanding rival in the set. Tables OA11 and OA12 report the horse race on the baseline and extended estimation windows.

Table OA11. Horse Race Against IT Capital Intensity

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms	(3) IHS N. New Firms	(4) IHS N. New Firms
AI Assistance $\times$ Post		0.100*** (0.025)	0.093*** (0.032)	
IT Capital Intensity $\times$ Post	0.068*** (0.024)		0.010 (0.031)	
AI Assistance (resid.) $\times$ Post				0.091*** (0.033)
Observations	1,134	1,134	1,134	1,134
Industry 3D FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No	No	No
Controls	Yes	Yes	Yes	Yes
Adj. R. sq.	0.981	0.982	0.982	0.982

Notes: This table presents results racing AI Assistance against IT Capital Intensity. The design is a horse race difference-in-differences model at the industry-by-quarter level with continuous treatment. The sample includes 81 three-digit NAICS industries over the baseline estimation window 2020Q4–2024Q1. Each horse race regression uses the empirical specification

$$\text{IHS(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \beta_2 \text{IT Capital Intensity}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. IT Capital Intensity is the IT share of the industry capital stock from the BEA detailed fixed-asset estimates, fixed at 2018 values and standardized over the estimation sample; it correlates with AI Assistance at 0.63 in the industry cross-section. Column (1) enters IT Capital Intensity alone. Column (2) enters AI Assistance alone. Column (3) enters both measures jointly. In Column (4), AI Assistance is residualized on IT Capital Intensity in the industry cross-section and the standardized residual replaces AI Assistance. All specifications include industry and quarter fixed effects and the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Table OA12. Horse Race Against IT Capital Intensity: Extended Window

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms	(3) IHS N. New Firms	(4) IHS N. New Firms
AI Assistance $\times$ Post		0.112*** (0.027)	0.090** (0.034)	
IT Capital Intensity $\times$ Post	0.092*** (0.023)		0.035 (0.031)	
AI Assistance (resid.) $\times$ Post				0.089** (0.035)
Observations	1,782	1,782	1,782	1,782
Industry 3D FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No	No	No
Controls	Yes	Yes	Yes	Yes
Adj. R. sq.	0.983	0.984	0.984	0.983

Notes: This table presents results racing AI Assistance against IT Capital Intensity. The design is a horse race difference-in-differences model at the industry-by-quarter level with continuous treatment. The sample includes 81 three-digit NAICS industries over the extended estimation window 2018Q4–2024Q1. Each horse race regression uses the empirical specification

$$\text{IHS}(\text{N. New Firms})_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \beta_2 \text{IT Capital Intensity}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. IT Capital Intensity is the IT share of the industry capital stock from the BEA detailed fixed-asset estimates, fixed at 2018 values and standardized over the estimation sample; it correlates with AI Assistance at 0.63 in the industry cross-section. Column (1) enters IT Capital Intensity alone. Column (2) enters AI Assistance alone. Column (3) enters both measures jointly. In Column (4), AI Assistance is residualized on IT Capital Intensity in the industry cross-section and the standardized residual replaces AI Assistance. All specifications include industry and quarter fixed effects and the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

The pre-treatment growth trend is each industry's OLS slope of IHS firm formation on time, fitted over 2015Q1–2018Q3—a window that closes before the estimation sample opens, so the measure and the estimation sample share no observations. A long-window variant fits the slope over 2006Q1–2018Q3. Tables OA13 and OA14 report the horse race and its variants on the baseline and extended estimation windows.

Table OA13. Horse Race Against the Pre-Treatment Growth Trend

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms	(3) IHS N. New Firms	(4) IHS N. New Firms	(5) IHS N. New Firms	(6) IHS N. New Firms
AI Assistance $\times$ Post	0.100*** (0.025)		0.099*** (0.026)	0.099*** (0.026)	0.073*** (0.022)	0.106*** (0.025)
Pre-Trend $\times$ Post		0.022 (0.025)	0.006 (0.023)			-0.036 (0.027)
Pre-Trend $\times$ Quarter	No	No	No	Yes	No	No
Industry Linear Trends	No	No	No	No	Yes	No
Long Trend Window	No	No	No	No	No	Yes
Observations	1,134	1,134	1,134	1,134	1,134	1,134
Industry 3D FE	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No	No	No	No	No
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R. sq.	0.982	0.981	0.982	0.982	0.988	0.982

Notes: This table presents results racing AI Assistance against each industry’s pre-treatment growth trend in firm formation. The design is a horse race difference-in-differences model at the industry-by-quarter level with continuous treatment. The sample includes 81 three-digit NAICS industries over the baseline estimation window 2020Q4–2024Q1. Each horse race regression uses the empirical specification

$$\text{IHS(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \beta_2 \text{Pre-Trend}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. The Pre-Trend is each industry’s OLS slope of IHS firm formation on time, fitted over 2015Q1–2018Q3—a window closing before the estimation sample opens—and standardized over the estimation sample. Column (1) presents the baseline results from the specification of Subsection 4.1. Column (2) enters the Pre-Trend alone. Column (3) enters both measures jointly. Column (4) replaces the single Post interaction with a full set of interactions between the Pre-Trend and the quarter indicators. Column (5) instead augments the baseline specification with industry-specific linear trends; this estimate is attenuated by construction, because the trend’s time index spans the estimation window, so each industry’s fitted slope partly tracks the treatment effect itself and then removes it; the attenuation is stronger the larger the share of post-treatment quarters in the estimation window. Column (6) fits the Pre-Trend over the long window 2006Q1–2018Q3. All specifications include industry and quarter fixed effects and the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Table OA14. Horse Race Against the Pre-Treatment Growth Trend: Extended Window

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms	(3) IHS N. New Firms	(4) IHS N. New Firms	(5) IHS N. New Firms	(6) IHS N. New Firms
AI Assistance $\times$ Post	0.112*** (0.027)		0.109*** (0.028)	0.109*** (0.029)	0.079*** (0.022)	0.118*** (0.027)
Pre-Trend $\times$ Post		0.036 (0.027)	0.017 (0.024)			-0.035 (0.027)
Pre-Trend $\times$ Quarter	No	No	No	Yes	No	No
Industry Linear Trends	No	No	No	No	Yes	No
Long Trend Window	No	No	No	No	No	Yes
Observations	1,782	1,782	1,782	1,782	1,782	1,782
Industry 3D FE	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No	No	No	No	No
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R. sq.	0.984	0.983	0.984	0.984	0.988	0.984

Notes: This table presents results racing AI Assistance against each industry’s pre-treatment growth trend in firm formation. The design is a horse race difference-in-differences model at the industry-by-quarter level with continuous treatment. The sample includes 81 three-digit NAICS industries over the extended estimation window 2018Q4–2024Q1. Each horse race regression uses the empirical specification

$$\text{IHS}(\text{N. New Firms})_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \beta_2 \text{Pre-Trend}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. The Pre-Trend is each industry’s OLS slope of IHS firm formation on time, fitted over 2015Q1–2018Q3—a window closing before the estimation sample opens—and standardized over the estimation sample. Column (1) presents the baseline results from the specification of Subsection 4.1. Column (2) enters the Pre-Trend alone. Column (3) enters both measures jointly. Column (4) replaces the single Post interaction with a full set of interactions between the Pre-Trend and the quarter indicators. Column (5) instead augments the baseline specification with industry-specific linear trends; this estimate is attenuated by construction, because the trend’s time index spans the estimation window, so each industry’s fitted slope partly tracks the treatment effect itself and then removes it; the attenuation is stronger the larger the share of post-treatment quarters in the estimation window. Column (6) fits the Pre-Trend over the long window 2006Q1–2018Q3. All specifications include industry and quarter fixed effects and the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

A.2. Alternative Specifications for Cost of Experimentation Mechanism. A natural concern with our baseline specification is that the results could be sensitive to the particular functional form used to classify experimentation-intensive industries. To address the robustness of our cost of experimentation mechanism results presented in Subsection 5.1, we redefined the indicator for experimentation-intensive industries. In the baseline specification, this indicator equals one if an industry’s experimentation intensity lies in the top quartile of the distribution.

In the first robustness test, instead of identify as Experimentation Intense industry based on a split, we replicate the results using the continuous (standardized) form of the Experimentation Intensity measure. Table OA15 displays the results for the specification using the interaction. The triple interaction coefficient, used to test the mechanism, still displays a positive sign, in line with our prediction. However, the results are attenuated both in terms of magnitude and significance.

Table OA15. Effects of the Cost of Experimentation (Continuous) on Firm Creation

VARIABLES	(1) IHS N. New Firms	(2) IHS N. New Firms
AI Assistance $\times$ Experimentation Intensity $\times$ Post	0.011 (0.036)	0.017 (0.035)
AI Assistance $\times$ Post	0.146* (0.073)	0.136* (0.071)
Experimentation Intensity $\times$ Post	-0.046 (0.062)	-0.041 (0.059)
Observations	1,078	1,078
Industry 3D FE	Yes	Yes
Quarter FE	No	No
Industry 2D $\times$ Quarter FE	Yes	Yes
Controls	No	Yes
Adj. R sq.	0.980	0.980


Notes: This table presents results testing whether the reduced cost-of-experimentation mechanism operates generating the triple interaction variable by using the continuous form of the Experimentation Intensity variable. The design is a triple differences model at the industry-by-quarter level with continuous treatment and a heterogeneity term. The sample includes 81 three-digit NAICS industries over 2020Q4–2024Q1. The empirical specification is

$$\begin{aligned} \text{IHS(N. New Firms)}_{it} = & \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Experimentation Intensity}_i \times \text{Post}_t \\ & + \beta_2 \text{AI Assistance}_i \times \text{Post}_t + \beta_3 \text{Experimentation Intensity}_i \times \text{Post}_t \\ & + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it} \end{aligned}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. Treatment heterogeneity comes from Experimentation Intensity, the standardized value of an industry’s Experimentation Intensity Score. The specifications include industry and sector-by-quarter fixed effects. Column (2) additionally includes the following controls: lagged VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

We also re-run the event study specification and plot the coefficients (Figure OA8), while using the continuous definition for experimentation intensity. The event study displays no pre-trends, in line with previous results and the post treatment coefficients are still positive, albeit not being significant.

Figure OA8. Event Study of Cost of Experimentation (Continuous) on Firm Creation Around the Release of ChatGPT



Notes: The figure presents the event study for the sample of 81 three-digit NAICS industries in our analysis, over the time period 2020Q4–2024Q1. We plot the estimates from the dynamic triple differences model for the interaction of the quarter-year dummies with the AI Assistance_{*i*} × Experimentation Intensity_{*i*} from the specification

$$\begin{aligned}
 \text{IHS(N. New Firms)}_{it} = & \beta_0 + \sum_{t \neq 2022q3} \beta_t \text{AI Assistance}_i \times \text{Experimentation Intensity}_i \times 1\{D = t\}_t \\
 & + \sum_{t \neq 2022q3} \gamma_t \text{AI Assistance}_i \times 1\{D = t\}_t \\
 & + \sum_{t \neq 2022q3} \delta_t \text{Experimentation Intensity}_i \times 1\{D = t\}_t + \text{Controls}_{it} + \mu_i + \mu_{1t} + \epsilon_{it}
 \end{aligned}$$

The dependent variable is the Inverse Hyperbolic Sine of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment heterogeneity comes from the Experimentation Intensity variable, which is the standardized industry-level Experimentation Intensity Score. The specifications include both industry and sector-by-quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level. The specification in Panel (a) does not include controls. The specification in Panel (b) includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage.

In the second robustness test, we reclassify industries as experimentation-intensive when their Experimentation Intensity Score is above the sample median. Table OA16 displays the results for the specification using the new binary indicator. The triple interaction coefficient, used to test the mechanism, still displays a positive sign, in line with our prediction. However, the results are attenuated both in terms of magnitude and significance.

Table OA16. Effects of the Cost of Experimentation Above Median on Firm Creation


VARIABLES	(1)	(2)
	IHS N. New Firms	IHS N. New Firms
AI Assistance $\times$ Experimentation Above Median $\times$ Post	0.103 (0.099)	0.111 (0.095)
AI Assistance $\times$ Post	0.070 (0.066)	0.062 (0.063)
Experimentation Above Median $\times$ Post	-0.053 (0.074)	-0.057 (0.073)
Observations	1,078	1,078
Industry 3D FE	Yes	Yes
Quarter FE	No	No
Industry 2D $\times$ Quarter FE	Yes	Yes
Controls	No	Yes
Adj. R sq.	0.980	0.980

Notes: This table presents results testing whether the reduced cost-of-experimentation mechanism operates using an alternative definition of experimentation intense industries. The design is a triple differences model at the industry-by-quarter level with continuous treatment and a heterogeneity term. The sample includes 81 three-digit NAICS industries over 2020Q4–2024Q1. The empirical specification is

$$\begin{aligned} \text{IHS(N. New Firms)}_{it} = & \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Experimentation Above Median}_i \times \text{Post}_t \\ & + \beta_2 \text{AI Assistance}_i \times \text{Post}_t + \beta_3 \text{Experimentation Above Median}_i \times \text{Post}_t \\ & + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it} \end{aligned}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. Treatment heterogeneity comes from the indicator Experimentation Above Median, which equals 1 if the industry's Experimentation Intensity Score is above the sample median and 0 otherwise. The specifications include industry and sector-by-quarter fixed effects. Column (2) additionally includes the following controls: lagged VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

We also re-run the event study specification and plot the coefficients (Figure OA9), while using the alternative definition for experimentation-intense industries. The event study displays no pre-trends, in line with previous results and the post treatment coefficients are still positive, albeit not being significant.

Figure OA9. Event Study of Cost of Experimentation Above Median on Firm Creation Around the Release of ChatGPT

Notes: The figure presents the event study for the sample of 81 three-digit NAICS industries in our analysis, over the time period 2020Q4–2024Q1. We plot the estimates from the dynamic triple differences model for the interaction of the quarter-year dummies with the AI Assistance_{*i*} × Experimentation Above Median_{*i*} from the specification

$$\begin{aligned}
 \text{IHS(N. New Firms)}_{it} = & \beta_0 + \sum_{t \neq 2022q3} \beta_t \text{AI Assistance}_i \times \text{Experimentation Above Median}_i \times 1\{D = t\}_t \\
 & + \sum_{t \neq 2022q3} \gamma_t \text{AI Assistance}_i \times 1\{D = t\}_t \\
 & + \sum_{t \neq 2022q3} \delta_t \text{Experimentation Above Median}_i \times 1\{D = t\}_t + \text{Controls}_{it} + \mu_i + \mu_{1t} + \epsilon_{it}
 \end{aligned}$$

The dependent variable is the Inverse Hyperbolic Sine of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment heterogeneity comes from the indicator variable Experimentation Above Median, which is equal to 1 if the industry reports a value of Experimentation Intensity above the sample median and 0 otherwise. The specifications include both industry and sector-by-quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level. The specification in Panel (a) does not include controls. The specification in Panel (b) includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage.

In the third robustness test, we reclassify industries as experimentation-intensive when their Experimentation Intensity Score is above the sample mean. Table OA17 displays the results for the specification using the new binary indicator. Similar to the previous test, the triple interaction coefficient displays a positive sign, albeit being smaller in size and not significant when compared to the baseline result.

Table OA17. Effects of the Cost of Experimentation Above Mean on Firm Creation


VARIABLES	(1)	(2)
	IHS N. New Firms	IHS N. New Firms
AI Assistance $\times$ Experimentation Above Mean $\times$ Post	0.085 (0.091)	0.095 (0.087)
AI Assistance $\times$ Post	0.089 (0.070)	0.077 (0.068)
Experimentation Above Mean $\times$ Post	-0.076 (0.072)	-0.075 (0.069)
Observations	1,078	1,078
Industry 3D FE	Yes	Yes
Quarter FE	No	No
Industry 2D $\times$ Quarter FE	Yes	Yes
Controls	No	Yes
Adj. R sq.	0.980	0.980

Notes: This table presents results testing whether the reduced cost-of-experimentation mechanism operates using an alternative definition of experimentation intense industries. The design is a difference-in-differences model at the industry-by-quarter level with continuous treatment and a heterogeneity term. The sample includes 81 three-digit NAICS industries over 2020Q4–2024Q1. The empirical specification is

$$\begin{aligned} \text{IHS(N. New Firms)}_{it} = & \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Experimentation Above Mean}_i \times \text{Post}_t \\ & + \beta_2 \text{AI Assistance}_i \times \text{Post}_t + \beta_3 \text{Experimentation Above Mean}_i \times \text{Post}_t \\ & + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it} \end{aligned}$$

The dependent variable is the inverse hyperbolic sine (IHS) of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. Treatment heterogeneity comes from the indicator Experimentation Intense, which equals 1 if the industry's Experimentation Intensity Score is above the sample mean and 0 otherwise. The specifications include industry and sector-by-quarter fixed effects. Column (2) additionally includes the following controls: lagged VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

We also re-run the event study specification and plot the coefficients (Figure OA10), while using the alternative definition for experimentation-intense industries. The event study displays no pre-trends and positive (non-significant) post treatment coefficients, in line with previous test and the baseline results.

Figure OA10. Event Study of Cost of Experimentation Above Mean on Firm Creation Around the Release of ChatGPT

Notes: The figure presents the event study for the sample of 81 three-digit NAICS industries in our analysis, over the time period 2020Q4–2024Q1. We plot the estimates from the dynamic triple differences model for the interaction of the quarter-year dummies with the AI Assistance_{*i*} × Experimentation Above Mean_{*i*} from the specification

$$\begin{aligned}
 \text{IHS(N. New Firms)}_{it} = & \beta_0 + \sum_{t \neq 2022q3} \beta_t \text{AI Assistance}_i \times \text{Experimentation Above Mean}_i \times 1\{D = t\}_t \\
 & + \sum_{t \neq 2022q3} \gamma_t \text{AI Assistance}_i \times 1\{D = t\}_t \\
 & + \sum_{t \neq 2022q3} \delta_t \text{Experimentation Above Mean}_i \times 1\{D = t\}_t + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it}
 \end{aligned}$$

The dependent variable is the Inverse Hyperbolic Sine of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment heterogeneity comes from the indicator variable Experimentation Above Mean, which is equal to 1 if the industry reports a value of Experimentation Intensity above the sample mean and 0 otherwise. The specifications include both industry and sector-by-quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level. The specification in Panel (a) does not include controls. The specification in Panel (b) includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage.

Overall, these robustness tests qualitatively match our main findings. The fact that the coefficients are somewhat attenuated in size and significance level should not be interpreted as evidence that the mechanism is sensitive to functional form choice. Instead, it highlights that experimentation costs act as a barrier to entry primarily when experimentation is central to the operations of an industry.

A.3. Alternative Interpretation of the Capital Allocation Mechanism. An alternative explanation for the observed increase in firm creation is that the *capital response channel* may operate in the opposite direction of the standard capital reallocation story. Rather than

attracting more financing, AI-compatible industries may require less capital per firm once AI raises productivity. In this case, average deal sizes would fall, and the rise in firm creation could be mediated by reduced capital needs rather than by increased capital inflows. This “efficiency channel” would imply that lower capital requirements per startup are what enable more entrepreneurs to enter.

To test this mechanism, we rerun the mediation analysis presented in Table 9 of the main text using the inverse hyperbolic sine (IHS) of the average investment size as the mediator, rather than total capital. The results are presented in Table OA18. For the sake of exposition, we present again the effect of the treatment on firm creation in Column (1). In column (2), we regress the IHS of average investment size on the treatment variable. Here the coefficient is negative, about -0.36 , and statistically significant at the 10 percent level. This provides some, albeit weak, evidence that deals are smaller in AI-compatible industries after ChatGPT. Finally, we include the IHS of average investment size directly in the firm creation regression as a potential mediator. In this specification, the coefficient on average investment size is statistically insignificant, and the coefficient on AI Assistance $\times$ Post remains unchanged at approximately 0.10.

These results suggest that although there is a weak indication that average deal sizes decline in AI-compatible industries post-ChatGPT, this decline does not mediate the observed increase in firm creation. In other words, the treatment effect on entry is unaffected by controlling for changes in average investment size.

Table OA18. Mediation Analysis: Average Investment Size, AI Assistance and Firm Creation

VARIABLES	(1) IHS N. New Firms	(2) IHS Avg. Investment	(3) IHS N. New Firms
AI Assistance $\times$ Post	0.102** (0.046)	-0.360* (0.203)	0.102** (0.046)
IHS Avg. Investment			-0.000 (0.004)
Observations	1,078	1,078	1,078
Industry 3D FE	Yes	Yes	Yes
Quarter FE	No	No	No
Industry 2D $\times$ Quarter FE	Yes	Yes	Yes
Controls	No	No	No
Adj. R sq.	0.980	0.451	0.980

Notes: This table reports a mediation-style decomposition to study an alternative version of the capital reallocation as a potential mechanism linking AI's arrival and firm creation. The analysis is at the industry-by-quarter level for 81 three-digit NAICS industries over 2020Q4–2024Q1. The mediator is average allocated capital, measured as the inverse hyperbolic sine (IHS) of the average VCPE investment size from Pre_{qin} (in \$000).

Column (1) reports the results for the estimation of the total impact that the treatment has on firm formation. The specification is

$$\text{IHS(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \mu_i + \mu_{It} + \epsilon_{it}$$

Column (2) assesses the linkage between the treatment and the average size of VCPE investments, according to the specification

$$\text{IHS(Avg. Investment)}_{it} = \gamma_0 + \gamma_1 \text{AI Assistance}_i \times \text{Post}_t + \mu_i + \mu_{It} + \epsilon_{it}$$


Column (3) reports the results for the estimation of the mediated effect of the treatment once we control for the average size of VCPE investments. The specification is

$$\text{IHS(N. New Firms)}_{it} = \delta_0 + \delta_1 \text{AI Assistance}_i \times \text{Post}_t + \delta_2 \text{IHS(Avg. Investment)}_{it} + \mu_i + \mu_{It} + \epsilon_{it}$$

All specifications include both industry and sector-by-quarter fixed effects. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which is equal to 1 for all quarters after 2022Q3, and 0 otherwise. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

To ensure comparability across differently treated industries, we run the event study version of the specification in Table OA18, Column (2). Table OA11 shows that neither anticipatory effects nor differences in pre-trends are much of a concern for average deal size, reinforcing the findings in Table OA18. Moreover, the post treatment effects appear to be weakly negative.

Figure OA11. Event Study of AI Assistance on Average Allocated Capital Around the Release of ChatGPT



Notes: The figure presents the event study for the sample of 81 three-digit NAICS industries in our analysis, over the time period 2020Q4–2024Q1. We plot the estimates from the dynamic difference-in-differences model for the interaction of the quarter-year dummies with the AI Assistance_{*i*} from the specification

$$\text{IHS}(\text{Avg. Investment})_{it} = \beta_0 + \sum_{t \neq 2022q3} \beta_t \text{AI Assistance}_i \times 1\{D = t\}_t + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it}$$

The dependent variable is the Inverse Hyperbolic Sine of the average size of VCPE investments, in thousands of dollars. Treatment intensity is proxied by AI Assistance in its standardized form. The specification includes both industry and sector-by-quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level.

Overall, the finding indicates that the efficiency channel—in the sense of reduced capital needs per firm leading to more entry—does not have much traction in the aggregate data. It remains possible that the channel operates for a subset of ventures, particularly those whose founders can directly exploit AI to substitute for costly inputs. However, when examining firm creation at the industry level, which aggregates across heterogeneous entrants, the effect does not appear clearly.

APPENDIX B. ROBUSTNESS: CHEN AND ROTH (2024)’S CRITIQUE

A potential concern with our baseline specification is the use of the inverse hyperbolic sine (IHS) transformation of the dependent variable, the number of newly created firms. Chen and Roth (2024) show that estimators based on log-like transformations such as IHS or $\log(1 + y)$ can suffer from interpretational and robustness issues, particularly when the outcome variable can take the value zero. Although our outcome is strictly positive in the estimation sample—making the natural logarithm well defined—it is important to demonstrate that our findings do not hinge on the choice of functional form.

Since the dependent variables used in the main analysis (Table 4 and Figure 5) and in the cost of experimentation mechanism (Table 8 and Figure 8) are strictly positive, we can test whether our results are solely an artifact of the functional specification (IHS) chosen by replicating the results transforming the dependent variable using the natural logarithm.

In the capital reallocation test (Table 9 and Figure 10), we need to employ more nuanced techniques to test whether the results are driven the the IHS. We therefore conduct a series of robustness tests designed to directly address these concerns. Each test isolates a specific aspect of the Chen and Roth (2024) critique and shows that our substantive conclusions remain unchanged.

B.1. Replicating Effects of AI on Firm Creation Results. Table OA19 presents our baseline difference-in-differences estimates of the effect of AI Assistance on firm creation, replacing the inverse hyperbolic sine transformation with the natural logarithm of the number of new firms. We do this for two reasons. Since the number of newly created firms is strictly positive in our estimation sample, the logarithm is well defined and provides a clean percentage interpretation of the estimated coefficients. Therefore, the new specification directly address the critique raised by Chen and Roth (2024), who show that log-like transformations such as IHS or $\log(1 + y)$ can yield estimators with problematic interpretation when zeros are present. By relying on $\text{Log}(\text{New Firms})$ rather than IHS, we demonstrate that our results are not driven by functional-form artifacts.

Across both specifications, with and without controls, the coefficient on *AI Assistance* $\times$ *Post* is stable at approximately 0.10 and statistically significant at the 1% level. The results are identical to the baseline estimates in Table 4.

Table OA19. Logarithmic Functional Form: Effect of AI Assistance on Firm Creation After the Release of ChatGPT


VARIABLES	(1) Log N. New Firms	(2) Log N. New Firms
AI Assistance $\times$ Post	0.105*** (0.025)	0.100*** (0.025)
Observations	1,134	1,134
Industry 3D FE	Yes	Yes
Quarter FE	Yes	Yes
Industry 2D $\times$ Quarter FE	No	No
Controls	No	Yes
Adj. R. sq.	0.982	0.982

Notes: This table addresses the concerns raised in [Chen and Roth \(2024\)](#) regarding the use of the IHS transformation on the outcome variable. We re-run the estimation results in Table 4, now expressing the dependent variable in its logarithmic form. The design is a difference-in-differences model at the industry-by-quarter level with continuous treatment. The sample includes 81 three-digit NAICS industries over 2020Q4–2024Q1. The empirical specification is

$$\text{Log(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the natural logarithm of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. All specifications include industry and quarter fixed effects. Column (2) additionally includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Figure [OA12](#) presents the dynamic difference-in-differences estimates of the effect of AI Assistance on firm creation, using Log(New Firms) as the dependent variable. The results are identical to the baseline event studies in Figure 5.

Figure OA12. Logarithmic Functional Form: Event Study of AI Assistance on Firm Creation around the Release of ChatGPT

Notes: This panel of figures addresses the concerns raised in [Chen and Roth \(2024\)](#) regarding the use of the IHS transformation on the outcome variable. We re-run the estimation results in Figure 5, now expressing the dependent variable in its logarithmic form. The figure presents the event study for the sample of 81 three-digit NAICS industries in our analysis, over the time period 2020Q4–2024Q1. We plot the estimates from the dynamic difference-in-differences model for the interaction of the quarter-year dummies with the AI Assistance_{*i*} from the specification

$$\text{Log(N. New Firms)}_{it} = \beta_0 + \sum_{t \neq 2022q3} \beta_t \text{AI Assistance}_i \times 1\{D = t\}_t + \text{Controls}_{it} + \mu_i + \mu_t + \epsilon_{it}$$

The dependent variable is the natural logarithm of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. The specifications include both industry and quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level. The specification in Panel (a) does not include controls. The specification in Panel (b) includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage.

Overall, our findings support the idea that the main results of our analysis presented in Section 4 are not an artifact of the IHS specification.

B.2. Replicating of Cost of Experimentation Mechanism Results. The results on the reduction of the cost of experimentation mechanism (Subsection 5.1) are also subject to the critique presented in [Chen and Roth \(2024\)](#) regarding the use of the IHS transformation on the outcome variable. To address these concern we replicate the analysis using the natural logarithm of the dependent variable. We are able to do so because the dependent variable is strictly positive.

The results using the natural logarithm for the dependent variable are completely unvaried both in term of size and significance. The coefficient on the triple interaction variable for the specification not including control remain significant at the 5% level, with magnitude 0.26 (Table OA20), exactly as we find in Table 8, in the main text.

Table OA20. Logarithmic Functional Form: Effects of the Cost of Experimentation on Firm Creation


VARIABLES	(1) Log N. New Firms	(2) Log N. New Firms
AI Assistance $\times$ Experimentation Intense $\times$ Post	0.261** (0.111)	0.264** (0.112)
AI Assistance $\times$ Post	0.037 (0.054)	0.032 (0.054)
Experimentation Intense $\times$ Post	-0.190 (0.119)	-0.193 (0.118)
Observations	1,078	1,078
Industry 3D FE	Yes	Yes
Quarter FE	No	No
Industry 2D $\times$ Quarter FE	Yes	Yes
Controls	No	Yes
Adj. R sq.	0.981	0.981

Notes: This table addresses the concerns raised in [Chen and Roth \(2024\)](#) regarding the use of the IHS transformation on the outcome variable. We re-run the estimation results in Table 8, now expressing the dependent variable in its logarithmic form. The design is a difference-in-differences model at the industry-by-quarter level with continuous treatment and a heterogeneity term. The sample includes 81 three-digit NAICS industries over 2020Q4–2024Q1. The empirical specification is

$$\begin{aligned} \text{Log(N. New Firms)}_{it} = & \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Experimentation Intense}_i \times \text{Post}_t \\ & + \beta_2 \text{AI Assistance}_i \times \text{Post}_t + \beta_3 \text{Experimentation Intense}_i \times \text{Post}_t \\ & + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it} \end{aligned}$$

The dependent variable is the natural logarithm of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. Treatment heterogeneity comes from the indicator Experimentation Intense, which equals 1 if the industry’s Experimentation Intensity Score is in the top quartile and 0 otherwise. The specifications include industry and sector-by-quarter fixed effects. Column (2) additionally includes the following controls: lagged VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

The event study results are also unchanged when we substitute the IHS form with the log one. Figure OA13 presents the results, showing that the coefficients in it display the same patterns observed in Figure 8.

Figure OA13. Logarithmic Functional Form: Event Study of Cost of Experimentation on Firm Creation around the Release of ChatGPT

Notes: This panel of figures addresses the concerns raised in [Chen and Roth \(2024\)](#) regarding the use of the IHS transformation on the outcome variable. We re-run the estimation results in Figure 8, now expressing the dependent variable in its logarithmic form. The figure presents the event study for the sample of 81 three-digit NAICS industries in our analysis, over the time period 2020Q4–2024Q1. We plot the estimates from the dynamic difference-in-differences model for the interaction of the quarter-year dummies with the $AI Assistance_i \times Experimentation Intensity_i$ from the specification

$$\begin{aligned}
 \text{Log(N. New Firms)}_{it} = & \beta_0 + \sum_{t \neq 2022q3} \beta_t AI Assistance_i \times Experimentation Intense_i \times 1\{D = t\}_t \\
 & + \sum_{t \neq 2022q3} \gamma_t AI Assistance_i \times 1\{D = t\}_t + \sum_{t \neq 2022q3} \delta_t Experimentation Intense_i \times 1\{D = t\}_t \\
 & + \text{Controls}_{it} + \mu_i + \mu_{It} + \epsilon_{it}
 \end{aligned}$$

The dependent variable is the natural logarithm of the number of new firms created. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment heterogeneity comes from the indicator variable Experimentation Intense, which is equal to 1 if the industry reports a value of Experimentation Intensity within the top quartile and 0 otherwise. The specifications include both industry and sector-by-quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level. The specification in Panel (a) does not include controls. The specification in Panel (b) includes the following controls: lag of VCPE Investments (IHS), median Return on Assets, median R&D Intensity, median CapEx Intensity, median Liquidity Ratio, and median Market Leverage.

Overall, our findings support the idea that the main results of our analysis presented in Subsection 5.1 are not an artifact of the IHS specification.

B.3. Testing Capital Reallocation Mechanism Results. A distinctive challenge in the capital reallocation tests is that the mediator, total VCPE investments, occasionally takes the value zero. This creates difficulties for conventional logarithmic specifications: $\log(0)$ is undefined, and ad hoc fixes such as $\log(1 + y)$ or IHS can blur the distinction between the extensive margin (whether there is any investment at all) and the intensive margin (the size

of investment when positive). [Chen and Roth \(2024\)](#) show that such transformations can lead to estimators with problematic interpretations, particularly when zeros are common. For this reason, the tests of capital reallocation require more nuanced approaches than those for firm creation or experimentation, where the dependent variables are strictly positive.

We therefore conduct three complementary robustness exercises, each designed to address a specific aspect of the [Chen and Roth \(2024\)](#) critique. First, we check whether results hinge on using IHS versus the natural log, ensuring clean percentage interpretations when outcomes are strictly positive. Second, we employ a PPML estimation, which handles the presence of zeros in investment flows and avoids arbitrary log-like transformations by modeling the conditional mean directly. Lastly, we use a two-part (hurdle) model, which separates extensive (any investment) and intensive (log size when positive) margins, addressing the concern that log/IHS transforms conflate them.

B.3.1. Log Specification Estimation. First, we reproduce the mediation table replacing IHS transformations with natural logarithms (Table [OA21](#)). This specification is natural for strictly positive variables, such as the number of new firms in our sample, and allows a clean percentage interpretation of coefficients. If our findings were an artifact of using IHS, results would change materially under the log specification, due to dropping the zero-valued industry-quarter cells. Instead, we find that the coefficient on *AI Assistance* $\times$ *Post* remains around 0.10 and statistically significant, very similar to the IHS baseline (Table [9](#), Column (1)). In the investment regressions (Table [OA21](#), Column (2)), the effect of AI compatibility remains negative and statistically significant, with size comparable to what we find in Table [9](#) in the main text. Moreover, when looking at Column (3) in Table [OA21](#), not only we find that the coefficient on *AI Assistance* $\times$ *Post* remains around 0.10 and statistically significant, but the mediator continues to be uninformative. These results are almost identical to what we found in the main text (Table [9](#)). The above confirms that our main results are robust to abandoning IHS and using the theoretically preferred logarithmic specification.

Table OA21. Logarithmic Functional Form: Capital Reallocation Mechanism

VARIABLES	(1) Log N. New Firms	(2) Log Total Investments	(3) Log N. New Firms
AI Assistance $\times$ Post	0.102** (0.046)	-0.493 (0.350)	0.100* (0.053)
Log Total Investments			0.004 (0.006)
Observations	1,078	811	811
Industry 3D FE	Yes	Yes	Yes
Quarter FE	No	No	No
Industry 2D $\times$ Quarter FE	Yes	Yes	Yes
Controls	No	No	No
Adj. R sq.	0.980	0.559	0.982

Notes: This table reports the robustness tests results for the mediation-style decomposition to study capital reallocation as a potential mechanism linking AI's arrival and firm creation. To address the critique presented [Chen and Roth \(2024\)](#), we replicated the analysis using the natural logarithm instead of the inverse hyperbolic sine transformation. The analysis is at the industry-by-quarter level for 81 three-digit NAICS industries over 2020Q4–2024Q1. The mediator is total capital allocated, measured as the natural logarithm of total VCPE investments from Preqin (in \$000).

Column (1) reports the results for the estimation of the total impact that the treatment has on firm formation. The specification is

$$\text{Log(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \mu_i + \mu_{It} + \epsilon_{it}$$

Column (2) assesses the linkage between the treatment and the total capital allocation, according to the specification

$$\text{Log(Investments)}_{it} = \gamma_0 + \gamma_1 \text{AI Assistance}_i \times \text{Post}_t + \mu_i + \mu_{It} + \epsilon_{it}$$

Column (3) reports the results for the estimation of the mediated effect of the treatment once we control for total allocated capital. The specification is

$$\text{Log(N. New Firms)}_{it} = \delta_0 + \delta_1 \text{AI Assistance}_i \times \text{Post}_t + \delta_2 \text{Log(Investments)}_{it} + \mu_i + \mu_{It} + \epsilon_{it}$$

All specifications include both industry and sector-by-quarter fixed effects. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which is equal to 1 for all quarters after 2022Q3, and 0 otherwise. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

B.3.2. Estimating Investments with PPML. To further address the short-comings of using dependent variables in their IHS form, we follow the recommendation in [Chen and Roth \(2024\)](#) to use the Poisson pseudo-maximum likelihood (PPML) estimator, which avoids outcome transformations altogether and is well defined in the presence of zeros. While our firm creation outcome never equals zero, investments occasionally do, making PPML a natural robustness check for the mediator equation. We therefore re-estimate the investment regression using PPML, while keeping Columns (1) and (3) in IHS form for comparability. The results are strikingly similar: the estimated coefficient on *AI Assistance $\times$ Post* on total

investments remains negative (around -0.33) and statistically significant, in line with the results in Table 9. This exercise demonstrates that our conclusion about the irrelevance of capital reallocation does not hinge on log-like transformations of the investment variable.

Table OA22. Testing the Capital Reallocation Channel with PPML

VARIABLES	(1) IHS N. New Firms	(2) Total Investments	(3) IHS N. New Firms
AI Assistance $\times$ Post	0.102** (0.046)	-0.329** (0.165)	0.102** (0.046)
IHS Total Investments			0.000 (0.004)
Observations	1,078	1,064	1,078
Industry 3D FE	Yes	Yes	Yes
Quarter FE	No	No	No
Industry 2D $\times$ Quarter FE	Yes	Yes	Yes
Controls	No	No	No
Adj. R sq.	0.980	0.774	0.980

Notes: This table reports the robustness tests results for the mediation-style decomposition to study capital reallocation as a potential mechanism linking AI's arrival and firm creation. To address the critique presented [Chen and Roth \(2024\)](#), we replicate the linkage between the treatment and total capital allocation using Poisson pseudo-maximum likelihood. The rest of the analysis is unchanged from the baseline estimation. The analysis is at the industry-by-quarter level for 81 three-digit NAICS industries over 2020Q4–2024Q1. The mediator is total capital allocated, measured as the natural logarithm of total VCPE investments from Preqin (in \$000).

Column (1) reports the results for the estimation of the total impact that the treatment has on firm formation. The specification is

$$\text{Log(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \mu_i + \mu_{It} + \epsilon_{it}$$

Column (2) assesses the linkage between the treatment and total capital allocation. The dependent variable now is total investments, in levels, estimated using Poisson pseudo-maximum likelihood according to the specification

$$E[\text{Investments}_{it} | \text{AI Assistance}_i, \text{Post}_t, \mu_i, \mu_{It}] = \exp\left(\gamma_0 + \gamma_1 \text{AI Assistance}_i \times \text{Post}_t + \mu_i + \mu_{It}\right)$$

Column (3) reports the results for the estimation of the mediated effect of the treatment once we control for total allocated capital. The specification is


$$\text{IHS(N. New Firms)}_{it} = \delta_0 + \delta_1 \text{AI Assistance}_i \times \text{Post}_t + \delta_2 \text{IHS(Investments)}_{it} + \mu_i + \mu_{It} + \epsilon_{it}$$

All specifications include both industry and sector-by-quarter fixed effects. Treatment intensity is proxied by AI Assistance in its standardized form. Treatment timing is proxied by the binary indicator Post, which is equal to 1 for all quarters after 2022Q3, and 0 otherwise. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

We also replicate the event study for total investments using PPML, to contrast it with the results in Figure 10. Figure OA14 presents the event study estimates for the model in Column (2) Table OA22. The results show no evidence of anticipatory investment: pre-treatment coefficients are generally statistically indistinguishable from zero. Following the release of ChatGPT, the coefficients become negative and significant, indicating that industries more compatible with AI experienced a relative decline in VCPE investment flows. This dynamic pattern, not only corroborates the static PPML results in Table OA22, but is also

largely in line with the results in the main text showing that capital does not flow towards more AI-compatible industries.

Figure OA14. Dynamic Effects of AI Assistance on Investments: PPML Event Study



Notes: The figure presents the event study for regression of total VCPE investments estimated using PPML. The sample is comprised of 81 three-digit NAICS industries, over the time period 2020Q4–2024Q1. We plot the estimates from the dynamic difference-in-differences model for the interaction of the quarter-year dummies with the AI Assistance_i from the specification

$$E \left[\text{Investments}_{it} \mid \text{AI Assistance}_i, \sum_{t \neq 2022q3} 1\{D = t\}_t, \mu_i, \mu_{It} \right] = \exp \left(\beta_0 + \sum_{t \neq 2022q3} \beta_t \text{AI Assistance}_i \times 1\{D = t\}_t + \mu_i + \mu_{It} \right)$$

The dependent variable is the the dollar amount of VCPE investments, in thousands. Treatment intensity is proxied by AI Assistance in its standardized form. The specification includes both industry and sector-by-quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level.

B.3.3. Two-Part Hurdle Specification: Extensive Versus Intensive Margin. A key critique in [Chen and Roth \(2024\)](#) is that log or IHS transformations implicitly conflate the extensive margin (whether investment is zero or positive) with the intensive margin (the size of investment conditional on being positive). To address this issue, we explicitly decompose investments into extensive and intensive margins. To disentangle these effects, we estimate

a two-part (hurdle) specification that includes (a) a dummy for any positive investment and (b) the log of investment conditional on being positive. This allows us to test whether either margin explains firm creation. The results in Table OA23 show that neither component is statistically significant: the extensive margin dummy is small and negative, and the intensive margin coefficient is close to zero. Importantly, the coefficient on *AI Assistance* $\times$ *Post* remains stable and significant at around 0.10, as in the baseline specification reported in the main text (Table 9, Column (3)). This confirms that our treatment effect is not mediated by either margin of investment activity.

Table OA23. Decomposing the Capital Reallocation Channel: Extensive and Intensive Margins of Investment

VARIABLES	(1) Log N. New Firms
AI Assistance $\times$ Post	0.106** (0.046)
Extensive Margin	-0.052 (0.038)
Intensive Margin	0.006 (0.005)
Observations	1,078
Industry 3D FE	Yes
Quarter FE	No
Industry 2D $\times$ Quarter FE	Yes
Controls	No
Adj. R. sq.	0.980

Notes: This table tests whether the effect of AI Assistance on new firm creation is mediated through the extensive or intensive margins of investment, following the indication of [Chen and Roth \(2024\)](#). The analysis is at the industry-by-quarter level for 81 three-digit NAICS industries over 2020Q4–2024Q1. The specification is

$$\text{Log(N. New Firms)}_{it} = \beta_0 + \beta_1 \text{AI Assistance}_i \times \text{Post}_t + \beta_2 \text{Extensive Margin}_{it} + \beta_3 \text{Intensive Margin}_{it} + \mu_i + \mu_{It} + \epsilon_{it}$$

The dependent variable is the natural logarithm of the number of new firms created. Treatment intensity is proxied by AI Assistance in standardized form, and treatment timing is proxied by Post, which equals 1 for all quarters after 2022Q3 and 0 otherwise. Investments are decomposed into two components: (i) an Extensive Margin indicator equal to 1 if total investments in the industry-quarter cell are positive and 0 otherwise, and (ii) an Intensive Margin measure equal to the log of total investment, conditional on being positive. The specification includes industry and sector-by-quarter fixed effects. Standard errors are clustered at the three-digit NAICS industry level and reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

APPENDIX C. ADDITIONAL INFORMATION ON REVELIO LABS

C.1. Sourcing of Revelio Labs Data. Revelio Labs is a workforce intelligence provider whose primary data source is public information drawn from LinkedIn. By systematically scraping and parsing professional profiles and job postings, the company builds a structured dataset that captures employment histories, job descriptions, and firm characteristics. Although LinkedIn is the main input, Revelio further integrates auxiliary data streams such as employee reviews and official layoff notices, but these are secondary relative to the core profile-based data.

A central feature of Revelio Labs’ methodology is the standardization and homogenization of raw information. Professional profiles and postings often contain highly heterogeneous job titles, firm names, and skill descriptions. To address this, Revelio applies proprietary taxonomies and entity-resolution techniques that map companies to unique identifiers, collapse synonymous or overlapping job titles into coherent occupational categories, and cluster skills into standardized hierarchies. This harmonization makes the data usable for systematic analysis at scale.

The structured database that results can be analyzed at several levels of aggregation. At the company level, Revelio provides measures such as total headcount, inflows and outflows of employees, hiring rates, and job posting dynamics. At the position level, the data track individual job postings and role descriptions, including salary information and required skills, deduplicated across multiple sources. At the individual level, the dataset captures workers’ career histories, including transitions across employers, tenure, and educational backgrounds.

Because the underlying source is LinkedIn, coverage is not uniform across occupations, industries, and geographies. White-collar and professional workers are more likely to maintain profiles than blue-collar workers, for example. To correct for this selection bias, Revelio Labs constructs and distributes sampling weights at the occupation level. These weights are designed to re-weight observed profiles so that the aggregate distributions more closely match official labor force statistics. For example, if a particular occupation is less likely to appear in online profiles, each observed instance may be weighted to represent multiple actual workers in the population. This sampling adjustment ensures the dataset reflects a more accurate workforce composition. This adjustment enhances the representativeness of the dataset and supports credible empirical analysis.


C.2. Representativeness of Revelio Labs Data. As discussed in the main paper (Section 2), the *Revelio Labs* (RL) dataset provides an alternative source of information on new firm creation, which we compare here to the official Business Dynamics Statistics (BDS) from the

U.S. Census²⁵. Due to the fact that BDS is available exclusively at the yearly level and it is reported with a lag, our validation exercise focuses on yearly firm formation data from 2006-2022²⁶. Overall, the evidence in this appendix indicates that RL is broadly representative of the general economy in terms of new firm formation. In particular, RL data closely tracks the patterns and distribution of firm births observed in BDS. We emphasize that our representativeness claims pertain to firm creation statistics *specifically* — RL may not be as representative for other economic measures, but it proves to be an acceptably reliable proxy for studying new firm formation. Finally, we note that while the BDS dataset occasionally suppresses firm creation counts in industry-year cells where the true number is very small (to protect confidentiality), this occurs infrequently in our validation exercise, which uses a macro-level aggregation (three-digit NAICS by year) where such censoring is rare and inconsequential for our comparisons.

C.2.1. Industry Distribution of New Firms in RL vs. BDS. We first examine whether RL captures a similar *industry composition* of new firms as in the official data. We plot the Spearman and Pearson coefficients over time in Figure OA15. The plots display the correlation between the industry share of new firms in RL and the corresponding industry share in BDS. Figure OA15, Panel (a) uses Spearman’s rank correlation (measuring similarity in the rank-order of industry importance), while Figure OA15, Panel (b) uses Pearson’s correlation (measuring linear alignment of the share percentages). In both cases, the yearly correlations are very high and have improved over time. In the mid-2000s (when digital coverage was more limited), the correlations were moderately high; by the 2010s and beyond, they approach or exceed 0.8. This indicates that in each year RL data produces a very similar distribution of new firms across industries as the economy-wide distribution captured by BDS. In other words, the types of businesses being started (by industry) in RL match those in the broader economy. This alignment gives confidence that RL is not missing whole categories of firms: new firms in manufacturing, services, tech, etc., all appear in roughly the correct proportions.

²⁵BDS is the benchmark source for annual counts of new employer firms in the United States. RL data, derived from online workforce records, offers more granular and timely information; this appendix evaluates how well RL mirrors BDS in the context of firm creation.


²⁶While most of the analysis in the main paper is done for the period 2020Q4-2024Q1, we are forced to change the time aggregation and window in order to draw any sort of meaningful comparison.

Figure OA15. Yearly Correlation of Industry Shares (RL vs. BDS)

Notes: This panel of plots displays the Spearman (Panel (a)) and Pearson (Panel (b)) correlation measures between the industry share vectors of new firms in RL and BDS for each year (2006–2022). Each point on the curve reflects how closely the fraction of new firms in each industry, as recorded by RL, aligns with the corresponding fractions from BDS in that year.

Another way to assess industry coverage is to compare the share of new firms by industry directly between datasets. Figure OA16 presents a scatter plot of industry shares in BDS versus RL. Industries are plotted for the period 2008 to 2022. Each point represents a single industry, plotting its percentage share of total new firms in BDS on the horizontal axis against the same share in RL on the vertical axis. The points cluster tightly around the 45-degree line, confirming that the two sources give very similar prominence to each industry. Industries that account for a large share of new firm births in the official data also tend to account for a large share in RL, and vice versa. For instance, if an industry constitutes about 10% of all new firms according to BDS, RL typically records a very similar share for that industry.

Most industries lie close to this line, indicating that their share of new firm activity is nearly the same in RL as in BDS. Industries above the line are represented slightly more in RL than in BDS, whereas industries below the line have a somewhat smaller share in RL. For example, RL shows a somewhat higher share of new firms in certain professional/technical sectors (above the line) and a slightly lower share in some traditional sectors like construction or food services (below the line). Overall, however, the deviations are minor and the high density of points near the diagonal demonstrates RL’s strong fidelity to the true industry distribution of new firms.

Figure OA16. Industry Share of New Firms: RL vs. BDS

Notes: This scatter plot compares the distribution of new firms across industries in RL and BDS. Each dot is an industry-year observation (using a consistent industry classification in both datasets), with the horizontal coordinate being that industry’s share of total new firms in BDS and the vertical coordinate its share in RL. The dashed 45-degree line represents equality.

While the overall industry distribution in RL is close to that of BDS, there are some intuitive differences in the relative shares of specific sectors. Table OA24 lists the top 10 industries (by share of total new firms) in the RL data and compares them with the top 10 in BDS. We see that many of the same industries dominate new firm formation in both datasets (indeed, six industries appear in the top ten of both RL and BDS). In particular, Professional, Scientific & Technical Services (NAICS 541) is the largest category of new firms in both datasets; however, RL attributes an even larger fraction of startups to this sector (about 24.8% of all new firms) compared to BDS (15.3%). This suggests that RL may over-sample new firms in highly knowledge-oriented or white-collar sectors, likely because those firms (and their employees) have a greater digital footprint. Conversely, some industries that rank highly in the official data have a somewhat smaller presence in RL’s top ranks. For example, NAICS 722 (Food Services & Drinking Places, which includes restaurants and

bars) accounts for around 11.3% of new firms in BDS (making it the second-largest category in BDS), but constitutes only about 5.8% of new firms in RL. Similarly, construction-related industries (e.g., NAICS 236 and 238, various construction contractors) and certain local service industries (e.g., Repair & Maintenance, NAICS 811) feature prominently in BDS's startup activity but are less represented in RL's top ten. These differences are understandable: new restaurants, construction firms, and repair shops are often smaller and less likely to have employees with online professional profiles, meaning RL is more prone to miss them. Nonetheless, RL still captures many new firms in these sectors (they do appear in RL, just at slightly lower ranks), and importantly, the broad composition is aligned. The rank-order correlation of industry importance between RL and BDS remains extremely high (as seen in Figure OA15), even if the exact shares differ a slightly for certain industries.

Table OA24. Top 10 Industries by New Firm Share: RL vs. BDS

Spot	Revelio Labs	Business Dynamics Statistics
1	541 (24.751%)	541 (15.283%)
2	611 (8.099%)	722 (11.250%)
3	722 (5.783%)	238 (7.878%)
4	561 (5.518%)	621 (6.962%)
5	621 (4.276%)	561 (6.459%)
6	531 (3.271%)	531 (5.696%)
7	522 (2.466%)	236 (4.912%)
8	512 (2.238%)	812 (4.454%)
9	423 (2.103%)	484 (3.111%)
10	812 (2.000%)	811 (3.072%)


Notes: This table compares the leading industries for new firm formation in RL and BDS. For each rank (1 = highest share of new firms), we report the NAICS industry code and its percentage share of total new firms in that dataset.

To elaborate further, we also verify that RL adequately represents the industries that are particularly relevant to our study on AI and entrepreneurship. The main paper focuses on the impact of AI, which may vary by industry, so it is important that RL provides good coverage in both highly AI-exposed sectors and less AI-exposed ones. Figure OA17 and Figure OA18 address this by comparing RL and BDS in the subsets of industries with the highest and lowest *AI Assistance* scores, respectively (these scores, defined in the main text, measure the potential for AI to assist with tasks in each industry). Figure OA17 looks at the top industries (those most likely to benefit from AI). We focus on the top industries in this AI-relevance ranking to ensure that RL appropriately captures new firm activity in the sectors most pertinent to our analysis of AI spillovers. The comparison reveals that RL and BDS assign very similar shares to each high-AI industry. For example, if an AI-intensive industry constitutes $X\%$ of all new firms in BDS, the corresponding share in RL is also


around $X\%$, with only minor deviations. This indicates that even in the technologically advanced sectors where one might particularly want to track entrepreneurship (due to AI's prominence), RL provides a representative sample of new firm creation.

Figure OA18 covers the bottom industries (those least likely to be affected by AI in the near term). Generally, the shares are comparable, indicating that RL does capture new firm formation even in sectors that are not technology-intensive. There are some modest gaps for certain industries in this group: for instance, the Food Services industry (a low-AI sector) constitutes a smaller share of new firms in RL than in BDS (consistent with Table OA24, where NAICS 722 is underrepresented in RL relative to BDS). Likewise, some construction and manual-service industries show lower shares in RL. These differences reflect the limitations of RL in covering very traditional sectors, but importantly they are not large enough to change the overall picture. Even for industries where AI plays little role, RL data still registers significant numbers of new firms and tracks the general distribution of new firm activity seen in the official data. This suggests that our use of RL data is not omitting entire classes of low-tech entrepreneurship.


All in all, in both groups, RL's representation of new firms closely matches that of BDS.

Figure OA17. New Firm Shares in Top AI Assistance Industries

(a) Web Search Portals, Libraries, Archives, and Other Information Services (b) Computing Infrastructure Providers, Data Processing, Web Hosting, and Related Services




(c) Professional, Scientific, and Technical Services (d) Computer and Electronic Product Manufacturing



(e) Broadcasting and Content Providers (f) Telecommunications

Notes: This figure considers the set of industries with the highest AI Assistance Scores (i.e., industries where a large portion of occupational tasks can potentially be augmented by AI, as defined in the main paper). It shows the share of total new firms accounted for by each of these AI-intensive industries, as recorded in BDS (blue bars) and in RL (red bars).


Figure OA18. New Firm Shares in Bottom AI Assistance Industries

Notes: This figure examines industries with the lowest AI Assistance Scores (industries where AI is expected to have minimal immediate impact on tasks). We compare the share of new firms in these low-AI industries between BDS (blue) and RL (red).

C.2.2. New Firm Counts and Trends: RL vs. BDS. Beyond industry composition, we assess whether RL can replicate the overall volume and trends of new firm creation found in official data. In the main text, we highlighted that RL data can be used to analyze year-to-year changes in startup activity. Here we confirm that the ups and downs in new firm counts

observed in BDS are reflected in RL as well. Figure OA19 presents the correlation between RL and BDS in terms of the *number of new firms by industry*, year by year. In these plots, we take each industry’s count of new firms in a given year and measure how well RL’s counts align with BDS’s. The patterns are similar to the share-based correlations shown earlier. The results indicate that RL is capturing not just the proportions but the actual magnitudes of variation in new firm births across industries.

Figure OA19. Yearly Correlation of New Firm Counts by Industry (RL vs. BDS)



Notes: This panel of plots displays the Spearman (Panel (a)) and Pearson (Panel (b)) correlation measures between the industry new firm count vectors of new firms in RL and BDS for each year (2006–2022). Each point on the curve reflects how closely the number of new firms in each industry, as recorded by RL, aligns with the corresponding count from BDS in that year.

To quantify the relationship between RL and BDS new firm counts more directly, we regress the RL counts on the BDS counts. Table OA25 reports the results of this exercise, using a panel of industry-by-year observations. Column(1) shows a simple regression of $\log(\text{New Firms RL})$ on $\log(\text{New Firms BDS})$ (with no fixed effects). The coefficient is about 0.57 and highly significant, implying that a 10% increase in the number of new firms in the official data is associated with roughly a 5.7% increase in the number of new firms captured by RL. This coefficient being below 1.0 is expected, since RL does not capture every new firm—for example, if BDS counts double (often driven by many very small firms), RL counts will increase but by a somewhat smaller proportion, reflecting its incomplete coverage of the smallest or least visible startups. In Column (2), we add industry and year fixed effects to account for any systematic level differences across industries and trends over time. With these controls, the coefficient falls to about 0.21, suggesting that within a given industry and year, deviations in RL from the average are still positively correlated with deviations in BDS, but much of the variance is absorbed by fixed effects. The very high R^2 (0.992) in Column (2) indicates that once we control for industry and year, RL’s counts line up almost perfectly

with BDS for explaining the remaining variation. In practical terms, this means that apart from a stable industry-specific under- or over-counting factor, RL captures fluctuations in new firm creation very reliably. This is reassuring information, as all specifications employed in the main text leverage both industry and time fixed effects, which will absorb these immutable discrepancies.


Table OA25. Regression of RL New Firm Counts on BDS New Firm Counts

VARIABLES	(1) Log N. New Firms (RL)	(2) Log N. New Firms (RL)
Log N. New Firms (BDS)	0.572*** (0.014)	0.209*** (0.018)
Observations	1,273	1,273
Industry 3D	No	Yes
Year FE	No	Yes
Adj. R sq.	0.569	0.992

Notes: This table reports OLS regression results assessing the relationship between the number of new firms recorded in RL and in BDS. The unit of observation is an industry-year cell, where industry is measured at NAICS three-digit NAICS. The time frame is 2006–2022. Column (1) reports the results from the simple regression, while in Column (2) we include both industry and year fixed effects. Robust standard errors are reported in parentheses. Coefficients are statistically significant at the 10% (*), 5% (**), and 1% (***) levels.

Finally, Figure OA20 provides a visual summary of the correspondence between RL and BDS new firm counts. It plots the logged number of new firms in RL against the logged number of new firms in BDS for each industry-year observation, along with a fitted regression line. The points lie close to the fitted line, reaffirming a tight relationship. However, the slope of the fitted line is clearly less than 1 (consistent with the coefficient estimates in Table OA25), which is visually apparent, as many points fall below the 45-degree line (shown for reference). Points below the diagonal mean that for those observations RL reports fewer new firms than BDS. This systematic gap again reflects the partial coverage of RL: for a given industry and year, RL typically counts somewhat fewer startups than exist in reality. Importantly, though, the alignment is strong: the correlation between log counts is high, and there are no extreme outliers where RL indicates a booming startup sector while BDS does not, or vice versa. Every industry-year with high startup activity in BDS is also one of the higher ones in RL.

Figure OA20. Relationship Between RL and BDS New Firm Counts (Industry-Year)



Notes: This scatter plot compares the log number of new firms recorded in RL and BDS across industries and years. Each dot represents a particular industry in a particular year, with the horizontal axis showing the log of the number of new firms in that industry-year according to BDS and the vertical axis showing the log of the number according to RL. The solid line is the OLS fit from Column (1) of Table OA25, and the 45-degree line (dashed) is shown for reference.

Summary and Caveats. In summary, the RL data on new firm creation exhibits patterns that mirror those found in the official BDS statistics. It captures the same cyclical and trend movements in startup activity over time, and it reproduces the cross-sectional distribution of new firms across industries with a high degree of fidelity. These findings support the notion that RL is—to a good degree—a suitable and representative dataset for studying firm creation in the United States. Researchers can leverage RL data to analyze entrepreneurial dynamics and infer insights similar to those obtained from traditional sources, with the added advantages that RL data are more granular (e.g., providing micro-level information on firm and worker characteristics) and more up-to-date (at the time of our study, BDS data are available only with a lag of about two years, whereas RL provides information in near real-time).

One additional technical caveat concerns the BDS dataset itself: in a small number of cases, BDS withholds firm birth counts in cells where the number of new firms is positive but very small, typically for confidentiality reasons. This means that for a few rare industry-year combinations, our benchmark BDS data may report a zero even though some firm creation occurred. However, because our validation exercises are conducted at the three-digit NAICS by year level—an aggregation that is macro enough to minimize such suppression—these occurrences are infrequent and have negligible influence on the overall representativeness assessments documented here.

Lastly, it is important to note that RL’s representativeness has been demonstrated here specifically for the domain of new firm births. RL is effectively a large sample of the population of new firms, but not a complete census. Its coverage, as we have seen, is somewhat thinner for very small firms and certain low-tech industries. As a result, RL might not be fully representative if one’s research question is, say, to measure the absolute level of entrepreneurship in every niche of the economy or to capture metrics beyond firm creation (such as total employment or revenue growth, where official data would still be the gold standard). For our purposes, though, which center on the rates and patterns of firm creation (especially in relation to technological change), RL proves to be an excellent proxy. The evidence above – strong correlations, similar distributions, and aligned trends – all point to RL being a valid tool to study firm creation in lieu of or in complement to BDS. This underpins the analyses in the main paper, giving us confidence that the insights derived from the RL data reflect real-world entrepreneurial dynamics.

C.3. Measures Creation Using Revelio Labs. To construct the measure of firm creation, we proceed as follows. First, we extract all U.S. companies with both a valid NAICS code and a valid state identifier. Next, for each company, we retrieve all worker positions that were ever employed at that company, identified by the unique company identifier (`rcid`).

To estimate the date of establishment, we take the earliest employment start date among all positions associated with that company’s `rcid`. This establishment date is measured at the quarter-year level.

To determine the date of shutdown²⁷, we calculate, for each company, the total number of active workers in every quarter. If a company records two or more consecutive quarters with zero active workers, we define the shutdown date as the last quarter in which the company had at least one active worker.

Lastly, we calculate the number of new firms by collapsing by quarter and naics code.

Our main outcome variable in the analysis is a quarterly measure of the number of newly created firms at the industry level. In Subsection C.2, we demonstrated that when aggregated at the three-digit NAICS level, the share of relative share of new firms in each industry is

²⁷We use this piece of information in the firm level analysis of the paper presented in Section 6.

largely in line with official statistics. However, since RL is a subsample of the true population of US firms, it does not match the levels of firm creation of official statistics. To this end, we rescale these shares to match the official aggregate level of firm creation. The new measure is constructed by combining the new firm count (i) firm incorporations calculated from Revelio Labs (RL) data and (ii) official aggregate statistics on business applications from the U.S. Census Bureau, obtained through the FRED series `BUSAPPWNSAUS`.

The scaling procedure works as follows. Let i index industries and t index quarters. For each industry-quarter cell, we compute the RL share of incorporations as

$$s_{it} = \frac{\text{N. New Firms}_{it}^{\text{RL}}}{\sum_i \text{N. New Firms}_{it}^{\text{RL}}}$$

where $\text{N. New Firms}_{it}^{\text{RL}}$ is the raw count of incorporations observed in RL for industry i at time t . This object s_{it} represents the share of incorporations accounted for by three-digit NAICS industry i within quarter t .

From FRED we obtain the series `BUSAPPWNSAUS`, which measures the total number of business applications at the national level. We collapse this series to the quarterly frequency and denote it by

$$\text{N. New Firms}_t^{\text{FRED}}$$

which is the aggregate number of new firms in quarter t .

To obtain an industry-quarter level measure of firm creation that is both consistent with the aggregate time series and preserves the cross-industry variation from RL, we rescale the RL shares by the official aggregate:

$$\text{N. New Firms}_{it} = s_{it} \times \text{N. New Firms}_t^{\text{FRED}}$$

This rescaled measure, N. New Firms_{it} , can be interpreted as the number of new firms created in industry i and quarter t , consistent with the aggregate level of firm creation reported in official statistics. It effectively anchors the RL industry distribution to the true total volume of incorporations, correcting for the fact that RL does not capture all firms but provides reliable information on their relative allocation across industries.

APPENDIX D. GENERATION OF MEASURES USING LLMs

This section details the generation process of the LLM-generated measurements. First, we present the textual data sources and then report the prompts used in each of the LLM layers, which use ChatGPT 4.0 as AI model.

D.1. Textual Data Sources. Table OA26 lists the textual materials used to identify and classify entrepreneurial tasks. The corpus spans government guides, academic handbooks, and practitioner-oriented works that together describe the key stages of venture creation and early operations. For each source, we report its type, a brief description of its purpose, and how its content contributes to the compilation of task statements used in our analysis.

Table OA26. Textual Data Sources Used for Task Extraction

Source	Type	Description	Contribution to Task Compilation
<i>10 Steps to Start Your Business</i> (U.S. SBA, 2020)	Government guide	Official step-by-step primer for starting a small business in the U.S. (entity choice, licensing, financing basics).	Yields boilerplate, industry-agnostic tasks for formation, registration, licensing, EIN/taxes, basic plans, and early financing.
<i>Small Business Resource Guide</i> (IRS, 2021)	Government publication	Practical orientation to tax compliance and record-keeping for small businesses.	Adds concrete admin/finance tasks: bookkeeping, quarterly/annual filings, payroll setup, forms, compliance calendars.
<i>Guide to Starting and Operating a Small Business</i> (Michigan SBDC, 2015)	Government guide	Comprehensive small-business playbook (planning, licensing, finance, operations) for non-tech SMBs.	Provides generalized—but detailed—operational setup tasks (permits, insurance, vendor setup, local compliance).
<i>Disciplined Entrepreneurship: 24 Steps</i> (Aulet, 2013)	Academic handbook	Structured 24-step method for opportunity validation and venture design.	Rich source of validation tasks: segmentation, persona, beachhead market, TAM, pricing, go-to-market experiments.

(continued on next page)

(continued from previous page)

Source	Type	Description	Contribution to Task Compilation
<i>The Lean Startup</i> (Ries, 2011)	Startup methodology	Build–Measure–Learn cycles and MVP-driven product/market learning.	Supplies experimentation tasks: MVP definition, metric selection, cohort tests, iteration/pivot decisions.
<i>The Startup Owner’s Manual</i> (Blank & Dorf, 2012)	Practitioner manual	Customer development field guide (discovery, validation, creation, company building).	Adds discovery/validation tasks: interview protocols, hypotheses, experiment design, evidence-based model updates.
<i>HBR Entrepreneur’s Handbook</i> (HBR, 2018)	Practitioner handbook	Broad reference across planning, fundraising, teams, marketing, and operations.	Contributes cross-functional tasks (pitch materials, early hiring/processes, channel tests, basic dashboards).
<i>Prototype It</i> (Savoia, 2011)	Methodology (prototyping)	Ultra-fast, low-fidelity testing to de-risk ideas pre-MVP.	Yields micro-experiments: fake-door tests, landing pages, ad tests, concierge/Paper-MVP tasks.
<i>The Hard Thing About Hard Things</i> (Horowitz, 2014)	Practitioner reflections	Operating under uncertainty; leadership, org design, tough calls in scaling.	Provides managerial/operational tasks: hiring frameworks, org changes, crisis comms, decision logs, exec processes.
<i>The Startup Playbook</i> (Altman, 2015)	Practitioner playbook	Concise guidance from early-stage investing/operator perspective.	Adds founder-centric tasks: founder-market fit checks, fundraising prep, milestone setting, early growth priorities.
<i>Zero to One</i> (Thiel, 2014)	Entrepreneurial essay	Strategy for building monopolies via differentiated, defensible products.	Contributes strategic tasks: unique value articulation, defensibility mapping, market selection, moats/PMF hypotheses.

D.2. Workflow Infrastructure. This subsection details the generation process of the LLM-generated measurements. All layers of the workflow utilize ChatGPT 4.0 as AI model. For each

prompt employed, we first explain the task it is designed to achieve, and then report the exact form in which it is present to the AI virtual assistant.

D.2.1. LLM Layer 1: Task Extraction. The first layer of the workflow is responsible for identifying concrete, actionable entrepreneurial tasks from unstructured business texts. Its role is to convert narrative descriptions (such as guidance from entrepreneurship manuals, industry guides, or startup competition documents) into a structured set of tasks. A task is considered valid only if it is executable, measurable, specific, actionable, business-relevant, and time-bound. The model is prompted to reject abstract advice (e.g., “be innovative”) and retain only concrete actions (e.g., “file LLC registration documents with the state secretary”). Each extracted task is returned with a task name, a short description of the deliverable, the source text, and a rationale explaining why the item qualifies as a task. To ensure reliability, the pipeline validates that each extracted task contains the required fields and implements retry logic when the model output is empty or malformed. The output of this layer is a structured list of validated tasks with associated metadata, which are saved for use in subsequent layers. The exact system prompt used is reproduced below:

You are an expert in entrepreneurship and business operations, with deep knowledge of the CONCRETE, ACTIONABLE steps required to start and grow a new firm. You specialize in analyzing unstructured business texts and identifying EXECUTABLE tasks, processes, and requirements involved in launching a company.

Your task is to identify MULTIPLE CONCRETE, ACTIONABLE entrepreneurial tasks described in the text provided.

Extract as many distinct EXECUTABLE tasks as possible from the given text.

CRITICAL DEFINITION: A task MUST be a specific, measurable, executable action that an entrepreneur can perform within a defined timeframe. Examples:

GOOD: "File LLC registration documents with state secretary", "Conduct 50 customer interviews", "Create financial projections for first 18 months"

BAD: "Build strong relationships", "Focus on customer satisfaction", "Maintain work-life balance", "Be innovative"

STRICT TASK CRITERIA:

1. EXECUTABLE: Must be a concrete action with clear start/end points
2. MEASURABLE: Should have quantifiable outcomes or deliverables
3. SPECIFIC: Not abstract concepts, principles, or general advice
4. ACTIONABLE: Something an entrepreneur can actually DO, not just think about
5. BUSINESS-RELEVANT: Directly related to starting/operating a business
6. TIME-BOUND: Can be completed within a reasonable timeframe (days to months, not years)

WHAT TO EXTRACT vs WHAT TO REJECT:

EXTRACT: "Register business entity", "Obtain business license", "Open business bank account", "Hire first employee", "Launch MVP product"

REJECT: "Think strategically", "Be customer-focused", "Stay motivated", "Network effectively", "Embrace failure"

VALIDATION CHECKLIST for each task:

- Can I delegate this to someone with clear instructions? (YES = extract)
- Does this have a clear deliverable or outcome? (YES = extract)
- Is this a specific action rather than a mindset/principle? (YES = extract)
- Would this appear on a business startup checklist? (YES = extract)

Guidelines for identifying tasks:

1. Focus ONLY on EXECUTABLE ACTIONS, reject abstract concepts
2. Tasks must be concrete, measurable, and time-bound
3. Include tasks across different business stages (idea, validation, launch, growth)
4. Consider industry-specific operational tasks
5. EXTRACT EVERY EXECUTABLE TASK from the text, but REJECT all abstract advice
6. Use action verbs: "Create", "File", "Register", "Hire", "Launch", "Develop", "Conduct"

For each EXECUTABLE task you identify, provide:

1. A concise task name using strong action verbs (e.g., "Register Business Entity", "Conduct Market Research Survey")
2. A detailed description focusing on the specific deliverable/outcome (1-2 sentences)
3. The exact source text that mentions this executable task
4. A rationale explaining why this is a concrete, actionable entrepreneurial task

RESPONSE FORMAT

Provide ALL EXECUTABLE TASKS in a SINGLE string. Each task's details should be separated by a semicolon (;), and different tasks should be separated by a pipe (|).

Example format for multiple tasks:

task_name1; description1; source_text1; rationale1 | task_name2; description2; source_text2; rationale2 | task_name3; description3; source_text3; rationale3

Field order within each task block:

1. task_name
2. description
3. source_text
4. rationale

IMPORTANT:

- Be extremely selective - QUALITY over quantity
- REJECT abstract concepts, principles, mindsets, and general advice
- ONLY extract concrete, executable actions with clear deliverables
- Avoid duplicating tasks
- Focus on ACTIONABLE entrepreneurial steps, not business philosophy
- Each task block MUST contain exactly 4 fields separated by semicolons (;)
- Separate different task blocks with a pipe (|)
- Ensure consistent field order within each block

WARNING

THE NUMBER OF ITEMS IN EACH PIPE-SEPARATED LIST MUST BE THE SAME (match the number of tasks extracted)

REMEMBER: If it's not a concrete, executable action with a clear deliverable, DON'T extract it!

D.2.2. *LLM Layer 2: Task Deduplication.* The second layer of the workflow standardizes task names and consolidates semantically similar tasks across documents. Its objective is to ensure that the set of entrepreneurial tasks is not only actionable but also consistently represented, avoiding duplication and inconsistencies in naming. Each batch of tasks produced by the extraction layer is examined to identify potential overlaps, and similar tasks are merged under a standardized name with a consolidated description. Tasks deemed too abstract or non-executable are flagged for removal. The output is a cleaned, deduplicated set of tasks with standardized names, clear descriptions, and metadata that preserve their source information. This ensures comparability across sources and prepares the task list for subsequent industry relevance checks. The exact system prompt used is reproduced below:

You are an expert in entrepreneurship and business analysis specialized in EXECUTABLE task consolidation and naming standardization.

Your job is to analyze a group of CONCRETE, ACTIONABLE entrepreneurial tasks, consolidate similar ones, and standardize their names and descriptions. You must examine BOTH the name AND description of each task to identify semantic similarities.

CRITICAL: Only work with EXECUTABLE, CONCRETE tasks. If any task appears to be abstract advice, principles, or general concepts rather than actionable steps, flag it for removal.

TASK VALIDATION CRITERIA:

KEEP: Tasks that are specific, measurable, executable actions with clear deliverables

REMOVE: Abstract concepts, general advice, principles, mindsets, or vague suggestions

GUIDELINES FOR TASK CONSOLIDATION AND STANDARDIZATION:

- Use concise, imperative verb phrases for task names (e.g., "Register Business Entity", "Conduct Customer Interviews", "File Tax Documents")
- Ensure consistent naming convention (capitalization, tense, etc.)
- Maintain specific, actionable language that clearly describes the EXECUTABLE task
- Group similar EXECUTABLE tasks using consistent naming patterns
- Always begin task names with a strong action verb in the imperative form
- Create consolidated descriptions that capture the full scope of similar EXECUTABLE tasks
- Preserve important operational details from all similar tasks in the consolidated description
- Focus on concrete deliverables and measurable outcomes

IMPORTANT CONSOLIDATION INSTRUCTIONS:

- Pay careful attention to both task names AND descriptions when identifying similar EXECUTABLE tasks
- Two tasks might have different names but describe essentially the same ACTIONABLE activity - these should be consolidated
- Consider semantic similarity of EXECUTABLE actions, not just exact matching of words
- Look for tasks that serve the same business purpose through concrete actions
- Subtasks of a larger process should maintain their specificity while following consistent naming patterns
- REJECT any tasks that are not concrete, executable actions

If two or more EXECUTABLE tasks are essentially the same (based on BOTH their names AND descriptions):

1. Standardize their names to be the same using action verbs

2. Create a consolidated description that combines the most important OPERATIONAL elements from all similar tasks
3. Make sure the consolidated description focuses on specific deliverables and measurable outcomes
4. Ensure the result is still a concrete, actionable task

For tasks that are unique EXECUTABLE actions, ensure their names are clear and follow the standardized format.

IMPORTANT: Review the list of previously standardized task names provided and try to maintain consistency with those naming patterns when relevant. This ensures that similar EXECUTABLE tasks across all batches use the same standardized names. Focus on consolidating similar EXECUTABLE tasks WITHIN the provided batch.

QUALITY CONTROL: If any task in the batch appears to be abstract advice rather than an executable action, mark it for removal in your response by setting `new_name` to "REMOVE_ABSTRACT_TASK".

D.2.3. *LLM Layer 3: Assessing Industry Relevance and Adapting Task Description to Industry.*

The third layer evaluates whether each actionable entrepreneurial task is relevant to a specific industry. Its purpose is to ensure that the task set reflects not only general entrepreneurial actions but also those that are operationally meaningful within a given sector. Each task is paired with a detailed industry description, and the model is asked to determine relevance using a binary (yes/no) assessment. For tasks deemed relevant, the model generates an industry-adapted description that tailors the generic task to sector-specific processes, terminology, and requirements. This adaptation accounts for industry regulations, operational workflows, and customer needs. Tasks judged irrelevant are removed from the pipeline. The output of this layer is a curated set of industry-specific, actionable tasks with standardized descriptions that capture sectoral context. The exact system prompt used is reproduced below:

You are an expert in industry analysis and entrepreneurship. Your task is to evaluate whether a specific EXECUTABLE entrepreneurial task is relevant to a given industry and create an industry-adapted description.

CRITICAL: You are evaluating CONCRETE, ACTIONABLE tasks only. These should be specific actions with clear deliverables, not abstract concepts or general business advice.

You will receive information about:

1. An EXECUTABLE task that entrepreneurs perform when starting or growing a business (specific actions with measurable outcomes)
2. An industry context with detailed industry description to evaluate the task within

You must:

1. Determine if this EXECUTABLE task is RELEVANT for entrepreneurs operating in this specific industry
2. If relevant, create an industry-adapted description that tailors the ACTIONABLE task specifically to that industry

ASSESSMENT PROCESS:

1. First, verify the task is concrete and executable (has clear deliverables, measurable outcomes, specific actions)
2. Analyze how this specific executable task applies to the industry context
3. Consider the specific operational characteristics, regulatory requirements, and business processes of the industry
4. Develop a chain of thought reasoning about WHY this EXECUTABLE task would or would not be operationally important in this industry
5. Provide a binary YES/NO decision on relevance for operational execution
6. If relevant, create a new task description that specifies HOW this task is executed in this specific industry context

GUIDELINES:

- Relevant tasks (YES) are EXECUTABLE actions that entrepreneurs in this industry typically need to perform operationally
- Irrelevant tasks (NO) are those that have minimal operational importance in this industry context
- Consider industry-specific regulations, operational processes, customer requirements, supplier relationships, and compliance needs
- Focus on OPERATIONAL RELEVANCE: How does this specific action apply to day-to-day business operations in this industry?
- An executable task can be generally important but still not operationally relevant to a specific industry's workflow
- Provide detailed reasoning that shows your analysis of both the EXECUTABLE task and industry operations
- The industry-adapted description should incorporate specific industry processes, terminology, regulatory requirements, and operational contexts

- Ensure the adapted description maintains the executable, actionable nature of the original task

RESPONSE FORMAT:

```
{
  "reasoning": "Your detailed chain of thought reasoning for why this
EXECUTABLE task is or is not operationally relevant to this industry,
focusing on specific business processes and requirements", "is_relevant":
true/false (true if operationally relevant, false if not),
  "industry_adapted_description": "For relevant tasks, provide a description
that specifies HOW this executable task is performed in this industry
context, including industry-specific processes, requirements, and outcomes.
For irrelevant tasks, leave this empty."
}
```

IMPORTANT:

- Be objective and consistent in your operational assessments
- Your reasoning must clearly demonstrate understanding of both the EXECUTABLE task and industry operations
- Make a clear binary decision based on operational relevance
- The industry-adapted description should be comprehensive, industry-specific, and maintain the actionable nature
- Focus on concrete operational requirements and industry-specific execution methods
- Use the detailed industry description provided to inform your task adaptation with specific operational context

D.2.4. *LLM Layer 4: Task Scoring.* The fourth layer of the workflow quantifies each validated task along three key dimensions. First, *Task Weight* captures how critical the task is for launching a new venture within the specific industry, accounting for regulatory, operational, and market requirements. Second, the *AI Assistance Score* evaluates the extent to which current large language models can meaningfully support entrepreneurs in executing the task. Finally, *Experimentation Intensity* measures the degree to which the task depends on trial-and-error, creative iteration, or uncertain feedback loops rather than standardized procedures. Each score is expressed on a 0–100 scale, with detailed rubrics ensuring consistency across tasks. The model also generates justifications for each assigned value, grounding the scoring in industry and task-specific reasoning. The output is a structured dataset of tasks with associated quantitative scores and explanations, forming the basis for constructing our industry-level measures of AI Assistance and Experimentation Intensity. The exact system prompt used is reproduced below:

You are an expert in entrepreneurship and business analysis with deep knowledge of industry-specific startup requirements.

Analyze the CONCRETE, EXECUTABLE entrepreneurial task provided and score it on the following dimensions using the detailed rubric below.

IMPORTANT: You are analyzing ACTIONABLE tasks only --- specific actions with clear deliverables, measurable outcomes, and concrete execution steps. These are NOT abstract concepts, principles, or general business advice.

SCORING DIMENSIONS & DETAILED RUBRIC

1. TASK WEIGHT (0--100): Industry-Specific Importance

Definition: How critical this task is for entrepreneurs in the given industry when launching a new firm.

- 0--25: Peripheral or optional
- 26--50: Helpful but not essential
- 51--75: Important for success or efficiency
- 76--100: Mission-critical or legally required

2. AI ASSISTANCE SCORE (0--100): LLM Capability Assessment

Definition: The extent to which current AI systems (e.g., GPT-4, Claude) can assist in completing the task.

- 0--25: Minimal AI help
- 26--50: Limited support (templates, guidance)
- 51--75: Substantial assistance (drafting, analysis, recommendations)
- 76--100: Near-autonomous execution with minimal human input

3. EXPERIMENTATION INTENSITY (0--100): Trial-and-Error Dependency

Definition: How much the task depends on iterative, uncertain, or experimental processes.

- 0--25: Highly standardized, predictable
- 26--50: Mostly procedural with some iteration
- 51--75: Substantial experimentation required
- 76--100: Highly experimental and creative, heavily dependent on iteration

###EXAMPLE 1

INPUT

#Task Name:

Register business

#Task Description:

Formally register the business entity with relevant local, state, or federal authorities, including selecting a legal structure and submitting required documentation.

#Industry Name:

Food Manufacturing

#Industry Description:

Industries in the Food Manufacturing subsector transform livestock and agricultural products into products for intermediate or final consumption. The industry groups are distinguished by the raw materials (generally of animal or vegetable origin) processed into food products. The food products manufactured in these establishments are typically sold to wholesalers or retailers for distribution to consumers, but establishments primarily engaged in retailing bakery and candy products made on the premises not for immediate consumption are included. Establishments primarily engaged in manufacturing beverages are classified in Subsector 312, Beverage and Tobacco Product Manufacturing.

OUTPUT

#Task Weight Score:

95

#Task Weight Justification:

Legal registration is required to operate officially, access banking, and hire staff---this is a near-universal startup task.

#AI Assistance Score:

60

#AI Assistance Score Justification:

AI can summarize requirements, suggest business structures, and help complete forms, but human oversight is needed for compliance and submission.

#Experimentation Intensity:

10

#Experimentation Intensity Justification:

This is a highly standardized process with clear steps and a predictable outcome.

###EXAMPLE 2

INPUT

#Task Name:

Develop brand identity

#Task Description:

Create a distinctive brand identity for the business, including its name, visual style, messaging, and overall tone to resonate with the target audience.

#Industry:

Furniture and Related Product Manufacturing

#Industry Description:

Industries in the Furniture and Related Product Manufacturing subsector make furniture and related articles, such as mattresses, window blinds, cabinets, and fixtures. Design and fashion trends play an important part in the production of furniture. The integrated design of the article for both aesthetic and functional qualities is also a major part of the process of manufacturing furniture.

OUTPUT

#Task Weight Score:

75

#Task Weight Justification:

Brand identity is crucial for furniture manufacturing where design and aesthetic appeal directly influence consumer purchase decisions and market differentiation in a competitive industry.

#AI Assistance Score:

70

#AI Assistance Score Justification:

AI can generate brand concepts, create visual designs, suggest messaging strategies, and provide market analysis, though human creativity and market intuition remain important for final decisions.

#Experimentation Intensity:

85

#Experimentation Intensity Justification:

Brand development requires extensive testing of different concepts, market feedback, iteration on visual elements, and refinement based on target audience response---highly experimental and creative process.

SCORING GUIDELINES

Industry Context Considerations:

- Always consider how industry-specific regulations, customer expectations, and market dynamics affect task importance
- Highly regulated industries (healthcare, finance, food) often have higher task weights for compliance-related activities
- Creative industries may score higher on experimentation intensity for design and innovation tasks
- B2B vs B2C differences significantly impact marketing and customer acquisition task weights

AI Capability Assessment:

- Consider current LLM limitations: cannot make phone calls, physical site visits, or handle confidential financial transactions
- AI excels at: research, content creation, data analysis, template generation, initial drafts
- AI struggles with: nuanced human judgment, regulatory compliance decisions, relationship building, complex negotiations

****Experimentation Assessment:****

- Linear tasks: legal filings, standard accounting procedures, basic market research
- Moderate experimentation: product development, marketing campaigns, pricing strategies
- High experimentation: R&D, new market entry, innovative business models, creative design work

###RESPONSE FORMAT

Your response should be a JSON with these field: "task name":"Name of the Task", "task description":"Detailed task description", "industry":"Industry used as context to evaluate the task, "task weight":"value assigned to task weight", "task weight justification":"Why task in industry of context was given specific value", "ai assistance score":"value assigned based on AI's ability to assist in task","experimentation intensity":"value assigned based on prominence of experimentation and trial and error when carrying out the task"

APPENDIX E. SURVEY OF NEW AND ASPIRING ENTREPRENEURS

This appendix provides the full details of the survey of new and aspiring entrepreneurs introduced in the main text. Subsection E.1 reproduces the survey instrument. Subsection E.2 describes how the raw survey responses and the LLM-extracted task data are transformed into the analysis-ready measures used in Subsection 4.4 of the main text. Subsection E.3 reports descriptive statistics on the final analytic sample of 1,036 respondents.

E.1. Survey Instrument. The survey was hosted on Qualtrics and disseminated through Dynata’s commercial respondent panel between November 2025 and May 2026. Dynata maintains an opt-in panel of respondents who are compensated for completing surveys. Respondents who do not satisfy the qualifying criteria are screened out at the beginning of the survey and do not complete it. To qualify, respondents must satisfy at least one of the following conditions: (i) they started a company from January 2023 onward, or (ii) they began actively working on a business idea from January 2023 onward. The final analytic sample comprises 1,036 respondents.

The survey is organized in eight blocks. Blocks 1 and 2 collect screening and demographic information, including business start status, the industry of the business or prospective business, age group, and educational attainment. Block 3 elicits open-ended views on how AI tools can help entrepreneurs. Blocks 4 and 5 ask respondents to rate the helpfulness of AI for, and report their actual use of AI in, twenty broadly defined entrepreneurial task categories. Block 6 focuses on six tasks specific to early-stage experimentation, asking respondents how much faster AI tools allow them to perform these tasks. Block 7 covers respondents’ overall reliance on AI tools and the counterfactual importance of AI availability for their entrepreneurial decisions. Block 8 collects optional final feedback.

The full text of the survey questions is reproduced below.

Study Number: IRB25-1533

Study Title: Understanding the Use of Artificial Intelligence Tools by Entrepreneurs

Researchers: Edoardo Marchesi, PhD Student, University of Chicago Booth School of Business; Sixun Tang, PhD Student, University of Chicago Booth School of Business

Description: We are researchers at the University of Chicago doing a research study about how entrepreneurs and aspiring entrepreneurs use artificial intelligence (AI) tools in their business activities. The purpose of this study is to better understand how AI technologies may affect entrepreneurship and business creation. Participation should take about 10 to 12 minutes. Your participation is voluntary.

Confidentiality: Your responses will be anonymous. The survey does not collect your name, email, or other identifying information. Data will be stored securely on the University of Chicago’s Box system, accessible only to the research team (Edoardo Marchesi and Sixun Tang). Results will only be reported in aggregate.

Contacts & Questions: If you have questions about this research, please contact: Edoardo Marchesi (+1 312-358-8174; emarche0@chicagobooth.edu). If you have any questions about your rights as a participant in this research, feel you have been harmed, or wish to discuss other study-related concerns with someone who is not part of the research team, you can contact the University of Chicago Social & Behavioral Sciences Institutional Review Board (IRB) Office by phone at (773) 702-2915, or by email at sbs-irb@uchicago.edu.

[Block 1: Screening]

Q1.1 Over the past *3 years*, have you *started a business* of any kind, regardless of its legal structure or whether it is still active?

- ☐ Yes ☐ No

[Shown only if Q1.1 = Yes]

Q1.2 When did you start your business? Input month and year (MM/YYYY) or use NA if business hasn't been launched yet.

[Open-ended]

[Shown only if Q1.1 = No]

Q1.3 Over the past *3 years*, have you *thought* about starting a business, *experimented* with business ideas, or *taken steps* to start a business?

- ☐ Yes ☐ No

[Shown only if Q1.3 = Yes]

Q1.4 When did you first begin seriously thinking about this business idea? Input month and year (MM/YYYY).

[Open-ended]

[Block 2: Demographics and business characteristics]

Q2.1 In which industry does (or will) your business primarily operate? If you are not sure which industry category is the best fit, look at the industry definitions in the link here.

[Dropdown menu of three-digit NAICS industries]

Q2.2 Is your business currently:

- ☐ In planning/pre-launch stage ☐ Never launched and idea was abandoned ☐ Active and operating ☐ Closed

Q2.3 Please select your age group.

[Dropdown: 21–25, 26–30, 31–35, 36–40, 41–45, 46–55, 56+]

Q2.4 What is the highest level of education you have completed?

- ☐ Less than high school ☐ High school diploma or GED ☐ Some college, no degree
- ☐ Associate degree (e.g., AA, AS) ☐ Bachelor's degree (e.g., BA, BS)
- ☐ Master's degree (e.g., MA, MS, MBA) ☐ Professional degree (e.g., JD, MD, DDS)
- ☐ Doctoral degree (e.g., PhD, EdD) ☐ Other

[Block 3: Open-ended views on AI and entrepreneurship]

Studies suggest that AI tools can help people become more productive in their day-to-day work. By AI tools, we mean software that uses artificial intelligence to assist with tasks such as communication, content creation, design, coding, data analysis, or decision-making. Examples include general-purpose assistants like ChatGPT, Claude, Gemini, or Llama, and domain-specific platforms like Jasper for writing, Lovable for app creation, or Midjourney for design. In the next questions, we want to know your opinion on how AI tools may affect entrepreneurs in the early stages of their business.

Q3.1 In your own words, what are your thoughts on how AI tools can help entrepreneurs like you create or grow new businesses?

[Open-ended]

[Block 4: AI helpfulness across twenty task categories]

On the next screens, you will see a list of tasks that entrepreneurs typically perform when starting and running a business. For each task, please indicate how helpful AI tools could be in completing it. Use the slider to give your rating, where 0 means AI provides no help at all and 100 means AI could fully complete the task on its own. If you have never used AI for a given task, please give your best estimate.

Q4.1 For each of the following tasks, indicate how much AI tools could help you complete the task. If you have never used AI for a given task, please provide your best estimate. On the scale, 0 means that AI tools provide no assistance in completing the task, and 100 means that AI tools can completely take over the task.

[Slider from 0 to 100 for each task]

(a) **Opportunity Identification and Market Understanding:** Figuring out what business opportunity to pursue and who your customers are — for example, identifying an unmet need or problem worth solving, researching who might pay for a solution, sizing up the competition, and validating whether your initial idea has real potential.

(b) **Product and Service Design:** Coming up with what your business will offer and developing it into something customers can actually use — for example, designing the features of a product, defining how a service will be delivered, building early versions to test, and refining the offering based on what you learn.

(c) **Business Model Development:** Figuring out how your business creates value and makes money — for example, deciding what to charge, choosing how to structure your pricing (such as a one-time fee, a subscription, or pay-per-use), identifying your main costs, and determining the different ways your business will generate revenue.

(d) **Strategic Planning and Decision Making:** Setting the overall direction for your business and making the key decisions that shape how it develops — for example, writing a business plan, choosing which markets to focus on, deciding when to change course, and determining how to allocate your time and resources.

(e) **Legal and Regulatory Compliance:** Making sure your business meets all legal requirements to operate — for example, choosing a legal structure, registering your business, obtaining licenses and permits, filing required government paperwork, and protecting your brand or inventions through patents and trademarks.

Q4.2 *[Same prompt as Q4.1, slider from 0 to 100 for each task]*

(a) **Financial Planning and Management:** Managing the financial side of your business to make sure it stays viable — for example, building financial projections, tracking your costs and revenues, managing your cash flow, setting up bank accounts, and making sure your tax obligations are in order.

(b) **Funding and Capital Acquisition:** Identifying how much money your business needs and securing the capital to get started and grow — for example, exploring personal savings, applying for a bank loan, pitching to investors, or researching government grants and funding programs.

(c) **Marketing and Customer Acquisition:** Marketing your product or service to attract and retain customers — for example, developing a marketing strategy, running advertising campaigns, trying different messages or channels to find what works, and building your online presence.

(d) **Sales and Revenue Generation:** Selling your product or service and getting paid — for example, handling customer transactions, collecting payments, converting interested prospects into paying customers, and setting up the processes your business needs to generate revenue consistently.

(e) **Operations and Process Management:** Running the day-to-day activities that keep your business functioning — for example, managing workflows, scheduling production or service delivery, improving how work gets done, and making sure your operations run efficiently and consistently.

Q4.3 *[Same prompt as Q4.1, slider from 0 to 100 for each task]*

(a) **Supply Chain and Procurement:** Acquiring the materials, products, or services your business needs to operate — for example, identifying and selecting suppliers, placing orders, managing inventory, and making sure inputs arrive on time and at the right cost.

(b) **Human Resources and Workforce Management:** Finding, hiring, and developing the people your business needs — for example, recruiting and onboarding new employees, setting compensation, designing training programs, managing performance, and making sure your team has the skills and support to do their jobs effectively.

(c) **Technology and Infrastructure Management:** Setting up and maintaining the technology your business relies on to operate — for example, building your website, setting up online payment systems, managing software tools, establishing communication systems, and making sure your digital infrastructure stays reliable.

(d) **Quality Assurance and Continuous Improvement:** Making sure your product or service meets the standards your customers expect, and continuously improving how your business operates — for example, testing before you launch, identifying what is not working, investigating the root causes of problems, and implementing changes to improve quality and reliability.

(e) **Risk Management and Compliance Monitoring:** Identifying and managing the risks your business faces and making sure it stays compliant with relevant rules and regulations over time — for example, assessing operational risks, monitoring regulatory changes, managing insurance, and developing plans to handle unexpected events.

Q4.4 *[Same prompt as Q4.1, slider from 0 to 100 for each task]*

(a) **Customer Relationship Management:** Maintaining and improving your relationship with existing customers — for example, responding to complaints, tracking whether customers are coming back, understanding how they use your product or service, and making changes to improve their experience and keep them loyal.

(b) **Data Management and Analytics:** Tracking how your business is performing and using that information to make better decisions — for example, monitoring your sales, measuring customer behavior, generating reports, and analyzing data to understand what is working and where there is room to improve.

(c) **Partnership and External Relationship Management:** Building and managing relationships with the external parties your business depends on — for example, identifying and negotiating with suppliers and partners, managing vendor relationships, coordinating with distributors or subcontractors, and maintaining strategic alliances.

(d) **Branding and Reputation Management:** Building and protecting how your business is perceived by customers and the public — for example, developing your brand identity, registering your business name or logo as a trademark, managing your reputation online, and making sure the way you present your business is consistent and recognizable.

(e) **Organizational Governance and Leadership:** Deciding how your business is structured and governed, and making sure it runs according to clear rules and responsibilities — for example, defining who is in charge of what, establishing decision-making policies, setting up the organizational structure, and creating the processes that keep the business accountable and well-managed.

[Block 5: Actual AI use across the same twenty task categories]

For the same tasks we just showed you, we'd like to know whether you have ever used AI when setting up and first running your company.

Q5.1–Q5.4 Please report whether you have ever used an AI tool to complete each of the following tasks.

[Yes/No for each of the twenty tasks listed in Q4.1–Q4.4]

[Block 6: Speed-up across six experimentation-related tasks]

On the next screens, you will see a list of tasks that entrepreneurs perform when testing and refining new ideas — for example, building prototypes, exploring pricing, or gauging customer reactions. For each task, please indicate how much faster you could complete it using AI tools compared to doing it without AI. Use the slider to give your rating, where 0 means AI provides no

speed-up at all and 100 means AI makes the task essentially instant. If you have never used AI for a given task, please give your best estimate.

Q6.1 For each of the following tasks, please indicate *how much faster* you could complete it using AI tools compared to doing it without AI. If you have never used AI for a given task, please provide your best estimate. The slider values are expressed as percentages (for example, 30 means AI allows you to complete the task 30% faster).

[Slider from 0 to 100 for each task]

(a) Figuring out what business to start or what problem to solve — for example, coming up with and discarding different ideas, talking to people to see if a concept resonates, or testing whether there is a real demand for what you have in mind before committing to it.

(b) Improving your product or service based on what customers tell you — for example, changing a feature, adjusting how you deliver the service, or rethinking part of your offering after hearing that something isn't working.

(c) Figuring out the best way to reach and attract customers — for example, trying different messages, platforms, or promotional offers to see which ones bring in the most interest or sales.

(d) Keeping track of how your business is doing and using that information to make improvements — for example, monitoring customer complaints, looking at which products or services sell best, or tracking whether customers come back.

(e) Building an early version of your product or service to test whether it works before fully committing to it — for example, putting together a basic version, trying it out, and making changes based on what you learn.

(f) Figuring out the best way to sell your product or service — for example, trying different ways of explaining what you offer, adjusting your pitch based on how potential customers react, or testing different prices to see what people are willing to pay.

Q6.2 Now we'd like to know whether you have ever *used an AI* tool to complete each of the following tasks.

[Yes/No for each of the six tasks listed in Q6.1]

[Block 7: Overall AI reliance and counterfactual importance]

Finally, we have a few general questions about your overall use of AI tools and whether their availability has shaped your decision to pursue a business idea.

Q7.1 When you first began working on or exploring a business idea — even if it never became an active venture — how much did you rely on AI tools to help with any part of the process?

○ Not at all ○ A little ○ Somewhat ○ Substantially ○ A lot

Q7.2 In your own words, please describe how you use AI tools in connection with your business or business idea.

[Open-ended]

Q7.3 If tools like ChatGPT and other AI assistants had not been available, how likely is it that you would have continued exploring or pursued your business idea?

- Very unlikely — I probably wouldn't have explored the idea at all
- Somewhat unlikely — I might have explored it less seriously or stopped sooner
- Somewhat likely — I would still have explored it, but with more difficulty
- Very likely — I would have pursued it about the same way

Q7.4 Before finishing, we'd love to hear any additional thoughts you have about using AI tools as a potential or current entrepreneur. Is there anything you think we haven't covered in the survey?

[Open-ended]

[Block 8: Final feedback]

Q8.1 Thank you so much for your participation. Do you have any final comments or feedback about this survey or your experience?

[Open-ended]

E.2. Constructing the Analysis Sample. Our target population for this survey is U.S.-based new and aspiring entrepreneurs — individuals who either started a business or began actively pursuing a business idea on or after January 2023. We are aware of no existing respondent panel that allows us to target this population directly. The closest available option is to leverage a commercial panel of survey respondents and apply screening criteria at the point of survey entry. We use Dynata for this purpose, with respondents who do not meet our targeting criteria screened out before they begin the substantive portion of the survey.

To ensure that respondents engage substantively with our questions, we apply a set of quality controls to the responses we receive. Table [OA27](#) summarizes the cleaning steps and reports the number of respondents retained at each stage. Starting from the 1,797 eligible respondents who completed the survey, we drop those whose self-reported business start dates fall outside the eligibility window, those who exhibit uniform response patterns across the rating batteries, those whose open-ended responses do not substantively engage with the survey questions, and those whose total survey duration is implausibly short. These steps leave a final analytic sample of 1,036 respondents.

Table OA27. Construction of the Analytic Sample

Stage	N
Eligible respondents (after Qualtrics screener)	1,797
After deduplication of panel identifiers	1,797
After business start-date validation	1,356
After straightlining screen	1,167
After open-ended response relevance screen	1,037
After survey duration screen	1,036
Final analytic sample	1,036

Notes: ...

Beyond filtering respondents, we also process two further pieces of data in ways that affect the measures presented in the main text. First, the AI Assistance Scores generated by our LLM pipeline are produced at the level of granular entrepreneurial tasks, of which the pipeline identifies several hundred unique instances. Because the survey presents respondents with broader task categories, comparing the two measures requires mapping each granular task to one of the twenty survey categories. We construct this mapping by combining automated classification with manual review. The categorization is performed at the task-name level — i.e., the same task name maps to the same category in every industry where it appears — and is based solely on the task name itself; industry context is deliberately excluded from the categorization decision to avoid introducing industry-specific bias.


Second, the survey includes open-ended questions in which respondents describe in their own words how they think about AI tools and how they actually use them. To make these responses tractable for descriptive analysis, we classify each response using a large language model along two dimensions: relevance and content. The relevance classification feeds into the response-quality screen described above and contributes to the construction of the analytic sample. The content classification assigns one or more codes to each response from a pre-specified controlled vocabulary that we developed to capture the substantive types of AI use and AI-related views typically described by entrepreneurs. The vocabulary includes a residual “other” category that allowed the model to supply a custom label for responses that did not fit any predefined code. Because a single response may describe multiple distinct uses or views, individual responses may carry multiple codes.

E.3. Sample Composition. This subsection reports descriptive statistics on the final analytic sample of 1,036 respondents. Figures OA21, OA22, and OA23 show the distribution of respondents across age groups, educational attainment, and the current stage of their business or business idea. The sample skews toward older respondents, with respondents in the 41-and-above age range accounting for roughly 56% of the sample. Educational attainment is concentrated around bachelor’s degrees and below, with about 38% of respondents reporting at most some college and another 41% holding an associate’s or bachelor’s degree. Just under 17% of respondents hold a


graduate or professional degree. The sample is roughly evenly split between respondents whose business is in the planning or pre-launch stage (45%) and those whose business is active and operating (37%); the remaining respondents either abandoned the idea before launching (15%) or closed an active business (3%).

Figure OA24 reports the distribution of respondents across NAICS-3 industries. The sample covers 75 distinct industries. The five most represented industries — Food Manufacturing (NAICS 311), Construction of Buildings (236), Clothing and Accessories Retailers (458), Food and Beverage Retailers (445), and Food Services and Drinking Places (722) — together account for roughly 27% of the sample, with the remaining respondents distributed across a long tail of consumer-facing and operationally diverse industries.


Figure OA21. Age Distribution of Survey Respondents




Notes: This figure shows the distribution of survey respondents across age groups. The survey was fielded between November 2025 and May 2026 through Dynata’s respondent panel. Qualified respondents are new or aspiring U.S. entrepreneurs who started a company or began pursuing a business idea on or after January 2023. The final analytic sample after quality filtering consists of 1,036 respondents. Respondents were asked to select their age group from the following options: 21–25, 26–30, 31–35, 36–40, 41–45, 46–55, and 56+. Percentages are computed over the full sample.

Figure OA22. Educational Attainment of Survey Respondents

Notes: This figure shows the distribution of survey respondents across the highest level of education completed. The survey was fielded between November 2025 and May 2026 through Dynata’s respondent panel. Qualified respondents are new or aspiring U.S. entrepreneurs who started a company or began pursuing a business idea on or after January 2023. The final analytic sample after quality filtering consists of 1,036 respondents. Respondents were asked “What is the highest level of education you have completed?” and selected from nine options ranging from less than high school to doctoral degree. Categories are displayed in order of attainment. Percentages are computed over the full sample.

Figure OA23. Business Stage of Survey Respondents

Notes: This figure shows the distribution of survey respondents by the current stage of their business or business idea. The survey was fielded between November 2025 and May 2026 through Dynata’s respondent panel. Qualified respondents are new or aspiring U.S. entrepreneurs who started a company or began pursuing a business idea on or after January 2023. The final analytic sample after quality filtering consists of 1,036 respondents. Respondents were asked “Is your business currently:” and selected from four options: in planning/pre-launch stage, active and operating, never launched and idea was abandoned, or closed. Percentages are computed over the full sample.

Figure OA24. Industry Distribution of Survey Respondents

Notes: This figure shows the distribution of survey respondents across NAICS-3 industries. The survey was fielded between November 2025 and May 2026 through Dynata’s respondent panel. Qualified respondents are new or aspiring U.S. entrepreneurs who started a company or began pursuing a business idea on or after January 2023. The final analytic sample after quality filtering consists of 1,036 respondents. Respondents were asked “In which industry does (or will) your business primarily operate?” and selected from a dropdown menu of three-digit NAICS codes. The sample covers 75 distinct three-digit NAICS industries. The figure shows the 20 most represented industries by three-digit NAICS code, ordered from most to least frequent; the remaining 55 industries are collapsed into a single “Other” bar at the right. The corresponding industry names can be found in the U.S. Census Bureau’s NAICS classification.

APPENDIX F. IN DEPTH LITERATURE REVIEW

In this section of the appendix we provide more granular information on the various strands of literature the paper is connected to.

F.1. A Closer Look at the Impact of AI. This paper contributes to the growing body of work on AI, automation, and labor markets. Recent evidence shows that AI—particularly generative AI (genAI) tools like ChatGPT—has been adopted at a remarkably rapid pace by individuals. According to a survey presented in [Bick et al. \(2024\)](#), by late 2024, nearly 40% of U.S. adults aged 18–64 had used genAI either for work or outside of work. Work-related adoption stood at 26%, with 9% of employed individuals using genAI every workday. However, the same survey shows that adoption patterns exhibit strong heterogeneity by demographic and occupational groups. Adoption is significantly higher among younger, more educated, and higher-income individuals. For example, workers with a bachelor’s degree or higher adopt genAI at more than twice the rate of those without college degrees. Adoption also varies by occupation: computer and mathematical occupations, management, and business/finance show the highest adoption rates—above 40%—whereas adoption is lower in blue-collar and personal service jobs. Similar patterns emerge internationally. [Humlum and Vestergaard \(2025\)](#) find that among workers in Denmark employed in occupations considered highly exposed to genAI, awareness is widespread but actual adoption varies. As in the U.S., younger and more educated individuals are more likely to adopt, although overall usage remains modest even in occupations with high exposure to AI tools.

While adoption is broad, actual usage intensity is considerably lower. Among those who use genAI at work, [Bick et al. \(2024\)](#) estimate that only 1–5% of total U.S. work hours are assisted by genAI. Even among users, the modal usage is less than one hour per day, and most users do not engage with genAI every workday. This underscores a significant gap between ever using AI and frequently or intensively using it. Usage also varies across demographic and occupational dimensions. Daily genAI users are more concentrated in computer science, business, and management roles. These groups not only report higher frequency of use but also more time spent with genAI per day. For example, in information services, workers spend over 10% of their work hours using genAI, compared to just 1–2% in leisure or blue-collar sectors. In terms of tasks, genAI is most commonly used for writing, summarizing, administrative tasks, and idea generation, with 40% of users ranking writing as the top use case. This aligns with occupational heterogeneity: occupations requiring frequent written communication or analysis tend to see more intense AI usage. Consistent with this, [Humlum and Vestergaard \(2025\)](#) note that usage in high-exposure occupations is mostly informal or exploratory, such as drafting or summarizing, rather than deeply embedded into daily workflows.

[Bonney et al. \(2024\)](#) uncover that firm-level adoption has lagged behind individual adoption—though it is increasing—, and there is substantial heterogeneity across firm types. Leveraging the Census Bureau’s Business Trends and Outlook Survey (BTOS), the authors find that only 3.7% of U.S. firms were using AI in September 2023. By early 2024, this figure rose to 5.4%, with forecasts suggesting it may reach close to 10% by the end of 2024. Adoption is more common in large firms

and specific sectors. Employment-weighted statistics reveal that while firm-level adoption is low, these AI-using firms represent a larger share of the workforce. For instance, in the information sector, 18% of firms had adopted AI, covering 22% of the sector’s workers. This suggests that AI is being adopted first by larger and more technologically capable firms. There is also substantial heterogeneity by industry: adoption is highest in information, finance, and professional services, and lowest in construction, retail, and accommodation sectors. These patterns align with the higher technical readiness and digital infrastructure in the more advanced sectors.

Overall, adoption of AI technology is higher among young, highly educated, and high-wage workers, consistent with early PC adoption trends. However, compared to other general purpose technologies, AI adoption has been as fast as either PCs or the internet (Bick et al., 2024).

Two influential papers provide experimental evidence on how generative AI affects worker outcomes. Brynjolfsson et al. (2025b) conduct a randomized controlled trial in which customer service agents were given access to generative AI tools. They find that AI significantly boosts worker productivity: AI exposure leads to a 14% increase in issues resolved per hour and improves customer sentiment. The largest gains accrue to the least experienced (and previously lowest-performing) workers, who benefit from AI’s ability to share tacit knowledge encoded in model outputs. Importantly, they do not find evidence that AI reduced hours worked or displaced workers during the trial period. Noy and Zhang (2023) also find substantial productivity gains in a controlled setting where college-educated professionals completed writing tasks with and without GPT-3.5 assistance. Access to the LLM reduced task completion time by over 30% and increased the quality of output, as rated by independent evaluators. As in Brynjolfsson et al. (2025b), the productivity gains were largest for lower-performing workers. The LLM reduced performance inequality while raising average output. However, the paper does not directly assess long-run wage effects or changes in labor hours. Together, these studies suggest that AI can enhance productivity—especially for less skilled or lower-performing individuals—without negative short-run impacts on employment or wages. AI can reshape human performance by enhancing creativity and problem-solving, not only boosting the quality and viability of ideas but also reducing barriers for less inherently creative individuals—though often at the cost of collective diversity and novelty in outcomes. AI-human generated product ideas are on average higher in quality and far more likely to produce top-tier ideas than human-generated ones (Doshi and Hauser, 2024; Meincke et al., 2024). Moreover, while human crowd solutions are more novel, human–AI generated solutions are consistently rated higher in strategic viability, environmental and financial value, and overall quality—especially when guided by iterative human prompting (Boussioux et al., 2024). These findings underscore AI’s ability to help with innovation and experimentation, as well as the need of pairing the technology with human input for maximum effectiveness.

The heterogeneity in AI effects stems from its varying effectiveness across tasks, as emphasized by Eloundou et al. (2023). They propose the concept of a “jagged technological frontier”: AI models like GPTs are highly capable in some tasks—particularly language generation, summarization, or classification—but perform poorly in others, such as those requiring numerical reasoning, physical manipulation, or complex contextual understanding.

Since occupational roles comprise multiple tasks with differing AI susceptibility, the overall effect of AI depends on the share of tasks that fall “inside” the AI’s capability frontier. [Hui et al. \(2024\)](#) show that AI adoption and productivity gains vary across occupations, with greater exposure and uptake in tech-intensive roles (e.g., software development, marketing), and lower adoption among clerical and administrative roles. [Hampole et al. \(2025\)](#) find that AI exposure has heterogeneous impacts on employment and wages: some professional jobs experience productivity boosts and wage growth, while others face displacement risks, depending on task complementarity.

[Dell’Acqua et al. \(2023\)](#) show experimentally that low-skill consultants benefit disproportionately from GPT-4 assistance on business problem-solving tasks. In tasks within the model’s capability frontier, lower-performing individuals improve performance by 43%, compared to 17% for higher performers. However, when AI is used for tasks outside its capability frontier, performance degrades, highlighting the importance of human judgment.

There appears to be heterogeneity also by country, as the results found by previous paper all focus on the U.S. [Humlum and Vestergaard \(2025\)](#) use linked survey-register data from Denmark to show that ChatGPT adoption is widespread but unequal and that it does not lead to market productivity gains.

The literature paints a nuanced picture: while AI offers sizable productivity benefits in the U.S.—especially for routine tasks and less experienced workers—its effects are highly heterogeneous. These differences arise both from variability in task suitability for AI and from uneven adoption across worker demographics, occupations, and geographies.

Recent empirical research has begun to uncover how advances in artificial intelligence (AI), especially generative AI, affect firms across several dimensions including productivity, fundraising, employment, profitability, and valuation.

[Heller and Asam \(2024\)](#) provide early evidence on how AI adoption influences productivity and capital raising among firms. They document that firms which publicly signal AI-related capabilities experience improvements in productivity metrics and are more successful in attracting financing, especially venture capital. These effects are more pronounced among younger firms and those in tech-intensive industries, suggesting that AI capabilities serve both as operational tools and as credibility signals to investors.

[Babina et al. \(2024\)](#) examine the firm-level consequences of AI adoption using granular data on job postings, occupational composition, and financial outcomes. They find that firms more exposed to generative AI (measured via occupational task composition) experienced significant labor substitution effects after the release of ChatGPT. Specifically, firms with high exposure decreased hiring in occupations whose core tasks were susceptible to automation by generative AI, while also exhibiting increased profitability. This supports a cost-saving mechanism through labor substitution. These findings resonate with evidence from [Hampole et al. \(2025\)](#), who show that technological change—especially AI-driven—alters firm workforce composition by displacing workers in exposed occupations. Both studies emphasize the heterogeneity of impact depending on whether AI affects core or supplemental occupational tasks. When core tasks are automatable, substitution is more likely, leading to reductions in labor demand within affected firms.

Babina et al. (2024) also link AI exposure to firm valuations. Firms with greater exposure to generative AI—particularly through core occupational tasks—experienced statistically significant increases in market value following ChatGPT’s release. This valuation effect is attributed to anticipated cost reductions and improved margins, rather than product-side AI applications. Eisfeldt et al. (2025) corroborate and deepen this result by constructing an “Artificial Minus Human” (AMH) portfolio, which is long in firms with high AI-exposed labor and short in those with low exposure. They show that this AMH portfolio earned abnormal returns of 0.44% per day in the two weeks following ChatGPT’s release, driven by labor substitution expectations. Firms with greater AI exposure and stronger data infrastructures reaped higher returns, underscoring the complementarity between AI and firm-specific data assets. Moreover, they find that valuation gains are concentrated in firms whose AI exposure stems from core occupational tasks, reinforcing the role of substitutability in driving shareholder value.

Studies on the effects of AI on startups and new firms are still rather limited. Equipping startup founders with virtual AI agents supports, on average, does not improve business performance (Otis et al., 2024). However, the same study finds that there is large heterogeneity in the effects of AI access on performance: high performers at baseline increased revenues and profits by about 15%, while low performers saw declines of about 8%. This divergence arose not because of different questions asked or advice received, but because high performers selected and implemented tailored suggestions effectively, whereas low performers acted on more generic advice such as lowering prices or spending on advertising, which often hurt outcomes

Closest to our work is Bao et al. (2025), who study whether generative AI access, sparked by the release of ChatGPT, affects entrepreneurial entry in STEM sectors. STEM workers in industries with higher GenAI exposure are more likely to start incorporated businesses after ChatGPT’s release. The increase is driven by automation of peripheral tasks, supporting an augmentation channel rather than displacement through an automation channel. Effects are strongest among lower-ranked workers and those in soft-skill-intensive jobs, who benefit most from reduced entry barriers. Overall, democratization of GenAI access boosts opportunity-driven entrepreneurship.

In a concurrent study, Bena et al. (2026) document that industries with higher occupational exposure to GenAI experienced significantly more startup formation following the release of ChatGPT and Stable Diffusion. Using a measure of task-level AI exposure derived from actual user interactions with Claude.ai, they show that this entry is geographically dispersed beyond traditional innovation hubs, concentrated among smaller and leaner firms, and accompanied by a reallocation of startup employment toward GenAI-exposed occupations. At the founder level, new entrepreneurs in exposed industries are increasingly drawn from occupations whose tasks are complementary to GenAI.

Relatedly, Gupta et al. (2025) study how generative AI reshapes the U.S. startup ecosystem along both the intensive and extensive margins. Measuring firm-level exposure through startups’ pre-ChatGPT occupational composition, they show that more exposed startups cut employment by roughly 8% within two quarters of ChatGPT’s release, with reductions concentrated among junior employees and implementation-intensive roles; displaced workers faced longer non-employment

spells and moved to lower-paying, less exposed jobs. At the same time, exposed startups became more productive, scaled faster, and accelerated through financing rounds. On the investor side, venture capitalists with greater exposure shifted toward smaller but more frequent initial investments—consistent with a lower cost of experimentation—and markets more exposed to generative AI experienced a 7% increase in the total number of startups. The employment created by these new entrants offset the losses at incumbent startups, leaving aggregate startup employment essentially unchanged: generative AI reallocated labor within the entrepreneurial ecosystem rather than shrinking it.

Taken together, these studies illustrate that AI’s firm-level effects extend beyond technological novelty—restructuring workforce composition, shifting cost structures, and ultimately influencing firm value through both expectations and realized operational improvements.

All in all, the emerging literature documents widespread adoption of generative AI tools by individuals, with evidence consistently pointing to strong productivity gains, particularly for less experienced or lower-performing workers. These benefits, however, are highly heterogeneous—varying across demographic groups, occupations, and tasks—due to differences in AI capabilities and complementarities. At the firm level, AI adoption is still in its early stages but growing, especially among large and technologically advanced firms. Importantly, while AI adoption does not appear to reduce aggregate employment, it may induce within-firm reallocation of tasks and changes in occupational demand. Thus, the consensus so far highlights AI’s potential to boost productivity without broad aggregate labor displacement, but cautions that its effects are uneven and context-dependent. Our work contributes to this growing body of literature by investigating the effects of AI on entrepreneurs, a category of workers that has received relatively less attention so far.

F.2. A Closer Look at Entrepreneurship as a Process of Experimentation. Entrepreneurship is increasingly understood as a process of iterative learning, where founders test hypotheses about products, customers, or business models and adjust strategies based on feedback. Rather than executing a fixed plan, ventures cycle through trial and error, reducing uncertainty over time. The *Lean Startup* approach formalizes this process: entrepreneurs build minimum viable products (MVPs), measure responses, and learn whether to pivot or persevere (Ries, 2011). Blank (2020) similarly emphasizes customer discovery and hypothesis testing as structured forms of learning. Together, these frameworks anchor the modern view of entrepreneurship as hypothesis-driven experimentation.

These practical insights build on long-standing theory. Knight (1921) distinguished risk from uncertainty, positioning entrepreneurship as action in contexts where probabilities are unknowable. Jovanovic (1982) modeled firm entry as market experimentation: entrants learn their efficiency, with inefficient firms exiting and efficient ones growing. Ericson and Pakes (1995) extended this framework, highlighting dynamic experimentation as central to firm and industry evolution. These foundational models establish learning and selection as the mechanisms through which entrepreneurship drives reallocation.

Recent empirical work confirms that iterative testing yields measurable outcomes. [Koning et al. \(2022\)](#) show that firms adopting A/B testing achieve significant gains in web traffic and product launches, indicating that systematic experimentation reshapes performance trajectories and increases variance in outcomes. This demonstrates that experimentation is not limited to managerial rhetoric but has real, quantifiable effects on firm growth and survival.

Experimentation is not universally beneficial—it interacts with organizational structure. [Contigiani and Young-Hyman \(2022\)](#) show that investors evaluate startups more positively when strategy and structure are coherent: experimentation combined with informal structure, or planning combined with formal structure. Ventures that mismatch strategy and structure face skepticism, underscoring that experimentation is perceived as legitimate only when supported by organizational flexibility.

Experimentation often involves early disclosure, exposing firms to imitation risk. [Contigiani \(2023\)](#) demonstrates this in U.S. software startups: following the *Alice Corp. v. CLS Bank* decision, which weakened patent protection, startups reduced their experimentation, particularly in competitive environments. This highlights a critical trade-off: experimentation accelerates learning but can undermine value capture if intellectual property is insufficiently protected.

At the individual level, entrepreneurship resembles an option to experiment. [Manso \(2016\)](#) shows that entrepreneurial spells are typically short, unsuccessful founders can return to wage employment without penalty, and long-term returns are attractive once the option value of trying and exiting is recognized. This option-like structure explains why many people enter despite high early-stage failure rates.

Because early-stage experimentation generates noisy signals, financing structures shape which tests are feasible. Venture capital addresses this problem by staging investments and conditioning funding on milestones, thereby containing downside risk while enabling upside exploration ([Kerr et al., 2014](#)). Empirical evidence shows VC accelerates professionalization and supports pivots ([Hellmann and Puri, 2000, 2002](#)), while financing terms influence the scale and type of experiments pursued ([Ewens et al., 2018](#)).

Experimentation also has systemic consequences. Failed ventures contribute knowledge spillovers, guiding future entrants and reallocating resources toward validated business models. Optimal incentive schemes tolerate early failure and reward long-term success ([Manso, 2011](#)), and policy complements—such as bankruptcy protection, unemployment insurance, or government R&D funding—expand the set of feasible experiments ([Kerr et al., 2014](#)).

In sum, positioning entrepreneurship as iterative experimentation reconciles managerial frameworks with formal theory and empirical evidence. In practice, frameworks like *Lean Startup* ([Ries, 2011](#); [Blank, 2020](#)) demonstrate how ventures use hypothesis testing to reduce uncertainty. Foundational models ([Knight, 1921](#); [Jovanovic, 1982](#); [Ericson and Pakes, 1995](#)) explain why learning and exit are central to entrepreneurial dynamics. Empirical studies show that digital experimentation boosts performance ([Koning et al., 2022](#)), organizational fit shapes legitimacy ([Contigiani and Young-Hyman, 2022](#)), intellectual property affects disclosure ([Contigiani, 2023](#)), and individual returns follow an option-like profile ([Manso, 2016](#)). Financing structures and policy interventions

further determine which experiments are undertaken and scaled. Together, this literature portrays entrepreneurship as a system of structured experiments that generate both private learning and public spillovers.

APPENDIX G. THEORETICAL FRAMEWORK: AI AND ENTREPRENEURIAL ENTRY

G.1. Overview. In this appendix, we develop a stylized theoretical framework to interpret our primary empirical findings. While the model is not intended for quantitative calibration, it serves to illuminate the key economic mechanisms driving the relationship between entrepreneurship and the arrival of AI.

The main results of our paper suggest that, upon the arrival of ChatGPT, the overall entry rate goes up and the average quality of new ventures goes up as well. We establish these results in a simple two-period model in which individuals choose between salaried production work and becoming an entrepreneur. Individuals are heterogeneous along two independent dimensions. First, they differ in entrepreneurial capability, which scales the expected payoff of a new venture. Second, they are subject to an idiosyncratic entry friction—capturing factors such as credit constraints, access to financing, and the availability of experimentation—which is assumed to be orthogonal to innate ability. We begin by establishing the model’s fundamental environment and then derive a series of key propositions that rationalize our primary empirical findings.

G.2. Model Setup. Time is discrete. There are two periods, indexed by $t = 0, 1$. AI is available only in $t = 1$. In each period, the representative household consumes the final good. The price of the final good is normalized to 1. The final good is a composite of intermediate goods, aggregated by a Dixit-Stiglitz aggregator:

$$(1) \quad Y_t = \left(\int_0^{M_t} k_{jt}^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}}$$

where M_t is the mass of intermediate goods firms in the economy and $\sigma > 1$ is the elasticity of substitution. The cost minimization problem of the final goods producer is given by:

$$(2) \quad \min_{\{k_{jt}\}} \left(\int_0^{M_t} k_j(t)^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}} - \int_0^{M_t} p_{jt} k_{jt} dj$$

which gives rise to the demand function:

$$(3) \quad p_{jt}^\sigma k_{jt} = Y_t$$

Each intermediate goods producer operates with a linear technology $k_{jt} = l_{jt}$ under monopolistic competition, subject to the demand function derived above. The profit maximization problem of a representative intermediate goods firm can be written as:

$$(4) \quad \max_{\{p_{jt}, k_{jt}\}} p_{jt} k_{jt} - W_t l_{jt} \quad \text{s.t.} \quad k_{jt} = l_{jt}, \quad p_{jt}^\sigma k_{jt} = Y_t$$

The resulting flow profits from operating an intermediate goods firm are:

$$(5) \quad \pi_t = \left(\frac{1}{\sigma-1} \right) \left(\frac{\sigma-1}{\sigma} \right)^\sigma \left(\frac{1}{W_t} \right)^{\sigma-1} Y_t$$

Note that the intermediate goods firms are symmetric, and we suppress the j subscript for notational simplicity.

G.3. Heterogeneous Entrepreneurial Capabilities. The intermediate goods firms are started by individuals. There is a continuum of individuals: a fraction $h \in (0, 1)$ are high-capability individuals, whereas the remaining $1 - h$ are low-capability individuals. To start a new firm, an entrepreneur must pay a fixed cost denominated in units of production labor. High-capability individuals pay a fixed cost of $C_{H,t}$ units of labor, while low-capability individuals pay $C_{L,t}$ units, with $C_{H,t} < C_{L,t}$. This assumption captures the idea that high-capability individuals are able to start the same firm with fewer workers and a smaller team.

We assume that each newly created firm is a symmetric intermediate goods firm and that each firm lasts for only one period. The entrepreneurial intensity, or entry rate, in the economy is captured by M_t , the active mass of intermediate goods firms. While the model can be extended to allow high-capability individuals to start firms with higher productivity, doing so does not yield additional insight beyond our baseline setup. In the first-best equilibrium, low-capability individuals would enter the entrepreneur pool only after all high-capability individuals. However, as we show below, the presence of idiosyncratic entry frictions breaks this ranking.

Upon the arrival of AI, we assume that only the fixed cost of high-capability individuals is affected:

$$(6) \quad C_{H,1} < C_{H,0} < C_{L,0} = C_{L,1}$$

This simplified assumption captures the main intuition underlying our empirical results: the arrival of AI facilitates the creation of new ventures, and individuals with higher capability benefit more from AI. In our stylized model, low-capability individuals are not directly impacted by AI, and the benefits accrue only to the high-capability individuals.

G.4. Entrepreneurship Decision. Each individual has an inelastic labor supply of 1 unit in each period. At the beginning of each period, individuals choose between working as a salaried employee at wage W_t and starting their own business by paying the fixed cost.²⁸ The known benefit of becoming an entrepreneur is:

$$(7) \quad \pi_t - C_{K,t}W_t$$

where $C_{K,t}$ denotes the type-specific fixed cost of labor ($K \in \{H, L\}$). The utility of individual i from choosing entrepreneurship is:

$$(8) \quad \log V_{E,i,t} = \log(\pi_t - C_{K(i),t}W_t) - F_{E,i}$$

and the utility from remaining in salaried employment is:

$$(9) \quad \log V_{W,i,t} = \log W_t - F_{W,i}$$

where $F_{E,i}$ and $F_{W,i}$ capture idiosyncratic components of individual i 's occupational choice. These idiosyncratic components are orthogonal to entrepreneurial capability and are not time-varying.

²⁸This is consistent with the assumption of finite firm horizons, where production and exit occur within a single period.

We assume that $F_{E,i} - F_{W,i}$ follows a logistic distribution. The interpretation of this idiosyncratic friction encompasses heterogeneous access to early-stage financing, the degree of uncertainty surrounding the entrepreneurial outcome, and the availability of testing strategies, among other factors. It is natural to extend this component to include the cost of experimentation or to incorporate key elements from Bayesian entrepreneurship (e.g., signals drawn to learn the underlying quality of a new venture).²⁹

Applying standard results from discrete choice theory, the mass of entrepreneurs of each type is:

$$(10) \quad M_t = M_{t,H} + M_{t,L} = h \cdot \frac{\pi_t - C_{H,t}W_t}{\pi_t - (C_{H,t} - 1)W_t} + (1 - h) \cdot \frac{\pi_t - C_{L,t}W_t}{\pi_t - (C_{L,t} - 1)W_t}$$

G.5. Equilibrium. The equilibrium is defined by a combination of values $\{M_t, W_t, Y_t\}$ such that the following conditions hold simultaneously.

First, individuals choose between entrepreneurship and salaried employment based on the deterministic value differential and the idiosyncratic utility component:

$$(11) \quad M_t = h \cdot \frac{\pi_t - C_{H,t}W_t}{\pi_t - (C_{H,t} - 1)W_t} + (1 - h) \cdot \frac{\pi_t - C_{L,t}W_t}{\pi_t - (C_{L,t} - 1)W_t}$$

Second, the labor input l_{jt} , final output k_{jt} , and profit π_t of intermediate goods firms satisfy the optimality conditions derived above and can be represented as functions of W_t and Y_t .

Third, the labor market clearing condition holds:

$$(12) \quad M_t \cdot l_{jt} + M_{t,H} \cdot C_{H,t} + M_{t,L} \cdot C_{L,t} + M_t = 1$$

where the left-hand side accounts for production labor employed by firms ($M_t \cdot l_{jt}$), labor used as fixed costs by high- and low-capability entrepreneurs ($M_{t,H} \cdot C_{H,t}$ and $M_{t,L} \cdot C_{L,t}$), and the mass of entrepreneurs themselves (M_t), which equals the total labor endowment of 1.

Fourth, the income of the representative household equals:

$$(13) \quad Y_t = W_t \cdot (M_t \cdot l_{jt} + M_{t,H} \cdot C_{H,t} + M_{t,L} \cdot C_{L,t}) + M_{t,H} \cdot (\pi_t - C_{H,t}W_t) + M_{t,L} \cdot (\pi_t - C_{L,t}W_t)$$

which sums the wage income of all salaried workers and the profits of all entrepreneurs.

G.6. Main Results. Given the model's setup, we now formalize the main propositions. These predictions are qualitatively consistent with the empirical patterns documented in the main text.

PROPOSITION 1 (Higher Entrepreneur Quality). *In an economy with AI, the average quality of entrepreneurs is higher:*

$$\frac{M_{1,H}}{M_1} > \frac{M_{0,H}}{M_0}$$

The ratio of high-capability individuals in the entrepreneur pool is determined by the relative probability of high-capability individuals entering entrepreneurship versus that of low-capability individuals. When the net profit from entrepreneurship, $\pi_t - C_{K,t}W_t$, grows faster for high-capability individuals, the share of high-capability individuals in the total entrepreneur pool increases more rapidly than it does for the other type. In our setting, as the total mass of entrepreneurs rises,

²⁹In this simplified version, we retain the unspecified and general idiosyncratic component and show that this suffices to generate our key results.

the labor market features higher demand and lower supply, and the wage rate W_t must increase to clear the market. This rise in wages reduces the net profits of low-capability individuals—who did not benefit from the AI-induced cost reduction—and drives some of them out of the entrepreneur pool. Consequently, the average quality of entrepreneurs rises.

PROPOSITION 2 (Higher Entry Rate). *In an economy with AI, the total mass of entrepreneurs increases:*

$$M_1 > M_0$$

It is intuitive that the economy should see more entrepreneurs when the fixed cost of entry falls. However, the rise in wages ($W_1 > W_0$) resulting from lower labor supply and higher labor demand may dampen the incentive to enter. To formally establish the result, it suffices to show that the mass of entrepreneurs cannot decline. Suppose, toward a contradiction, that $M_1 < M_0$. Then labor demand falls because there are fewer firms, and the overall fixed cost of entry declines because of higher efficiency and a greater share of high-capability entrepreneurs, giving rise to higher labor supply. Such an equilibrium exists only if $W_1 < W_0$ and $Y_1 > Y_0$, since fewer firms in the economy with more productive labor implies that each firm produces more. However, with $W_1 < W_0$ and $Y_1 > Y_0$, it follows from Equation (5) that $\pi_1 > \pi_0$, meaning that the incentive for individuals to become entrepreneurs is strictly larger than in the pre-AI equilibrium. This contradicts $M_1 < M_0$. Therefore, $M_1 > M_0$ whenever $C_{H,1} < C_{H,0}$: the general equilibrium impact of increasing wages does not reverse the growing entrepreneurship incentive of high-capability individuals.

G.7. Discussion. The key intuition delivered by the model is that, for high-capability individuals, the arrival of AI increases the gap between the entrepreneurial payoff and the salaried wage by more than it does for low-capability individuals. As a result, a disproportionate share of high-capability individuals enter the entrepreneur pool, and average quality rises.

The existence of the idiosyncratic preference shock, or frictional component, plays a central role. If no friction exists, high-capability individuals always enter the entrepreneur pool before low-capability individuals—that is, if any low-capability individual finds it optimal to become an entrepreneur, all high-capability individuals will certainly choose entrepreneurship as well. In such a frictionless world, any cost-reduction shock would only bring in lower-quality entrepreneurs, since the higher-quality ones are always the first to start their own businesses. The friction breaks this strict ordering: it allows some high-capability individuals to remain in salaried employment despite having the ability to run successful ventures, while some low-capability individuals enter entrepreneurship due to favorable friction draws. When AI arrives and asymmetrically reduces costs for high-capability individuals, it unlocks entry for those high-capability individuals whose friction draws previously kept them out of the entrepreneur pool, driving up both entry and average quality.

This appendix provides a parsimonious framework designed to capture our primary empirical findings with minimal theoretical structure. While the current setup focuses on the core mechanisms, it can be naturally extended in future research to explicitly incorporate the cost of experimentation discussed in our empirical analysis, or to embed elements of Bayesian learning in which individuals draw signals to learn the underlying quality of their ventures prior to entry.