

Generative AI and Entrepreneurship*

Abhinav Gupta[†], Franklin Qian[‡], Elena Simintzi[§] & Yifan Sun[¶]

February 24, 2026

This paper studies how Generative AI (Gen AI) is reshaping the U.S. startup ecosystem. Exploiting the release of ChatGPT, we show that startups with greater pre-release Gen AI task exposure reduced employment within two quarters, primarily among junior and implementation roles. Displaced workers experienced longer unemployment spells and moved to lower-paying but less exposed jobs. Conversely, exposed startups increased productivity, scaled faster, and accelerated through financing rounds. Venture capital shifted toward frequent, smaller investments, boosting new firm formation. Overall, incumbent contraction was offset by new firm formation, leaving aggregate employment unchanged but shifting composition to senior roles.

Keywords: Generative Artificial Intelligence, Entrepreneurship, Labor Markets, Venture Capital, Firm Formation

JEL Codes: L26, J23, J24, G24, M13

*We are grateful to seminar participants at Purdue University. All errors are our own.

[†]UNC Kenan-Flagler Business School, Email: abhinav.gupta@kenan-flagler.unc.edu

[‡]UNC Kenan-Flagler Business School, Email: franklin.qian@kenan-flagler.unc.edu

[§]UNC Kenan-Flagler Business School, CEPR, & ECGI, Email: elena.simintzi@kenan-flagler.unc.edu

[¶]University of Pennsylvania, Email: sunyifan@sas.upenn.edu

1 Introduction

Startups have long been recognized as engines of economic growth. Traditionally, the process of business dynamism, namely new firm formation, has been accompanied by strong job growth accounting for a large share of job creation in the U.S. economy (Haltiwanger et al., 2013; Decker et al., 2014; Adelino et al., 2017). At the firm level, the entrepreneurship literature has long documented the relationship, where successful startups are those that grow their employment and raise capital. However, the advent of Generative AI (Gen AI) has raised questions about whether this traditional link between startup growth and employment will persist. For example, Silicon Valley players have speculated about the emergence of “one-person unicorns”,¹ as Gen AI expands both the depth and breadth of what a single entrepreneur can accomplish. This paper provides new evidence on whether, and how, Gen AI is changing entrepreneurship.

Whether Gen AI disrupts the relationship between growth and employment for young firms is a priori ambiguous. This relation critically depends on two factors: adoption frictions and Gen AI’s impact on startup labor demand. Historically, smaller firms have lagged in technology adoption (Acemoglu et al., 2024; Bessen, 2020; Zolas et al., 2020), primarily due to limited access to capital and smaller economies of scale. However, Gen AI differs from earlier technologies in several respects. It is accessible through pay-per-use pricing, relies on pre-trained foundation models, and benefits from extensive open-source frameworks. These features may lower traditional scale barriers (Deshpande et al., 2023), enabling startups to adopt Gen AI with limited upfront investment and without proprietary data.

The impact of Gen AI on startups’ labor demand is also unclear. Startups often face cash constraints (Kerr et al., 2014; Chen and Ewens, 2025), and they may adopt Gen AI to substitute for labor and lower costs. At the same time, startups face persistent talent constraints (Beckman and Burton, 2008; Ouimet and Zarutskie, 2014; Sorenson et al.,

¹Sam Altman remarked, “*In my little group chat with my tech CEO friends, there’s this betting pool for the first year that there is a one-person billion-dollar company,*” and The Economist recently covered this issue <https://www.economist.com/business/2025/08/11/how-ai-could-create-the-first-one-person-unicorn?>.

2021; Gupta et al., 2024), which make each employee especially valuable. Moreover, startups’ smaller size, high task variability, and uncertain environments result in less formalized roles and greater adaptability,² suggesting that startups may use Gen AI to complement existing workers instead, reallocating labor toward higher-value tasks rather than reducing headcount.

Beyond intensive-margin adjustments within firms, Gen AI may also alter entry dynamics on the extensive margin, shaping total startup employment. Standard models of entrepreneurial dynamics suggest a tight link between startup growth, valuation, and external financing. If Gen AI alters the production function by reducing labor required by enhancing team productivity, its effects on firm value and financing demand may depart from historical patterns. Understanding how Gen AI reorganizes labor, productivity, and financing in startups thus presents a central empirical question, which this paper seeks to tackle.

To address this question, we assemble a new dataset that links information on high-growth startups, their workforce, and their financing outcomes. We identify startups using Crunchbase and Pitchbook, two of the most comprehensive databases that track firms likely to receive venture capital financing and cover a substantial share of U.S. high-growth entrepreneurship (Cook et al., 2023). We restrict attention to firms founded between 2018 and 2021³, ensuring that all startups in our sample were established before the introduction of ChatGPT.

The result is a quarterly panel of 94,789 startups, observed from their founding through 2025. On average, these firms employ 11 workers, raise \$3 million in external funding, and generate \$700,000 in annual revenue. Firm-level Gen AI exposure is constructed by integrating each firm’s pre-ChatGPT occupational distribution with occupation-specific estimates of Gen AI automation potential from Eisfeldt et al. (2023). Specifically, for each firm, we measure the share of job tasks exposed to Gen AI in the quarter before ChatGPT’s introduction (September to November 2022). This measure

²Ford and Slocum (1977) provides a synthesis of this literature.

³Appendix Table C5, Column 6, shows that our results are robust to defining startups using an alternative founding-year cutoff of 2016–2021.

captures the extent to which each startup’s workforce was exposed to advances in Gen AI. Roughly one-third of the tasks in our sample are exposed to Gen AI. This measure captures a much broader set of activities than just high-technology occupations, as roughly 16% of tasks involve technology-oriented work, as defined by Hecker (2005).

We find that the introduction of ChatGPT led to a substantial decline in employment at startups with high exposure to Gen AI. Startups fully exposed to Gen AI experienced an 8% reduction in employment, relative to the pre-treatment level. This effect emerged rapidly, within two quarters of the ChatGPT release, and persists till the end of our sample period. Importantly, we do not observe any pre-trends before the introduction of ChatGPT, supporting a causal interpretation.

A decomposition of employment changes reveals that reductions primarily affected less senior employees and those in implementation or execution roles, while the share of senior employees rose. Post-exit outcomes indicate that about half of the exiting workers faced delays of about six months in reemployment, and many took positions that were lower-paying but less exposed to Gen AI. These results suggest that Gen AI is replacing the least-experienced talent, who bear costs in transitioning to alternative roles.

We next examine which startups are best positioned to capitalize on Gen AI adoption. We explore three potential mechanisms: it can be that Gen AI adoption is driven by managerial skill, Venture Capital (VC) monitoring, or the technical capacity of the founder. Our evidence strongly supports the latter. Founders with coding experience or computer science backgrounds achieved the largest employment reductions. In contrast, we do not find any conclusive evidence for startups of serial entrepreneurs or VC-backed firms having a differential effect.

We validate that our Gen AI exposure measure captures actual Gen AI adoption using job-posting data. Firms with higher Gen AI occupational exposure exhibit larger increases in Gen AI-related postings. Our results are robust to alternative measures of Gen AI exposure. Our results are also not driven by broader trends in the tech

sector: the estimates remain robust when restricting to non-tech occupations, excluding tech industries or tech-heavy states such as California and New York. The results are not driven by firms pivoting to Gen AI-related products or by shifts in VC financing dynamics. We also find no significant effects from alternative explanations such as changes in job remoteness or offshorability.

We next examine firm-level sales and funding milestones to assess whether reductions in labor translated into firm-level efficiency gains. Using code contribution data from GitHub and sales data from Data Axle, we find that firms more exposed to Gen AI experienced a 64% increase in productivity as measured by GitHub contributions, along with a 20% higher likelihood of achieving top-quartile overall and per-employee sales relative to the pre-period mean. These patterns indicate that output rose even as headcount declined, with remaining employees taking on tasks previously performed by displaced workers.

At the same time, startups with greater Gen AI exposure were significantly more likely to achieve additional funding milestones, with a 21% increase over the pre-period mean in the probability of receiving VC funding. The joint increase in per-employee sales and funding is consistent with Gen AI altering the production function of labor, enabling each employee to generate more output. Overall, these findings indicate that Gen AI weakens the traditionally tight link between startup employment growth, productivity, and financing, allowing startups to scale more efficiently with fewer workers.

Having examined existing startups, we next turn to Gen AI's impact on the rate of new firm formation and the financing these new entrants receive. Following the introduction of ChatGPT, VC firms with higher ex-ante investments in markets (industry-state pairs) more exposed to Gen AI occupations experience a 12% decline in average initial funding size, accompanied by a 7% increase in the number of investments. Although the total investment value declined, the change was not statistically significant. These patterns suggest that Gen AI prompted more numerous

but smaller VC investments, boosting new startup formation.

We verify that these patterns also hold at the aggregate market level. Markets more exposed to Gen AI experienced a 6% increase in the total number of startups formed following ChatGPT’s release, relative to the pre-period. At the same time, the average size of both newly formed and incumbent startups declined, resulting in an overall 4% reduction in average startup size in the post period. Importantly, the additional employment generated by new firm entry largely offset the job losses among existing startups, leaving the net effect on total market-level employment close to zero. However, this aggregate null effect conceals meaningful heterogeneity in composition. In particular, Gen AI exposure increased the share of managerial and senior positions by 0.5 percentage points, indicating a reallocation in the structure of employment rather than a uniform contraction.

This paper contributes to the growing literature on Gen AI. Prior work shows that Gen AI can raise the valuations of large firms by reducing hiring (Eisfeldt et al., 2023). While Eisfeldt et al. (2023) focus on labor outcomes for large or public firms, the implications for startups remain theoretically ambiguous. Our paper fills this gap by providing the first evidence on how Gen AI reshapes entrepreneurship. We find that employment reductions are steepest among young firms that are less than five years old. Unlike public firms, where Gen AI primarily reduces hiring, startup employment declines arise mainly from worker separations. This pattern is consistent with their higher employment volatility (García-Trujillo et al., 2023) and with startup employees’ lower valuation of job security, which supports more short-term contracting (Sauermann, 2018). In other work, Brynjolfsson et al. (2025b,a) study Gen AI’s effects on labor productivity in call centers and on younger workers. We extend this work by linking the labor impact of Gen AI to entrepreneurship, showing how changes in Gen AI-induced productivity affect startup financing and new firm formation.

This paper also builds on the literature that examines how new technologies affect startup outcomes.⁴ Startups have historically been slower adopters of new technologies

⁴The broader literature in corporate finance studies the frictions that induce firms to adopt

and production processes than incumbent firms (Acemoglu et al., 2024; Bessen, 2020; Zolas et al., 2020). Foundational work by Schumpeter (1942) attributes this pattern to startups’ limited access to capital and smaller economies of scale relative to larger firms. However, new technologies offering pay-per-use access can mitigate these constraints. For instance, Ewens et al. (2018) show that Amazon Web Services reduced fixed capital expenditures and encouraged VC experimentation, increasing startup investments but weakening long-term performance. In contrast, Gen AI lowers variable labor costs and enhances productivity over the firm’s life cycle, which not only raises the number of investments but also improves outcomes for startups.

Traditionally, startup growth has been closely tied to growth in labor demand. Decker et al. (2016) document a strong positive correlation between business dynamism and labor market dynamism in the United States. Similarly, Haltiwanger et al. (2013); Decker et al. (2014); Adelino et al. (2017) show that new firms contribute disproportionately to job creation. Our findings suggest that Gen AI may be weakening this relationship. While Gen AI fosters the creation of new firms, it simultaneously displaces early-career employees, generating adjustment costs for these workers. Taken together, these results underscore a growing policy challenge, namely, how to promote start-up formation and technological innovation while mitigating the labor market costs borne by younger cohorts.

2 Data

2.1 Startup Data

The first step in analyzing the impact of Gen AI on startups is to identify the relevant universe of startup firms. We focus on *high-growth startups*, which are most likely to receive venture capital (VC) financing. Following Cook et al. (2023), who

technologies, along with the implications for such technology adoption. See, for example, Geng et al. (2022); Tuzel and Zhang (2021); Bena et al. (2022); Bena and Simintzi (Forthcoming); Ouimet et al. (2025); Ma et al. (2025).

show that Crunchbase⁵ and PitchBook⁶ cover a substantial share of U.S. high-growth entrepreneurship, we use these two sources to identify our sample of high-growth startups.

We construct a comprehensive database of startups by taking the union of firms listed in Crunchbase and PitchBook. For each firm, we record its founding year, location, industry, and fundraising activity over time. We restrict the sample to startups founded between 2018 and 2021, ensuring that all firms in our analysis were established before the introduction of ChatGPT⁷. Applying these filters, we obtain 116,869 unique startups from Crunchbase and 40,794 unique startups from PitchBook founded during this period in the US.⁸

2.2 Employment Data

To measure startup-level employment, we rely on individual-level data from Revelio Labs, which compiles publicly available resume information from LinkedIn. For each worker, we observe their job title, firm, location, and employment period. Revelio additionally classifies each position using 6-digit SOC occupation codes and assigns a seniority level across four categories. Job tenure is directly measured based on the recorded start date. The dataset also includes estimated salaries, generated using a regression model that predicts compensation based on role, seniority, company, and geography.⁹

We aggregate these individual-level observations to construct startup-level measures of employment, team composition (including occupations, job titles, seniority, tenure), and worker flows (hires and separations). For employees who depart, we track

⁵Crunchbase (established in 2007) is a global repository of information on companies, investors, and key individuals within the startup ecosystem. It records information on firm names, locations, industries, founders, and firm-level events such as funding rounds, IPOs, and acquisitions.

⁶PitchBook (also established in 2007) similarly provides detailed data on deal activity for private firms, including detailed information on investments, exits, funds, and investors for VC back firms.

⁷Appendix Table C5, Column 6, shows that our results are robust to defining startups using an alternative founding-year cutoff of 2016–2021.

⁸Approximately 75% of the startups covered in PitchBook also appear in Crunchbase.

⁹Revelio trains this model on over 200 million salaries drawn from job postings and publicly available labor certification applications, adjusting for country-level inflation to account for changes over time.

their subsequent positions to analyze post-shock employment outcomes. Additionally, we extract founder characteristics, including educational background and prior work experience.

We link Revelio’s employment data to startups listed in Crunchbase and PitchBook using normalized LinkedIn and startup homepage URLs. Appendix Table C1 reports the corresponding match rates. We successfully merge 87,549 of the 116,869 Crunchbase startups (a 75% match rate) and 28,330 of the 40,749 PitchBook startups (a 70% match rate). The combined employee–startup dataset includes 94,789 unique startups, which serves as our baseline sample.¹⁰

2.3 Startup Fundraising

Both Crunchbase and PitchBook provide detailed information on startup fundraising activities. For each financing event, they report the date, round type and series, amount raised, lead investor, number of participating investors, and, in some cases, the valuation. PitchBook additionally identifies non-lead syndicate members involved in each round. We combine data from both sources to construct a comprehensive dataset of startup financing, capturing the types of rounds and the total amount of funding raised.

2.4 Startup GitHub Contributions

We link our baseline startup sample to GitHub organizations and the repositories they own to obtain startup-level GitHub contributions. GitHub provides a granular, activity-based measure of software development output, enabling us to quantify startup productivity through contribution metrics.

We link our baseline startup database to organizations with active contribution events during the sample period using normalized URLs from LinkedIn, Twitter, Facebook, and the startup homepage. We also perform fuzzy matching based on

¹⁰A total of 21,090 startups appear in both PitchBook and Crunchbase, implying that PitchBook contributes 7,240 additional startups beyond those covered by Crunchbase.

company names using the rapidfuzz library and requiring a match score above 95%.¹¹ This procedure links 18,947 startups to GitHub, corresponding to a 20.0% match rate.

2.5 Startup Sales

We link our baseline startup sample to Data Axle (formerly Infogroup) to obtain startup-level sales data. Data Axle’s Historical Business dataset provides annual snapshots of U.S. businesses, locations, industries, and sales.

We perform fuzzy matching between our baseline startup database and Data Axle records based on company names, founding years, and locations, using the rapidfuzz library and requiring a match score above 95%.¹² This procedure links 43,250 startups to Data Axle, corresponding to a 45% match rate.

2.6 Measuring Startup Exposure to Gen AI

We estimate a startup’s exposure to Gen AI by leveraging the occupational-level exposure measures developed by Eisfeldt et al. (2023). For each firm, we obtain its ex-ante occupational composition at the six-digit Standard Occupational Classification (SOC) level from Revelio Labs, measured in the quarter immediately preceding the introduction of ChatGPT (September–November 2022). A firm’s exposure to Gen AI is constructed as the employment-weighted average of its occupations’ exposure scores. Specifically, the firm-level exposure to Gen AI is defined as

$$E_i = \sum_{o \in i} EmpShare_{i,o} \times E_o, \quad (1)$$

where $EmpShare_{i,o} = \frac{emp_{i,o}}{emp_i}$ is the employment share of occupation o in firm i , and E_o denotes the occupational-level exposure to Gen AI. This measure captures the extent to which a firm’s workforce is engaged in tasks more exposed to Gen AI technologies.

¹¹We implement fuzzy matching based on RapidFuzz weighted ratios, retaining only pairs with a similarity score of at least 95% to ensure high-quality matches.

¹²We implement fuzzy matching based on RapidFuzz weighted ratios, retaining only pairs with a similarity score of at least 95% to ensure high-quality matches.

2.7 Summary Statistics

Table 1 presents summary statistics for our sample. Startups in our sample have both a mean and median Gen AI task exposure score of 31%, with a standard deviation of 10 percentage points. This is slightly lower than the 35% task exposure to AI reported for publicly listed firms by Eisefeldt et al. (2023). The AI task exposure differs from exposure to tech-oriented occupations, defined following Hecker (2005). The mean tech-occupation exposure is 16%, consistent with the idea that occupations outside traditional tech roles are also replaceable by Gen AI.

Our main sample consists of startups founded between 2018 and 2021 and listed on Crunchbase or Pitchbook. This design allows us to focus on high-growth firms that are one to three years old. On average, these firms employ about 11 people. This is comparable to data from 2022 Business Dynamics Statistics of High Growth Firms (BDS-HG, as described in Kim et al. (2024)), who find that firms in the highest growth decile aged one to five years as of 2021 employ around 10 workers per firm.

We find that the average employee salary in our data is about \$120,000, which aggregates to about \$1.5 million at the firm level. This aligns closely with the average employee salary of \$ 105,000-110,000 for startups reported across publications¹³. The startups in our sample have raised an average of \$3.1 million in cumulative funding and generate approximately \$686,000 in annual sales. The funding figures are consistent with companies being near the median size of seed rounds (typically raised 1 to 3 years after founding), with an average size of around \$ 2.5 Million as of 2024¹⁴. Table A2 reports the distribution of these startups across industries and states. The Information Technology sector is the largest, accounting for 36.8% of all firms, consistent with numbers from the National Venture Capital Association and Pitchbook report, which reported about 37% of total VC funding going to the software industry. Geographically, the startups are concentrated in major economic hubs. California hosts the largest share

¹³<https://wellfound.com/hiring-data/i/startups-1?> reports an average of \$106,000 for startup employees, and <https://carta.com/data/startup-compensation-h1-2024/> reports a salary range of \$ 105,000 to 125,000 for mid-level startup employees

¹⁴<https://carta.com/uk/en/data/top-seven-states-seed-funding-2024/>

(21.9%), followed by New York (11.0%) and Texas (8.9%). Together, these three states account for over 41% of firms in our sample.

Finally, we examine employee-level outcomes. About 57% of employees hold senior positions. The average execution score, which measures hands-on implementation activities¹⁵, indicates that roughly 35% of employees are directly involved in implementation tasks. The average hiring rate is 7.4%, compared to an exit rate of 5.6% over the same period. The higher hiring-to-exit ratio is consistent with employment growth early in the life cycle of startups.

3 Impact on Existing Startups

3.1 Empirical Specification

We begin by examining the impact of Gen AI on the employment of existing startups. In our baseline analysis, regression estimations are run at the startup-quarter level between July 2021 and August 2025 (or annual level where quarterly outcomes are not defined) and take the following form:

$$Y_{it} = \beta E_i \times \text{Post}_t + \theta_i + \theta_{kt} + \theta_{st} + \gamma' X_i \times \theta_t + \epsilon_{it}, \quad (2)$$

where Y_{it} denotes the outcome of startup i in period t , and E_i measures the Gen AI exposure of startup i . Post_t is an indicator equal to one for periods following the introduction of ChatGPT in period T .¹⁶ We include startup fixed effects (θ_i) to absorb time-invariant differences across firms that may correlate with Gen AI exposure. To account for common time-varying shocks, we further include industry-by-time fixed effects (θ_{kt})¹⁷ and state-by-time fixed effects (θ_{st}). In our most granular specification, we additionally control for the time-varying influence of pre-determined firm characteristics

¹⁵We predict managerial score using a chat prompt to LLM API as detailed in Appendix A.1.

¹⁶For quarterly analysis, T corresponds to December 2022–February 2023. For yearly analysis, $T = 2023$.

¹⁷We define industries using 686 Crunchbase categories.

by interacting ex-ante firm characteristics (X_i), including founding year, pre-shock size, and pre-shock growth rate, with time fixed effects (θ_t). We verify the absence of pre-trends by estimating a dynamic version of our baseline equation as follows:

$$Y_{it} = \sum_{\tau \neq -1} \beta_{\tau} E_i \times \mathbf{1}\{t - T = \tau\} + \theta_i + \theta_{kt} + \theta_{st} + \gamma' X_i \times \theta_t + \varepsilon_{it}, \quad (3)$$

where the coefficients β_{τ} trace the evolution of outcomes for startups with different levels of Gen AI exposure relative to the period just before ChatGPT's introduction.

3.2 Impact on Employment

We examine changes in employment among more exposed firms following the introduction of ChatGPT using a Poisson pseudo-maximum likelihood (PPML) specification. Following Cohn et al. (2022) and Chen and Ewens (2025), we avoid the $\log(1 + y)$ transformation, as PPML handles zeros and heteroskedasticity more appropriately. Hence, our coefficients can be interpreted as semi-elasticities, i.e., in terms of percentage changes in the dependent variable. Table 2 presents the estimates. Columns (1)–(3) progressively add firm and time, industry–time, and state–time fixed effects, while Column (4) adds our most comprehensive set of controls with firm, industry–time, state–time, and pre-existing firm characteristics by time fixed effects. The estimated coefficients are consistently negative and of similar magnitude across Columns (1)–(4), confirming that the result is not sensitive to alternative controls. In our preferred specification in Column (4), the coefficient on Gen AI Exposure $\times$ Post is -0.078 , implying that a fully exposed startup (exposure = 1) experiences about a 7.5%¹⁸ decline in employment after the introduction of ChatGPT. In comparison, for a median-exposure startup (exposure = 0.405), the implied effect is a 3.1% decline¹⁹ in employment. We find that these reductions in employment led to a similar decrease in

¹⁸Calculated as $\exp(-0.078) - 1$.

¹⁹Calculated as $\exp(-0.078 \times 0.405) - 1$.

total salary bill for startups.²⁰

Figure 1 plots the corresponding event-study estimates from Equation 3, based on the same specification as Column (4). The pre-period coefficients are flat and statistically indistinguishable from zero in the quarters leading up to the release of ChatGPT, indicating that there is no evidence of anticipation. The employment effect appears sharply in the first quarter following the release. Employment declines monotonically thereafter, consistent with startups adjusting their labor force after adopting Gen AI. The decline persists through the end of the sample. The timing and persistence of these effects support a causal interpretation that Gen AI’s introduction led to immediate and sustained employment contractions among more exposed startups, rather than reflecting pre-existing trends or cyclical factors.

Columns (5)-(6) of Table 2 examine whether the employment decline documented earlier reflects changes in hiring or separations. The dependent variables are the hire rate (new hires relative to average employment) and the exit rate (separations relative to average employment). Columns (5) report estimates for hiring, while Columns (6) report estimates for exits.

The results show that hiring remains largely unchanged after ChatGPT’s introduction, whereas separations increase sharply among more exposed startups. In Column (5), the estimated post-period change in the hiring rate for fully exposed startups is small and statistically insignificant at about 3% relative to the mean, while in Column (6), the exit rate rises by roughly 25% of the mean exit rate. Thus, after ChatGPT’s release, more exposed startups experienced substantially higher employee turnover, with the magnitude of the increase in exits being roughly four times larger than that of new hires. This pattern indicates that the contraction in employment observed earlier primarily reflects elevated separations rather than a slowdown in hiring.

²⁰Internet Appendix Table C3 shows a similar reduction in the total salary bill of startups. The total salary is computed as the sum of predicted Revelio salaries for all startup employees. Revelio estimates employee-level salaries using a model trained on over 200 million salary observations, adjusted for inflation. The columns are defined in the same way as in Table 2. In the preferred specification in Column (4), total payroll among fully exposed startups declined by about 8.6% following the introduction of ChatGPT.

This increase in employee separations in startups following ChatGPT’s introduction is consistent with the idea that startups and small firms exhibit greater employment volatility (García-Trujillo et al., 2023) and may adjust labor more aggressively following shocks, due in part to weaker binding long-term contracts and employee motives that place less weight on job security (Sauermann, 2018).

3.3 Impact on Separations

To better understand the nature of these separations, Table 3 decomposes the exit effects by the duration of non-employment and the characteristics of workers’ subsequent roles.

Table 3 Panel A examines whether employees leaving more exposed startups were more likely to remain out of work for extended periods. Columns (1)–(4) report changes in exit rates for unemployment durations ranging from 1 to 12 months following separation. We find large and statistically significant effects up to six months after separation, but these effects begin to attenuate by the one-year mark. Specifically, among employees at fully Gen-AI-exposed startups, the probability of remaining unemployed for at least six months increases by 11% relative to the mean separation rate following the release of ChatGPT. By contrast, the corresponding effect at the one-year horizon is approximately 5% of the mean. Overall, these results suggest that workers experience displacement lasting roughly six months, but many appear to re-enter employment within a year of job loss.

Table 3 Panel B investigates the reallocation paths of departing employees. We classify exits based on whether the worker’s next observed position is in an occupation with higher or lower Gen AI exposure, with a higher or lower salary than before, or whether it involves founding a new firm. The results show that most transitions are toward less AI-exposed or lower-paying jobs. For fully exposed startups, exits to lower-exposure jobs rise by about 37% of the mean exit rate (Column (1)), while exits to higher-exposure jobs fall by roughly 26% of the mean exit rate (Column (2));

the difference is highly significant (t -test $p < 0.001$). Similarly, exits to lower-salary positions increase by approximately 21% of the mean exit rate (Column (3)), whereas transitions to higher-salary positions remain statistically unchanged in Column (4) (t -test $p = 0.003$). Finally, exits to become founders are statistically indistinguishable from zero, with an implied magnitude of only 1.5% for fully exposed startups, indicating little evidence that separations translate into entrepreneurship.

Taken together, these findings indicate that post-shock employment losses were largely non-voluntary: displaced workers at more Gen-AI-exposed startups faced longer non-employment spells and downward occupational mobility, consistent with involuntary displacement rather than voluntary job changes.

3.4 Impact on Employment Composition

Table 4 and Figure 2 examine how Gen AI exposure reshaped the composition of startup teams. We focus on four outcomes: the share of senior employees (those at the manager level or above), median employee tenure (years since starting the current role), the average execution score, and the average managerial score. The last two outcomes are based on a machine-learning measure of how implementation or management-intensive a role is, such as coding, analysis, or operations.²¹

Columns (1), (3), (5), and (7) include firm, industry-time, and state-time fixed effects, while Columns (2), (4), (6), and (8) additionally control for pre-existing firm characteristics by time fixed effects. The results show systematic and economically meaningful compositional shifts. Startups that are fully exposed to Gen AI increased the share of senior employees by roughly 3.1% of the mean (Column (2)) and median tenure by 8.7% of the mean of the firm’s median tenure (Column (4)), indicating that more exposed firms retained longer-tenured workers and more senior employees even as employment declined. The average execution score declined by about 2.2% of the mean (Column (6)), and the average managerial score increased by about 1.6% of the

²¹See Appendix A.1 for details on the construction of the Execution Score.

mean (Column (8)) for firms fully exposed to Gen AI after ChatGPT’s introduction. These patterns are robust across Columns (1), (3), (5), and (7) with fewer controls, confirming that they are not driven by model choice.

Figure 2 plots the dynamics of these effects: exposed startups reorganize their teams toward more senior, experienced, and less execution-focused employees. This pattern is consistent with selective downsizing—firms retain higher-skill workers better able to leverage or complement Gen AI tools, while trimming more junior or implementation-heavy positions. These findings are consistent with Brynjolfsson et al. (2025a), who also document that Gen AI adoption reduced new firm hiring in the ADP data, which is more representative of larger firms in the economy (Cajner et al., 2018). Our results show that similar patterns hold even among startups, though the effect in our setting is driven primarily by firm exits rather than hiring reductions among larger firms. They also echo the mechanism highlighted by Hampole et al. (2025), who show that tasks and roles with greater task variation adapt more effectively to AI change. Senior employees or those in managerial occupations typically span a broader set of tasks and therefore experience more heterogeneous AI exposure, allowing them to reallocate effort and capture productivity gains from Gen AI adoption.

3.5 Mechanisms and Robustness

The effects of Gen AI exposure may operate through distinct founder-level or VC-level mechanisms. Technically skilled founders may be better positioned to implement and integrate Gen AI tools, accelerating productivity adjustments but also deepening employment cuts. At the same time, founders with prior entrepreneurial experience may better cushion these adjustments through managerial know-how. Finally, it may be that startups with VC funding, and hence benefiting from VC monitoring, have more appetite or ability to adopt Gen AI to replace labor. We examine these potential channels by comparing heterogeneous responses across founder and investor characteristics, including technical background, prior entrepreneurial experience, and

pre-ChatGPT venture funding.

Table 5 Panel A tests whether the employment effects differ between startups led by founders with technical backgrounds, measured by whether the founder has a computer science (CS) degree or prior coding experience, and those led by non-technical founders. Columns (1)–(2) examine heterogeneity by CS degree, while Columns (3)–(4) focus on coding experience. Columns (2) and (4) include the full set of firm, industry–time, state–time, and firm-characteristics–by–time fixed effects.

The results show clear differences in magnitude. Among non-technical founders, the estimated post-ChatGPT employment decline for fully exposed startups is modest at around 3–5%. In contrast, for startups led by CS-degree founders, employment declines by roughly 19%, and for those led by founders with coding experience, the decline is about 14%. These findings suggest that founders with greater technical expertise adopted Gen AI more rapidly and reorganized production accordingly, leading to sharper immediate workforce contractions.

Figure 3 illustrates the dynamics of these patterns. Panel A plots event-study estimates by founder CS degree, and Panel B by coding experience. In both cases, startups led by technical founders exhibit no differential pre-trends before ChatGPT’s release but experience markedly larger employment declines immediately afterward. The gap between technical and non-technical founders persists through the post-shock period, consistent with sustained organizational restructuring among technically-led startups.

Panel B of Table 5 examines whether similar heterogeneity arises from other founder attributes. Columns (1)–(2) interact Gen AI exposure with an indicator for serial entrepreneurs. The estimated differences are economically small and statistically insignificant, indicating that the heterogeneity in employment effects is primarily driven by the founder’s technical skill rather than by prior entrepreneurial experience. Columns (3)–(4) interact Gen AI exposure with an indicator for whether the firm had received venture capital funding prior to ChatGPT’s release. We again find no evidence

of these results being driven by startups who have already raised VC funding.

3.6 Validation Tests

One concern is whether firms more exposed to Gen AI-related tasks actually adopt Gen AI technologies. To validate this, we examine how exposed startups' use of Gen AI evolved following ChatGPT's introduction. We use job posting data from the Revelio COSMOS database covering July 2021 to August 2025. The outcome variable is the share of Gen AI job postings, defined as the fraction of a startup's job postings that include Gen AI-related keywords such as ChatGPT, Claude, LangChain, Llama, or Falcon.²² These keywords capture both explicit adoption and demand for AI-complementary skills.

Table 6 reports estimates from Equation 2 using startup-year observations. Columns (1)–(4) progressively add firm and time, industry–time, state–time, and firm-characteristics–by–time fixed effects. Across all specifications, startups with higher pre-shock exposure to Gen AI display a significantly greater increase in Gen AI-related hiring after ChatGPT's release. In the preferred specification in Column (4), fully exposed startups increased the share of postings mentioning Gen AI by about 1.4 percentage points. The estimates are robust across columns, confirming that the results are not driven by alternative control sets. These results indicate that more Gen AI-exposed startups not only reorganized their internal workforce but also actively expanded external hiring for Gen AI-related roles. This pattern is consistent with rapid technology adoption, where firms simultaneously reduce headcount in substitutable tasks while seeking new workers with skills to integrate and deploy emerging AI tools.

Internet Appendix Table C4 provides direct evidence on the mechanisms underlying the employment effects of Gen AI by decomposing exposure into core and supplemental tasks. The logic follows Eisefeldt et al. (2023), who interpret exposure through core tasks as capturing occupations whose primary functions overlap with those automated by

²²The full list of keywords are reported in Appendix A.2.

large-language-model technologies, while supplemental-task exposure reflects ancillary functions (such as communication, documentation, and analysis) that can be augmented by Gen AI tools. If our estimates indeed capture the impact of task-level Gen AI, we should observe the strongest effects in occupations whose core tasks are exposed to Gen AI. By contrast, occupations that are primarily supplemented by Gen AI should exhibit limited effects or even modest employment gains.

Column (1) replicates our baseline estimate. Column (2) interacts Gen AI exposure with the share of supplemental tasks. We find that the magnitude of the employment contraction strengthens to roughly 14% for startups with only core tasks. For startups with fully supplemental tasks, the magnitude is about 39%²³ larger than that for core tasks, meaning that the overall effect is positive for purely supplemental tasks. Column (3) reverses the baseline, using supplemental tasks as the reference group and interacting with the share of core tasks. In this specification, employment among startups fully exposed to supplemental tasks increases by approximately 19%. These results align with the notion that Gen AI exposure has positive effects for occupations where the technology acts as a supplement, but negative effects where it substitutes for core tasks. On average, however, the results suggest that Gen AI primarily substitutes rather than complements labor.

3.7 Robustness Tests

We conduct a series of robustness checks to ensure that our findings are not driven by model specification, sample composition, or the construction of the Gen AI exposure measure.

First, we may be concerned that our results are driven by broader trends in the technology industry. Although all specifications include industry-by-time and state-by-time fixed effects to address this concern, Internet Appendix Table C5 takes an additional step. We re-estimate the baseline excluding technology-related occupations

²³Calculated as $\exp(0.326) - 1$.

in Column (1), technology-related industries in Column (2), and states with a high concentration of high-technology activity in Column (3). The coefficients show that the post-ChatGPT employment decline remains sizable and statistically significant even among startups outside the technology sector, keeping only non-technology occupations, and without states such as California and New York. These results indicate that employment contractions associated with Gen AI exposure are broad-based rather than concentrated in specific sectors or states.

Second, a potential concern is that our results may reflect changes in product markets, as some startups might have introduced new Gen AI-related products following ChatGPT’s release. To address this, Internet Appendix Table C5, Column (4), re-estimates the baseline excluding all startups with Gen AI-related descriptions, thereby removing firms that appear to have shifted their product focus after the introduction of ChatGPT. The results continue to show employment declines of similar magnitude to the baseline, suggesting that our findings are not driven by product-level changes.

Third, one concern is that our results might be confounded by concurrent shifts in the VC industry that are unrelated to Gen AI. For instance, startups with more Gen AI-replaceable tasks may also have been those with greater VC exposure and therefore more sensitive to changes in the interest rate environment as VCs pulled back on funding. Our baseline specification includes industry $\times$ time and location $\times$ time fixed effects to absorb many of these broader trends. In addition, we show that our findings are unchanged when we exclude all firms with any ex-ante VC financing, the firms for which this alternative explanation would be most relevant.

Fourth, a concern might be on the definition of startups in our sample. Column (6) uses an alternate sample of startups as those which were founded from 2016 onwards (instead of 2018). We find that our results are robust to this change in definition.

Fifth, a key concern is whether the results truly reflect the impact of Gen AI, or instead capture effects among firms that engaged in greater remote work²⁴, offshoring, or

²⁴In fact, Gupta et al. (2025) document the opposite effect, greater scaling for startups that go more

technology related occupations. Internet Appendix Table C6 addresses this concern by adding controls for occupational characteristics that could confound the relationship between Gen AI exposure and employment, each interacted with the post period. Following Dingel and Neiman (2020), Autor and Dorn (2013), and Hecker (2005), we control for each startup’s share of remote, offshorable, and technology-related occupations, respectively. If these alternative channels were driving the employment decline, we would expect significant coefficients on the respective characteristic $\times$ post interactions, along with a substantial attenuation of the Gen AI $\times$ post coefficient. Instead, the estimated post-ChatGPT employment decline for fully exposed startups remains approximately 6%, even when all three controls are included simultaneously. Moreover, the interaction coefficients on remote, offshorable, and technology-related shares are economically small and statistically insignificant.

Sixth, we might be concerned about our measure of Gen AI exposure. Internet Appendix Table C7 tests the sensitivity of the results to alternative definitions of Gen AI exposure. Callaway et al. (2024) document that the parallel trends assumption may be biased under continuous treatment in diff-in-diff and recommends using binary treatment to verify pre-trends. Column (1) uses a binary high-exposure indicator (above-median Gen AI exposure).²⁵ Column (2) employs only core tasks when calculating the Gen AI exposure for any startup. Column (3) replaces our baseline measure with the Gen AI applicability index from Tomlinson et al. (2025), which captures the potential for Microsoft Copilot integration based on task descriptions in O*NET. Across all three measures, the estimated post-ChatGPT employment effects remain negative and statistically significant, confirming that the results do not depend on any single exposure definition.²⁶

Finally, Internet Appendix Table C8 re-estimates the main models using log

remote

²⁵Internet Appendix Figure B3 replicates the baseline event study using a binary treatment that classifies startups above the median of Gen AI exposure as “high-exposure.” The dynamic pattern remains unchanged.

²⁶For a startup with median exposure from core tasks (0.305), the implied effect is -2.6%. For a startup with a median Gen AI applicability index from Tomlinson et al. (2025) (0.237), the implied effect is -5.5%. These magnitudes are broadly consistent with our baseline exposure (-3.1%).

specifications instead of Poisson pseudo-maximum likelihood. The results are qualitatively and quantitatively consistent: employment among fully exposed startups declines by roughly 10%, while sales per employee rise by about 25%. These findings confirm that the productivity and financing effects documented earlier are not artifacts of functional-form assumptions.

Overall, these robustness exercises confirm that the estimated effects of Gen AI exposure are stable across alternative specifications, exposure measures, and subsamples.

3.8 Comparison with Incumbent Firms

A natural question is whether these employment effects persist across firms of different vintages and sizes. We construct a sample of 464,877 U.S. firms on LinkedIn with non-missing information on industry, state, founding year, and active employees in the quarter immediately preceding ChatGPT’s introduction. We then estimate a triple-difference version of our baseline specification, where the key regressor is the triple interaction between firm-level Gen AI exposure, an indicator for the firm’s age bucket, and a post-ChatGPT dummy. The dependent variable is total firm employment, estimated using Poisson pseudo-maximum likelihood regressions, consistent with our baseline approach.

Internet Appendix Figure B1 presents the heterogeneous effects of Gen AI exposure on employment across firm age groups. The results show that startups with higher Gen AI exposure experienced significant employment declines following ChatGPT’s release. This negative effect weakens with firm age, turning slightly positive for firms that are 25 years or older.

Internet Appendix Figure B2 further examines heterogeneity by firm size. We again find large negative effects for young firms (1 to 5 years) across all size categories. For firms aged 6 to 25 years, the estimated effects are negative but statistically insignificant. Among the oldest firms, there is a modest positive effect for small (≤ 50 employees) and

mid-sized (51–1,000 employees) firms. In contrast, large and old firms (25+ years and >1,000 employees) exhibit significant negative employment effects of about 3%, a quarter in magnitude to the 12% effect observed for young (1–5 years) firms.

Two main takeaways emerge from this analysis. First, the negative employment effects of Gen AI are most pronounced for startups and young firms. In contrast, firms aged 5–25 years appear less responsive in adjusting their workforce composition. Second, large and old firms also reduce headcount, consistent with the decline in job postings among public firms documented in prior work (Eisfeldt et al., 2023). However, this effect is substantially stronger for startups, aligning with the notion that younger and smaller firms exhibit greater employment volatility (García-Trujillo et al., 2023) and may adjust labor more aggressively in response to shocks.

4 Startup Productivity & Financing

The preceding sections show that startups with greater exposure to Gen AI rapidly restructured their workforces, reducing headcount in Gen AI-substitutable tasks while expanding hiring for Gen AI-related roles. A natural question is whether these organizational changes translated into measurable gains in startup performance. Previous research links firm-level dynamism to joint increases in employment and productivity (Decker et al., 2014). In contrast, Gen AI may act as a technological shock to the production function of startups that simultaneously displaces labor and enhances efficiency, thereby breaking the traditional employment–productivity nexus. We now turn to the firm-level outcomes that capture these mechanisms most directly: performance and fund-raising.

Table 7 evaluates how Gen AI exposure affected startup performance. Panel A measures productivity using the number of GitHub contributions made to repositories directly owned by the startup. Prior work by Gupta et al. (2024) shows that GitHub contributions correlate strongly with other employee-level (such as promotions and publications) and firm-level (including Tobin’s q , and organizational capital)

productivity measures making them a useful proxy for coding output. While one concern is whether this additional coding translates into meaningful firm-level gains, Panel B addresses this by measuring productivity using the dollar value of startup sales. Columns (1)–(2) present results for overall productivity, and Columns (3)–(4) examine productivity per employee. Columns (2) and (4) report our preferred specifications, which include firm, industry–time, state–time, and firm-characteristics–by–time fixed effects.

The results indicate that Gen AI exposure led to substantial improvements in performance among fully exposed startups. Following the introduction of ChatGPT, the number of GitHub contributions for fully exposed firms increased by 64% (Column (2)), while contributions per employee rose by 75% relative to the pre-period mean. While large, these results are consistent with evidence from field experiments which find 55% increase in developer productivity following adoption of Gen AI coding tools (Peng et al., 2023; Gambacorta et al., 2024). Similarly, the probability of being in the top quintile of sales rose by about 25% (Column (2)), and the likelihood of being in the top quintile of sales per employee increased by roughly 20% relative to the mean (Column (4)), post the shock. Although each measure is individually noisy and available only for a subset of firms, together they suggest that highly exposed startups became both more productive and more efficient, consistent with the efficiency gains following the selective workforce restructuring documented earlier.

We next study fundraising outcomes to estimate whether the employee reductions resulted in better financing outcomes for the startup. Table 8 and Figure B4 present the corresponding estimates.

Table 8 reports three margins: the probability a startup receives any VC round (Columns (1)–(2)), the probability it receives a late-stage round (Columns (3)–(4)), and the cumulative amount raised estimated with Poisson pseudo-maximum likelihood (Columns (5)–(6)). Columns (2), (4), and (6) present our preferred specifications that include firm, industry–time, state–time, and firm-characteristics–by–time fixed effects.

The financing response is strong across all three outcomes for fully exposed startups. Following the ChatGPT shock, the likelihood of receiving any VC round rises by 21% of the mean (Column (2)), and the probability of a late-stage round increases by 59% of the mean (Column (4)), for startups fully exposed to Gen AI. Total cumulative capital raised grows by about 50% (Column (6)) for startups fully exposed to Gen AI.²⁷

Taken together, the findings indicate that Gen AI exposure enhanced firm-level productivity and the ability to raise capital. Startups that reorganized more aggressively not only maintained operations with fewer employees but also generated higher output per worker, consistent with complementarities between Gen AI tools and remaining human capital. In parallel, Gen AI exposure improved fundraising outcomes along both the extensive margin (a higher likelihood of securing any and late-stage rounds) and the intensive margin (larger cumulative amounts), suggesting that Gen AI facilitated faster scaling by enabling startups to achieve fundraising milestones more quickly.

5 Shifts in VC Financing Strategy: New Startups

The evidence so far shows that Gen AI has transformed the internal organization and financing capacity of existing startups, raising productivity and attracting new capital through efficiencies in employment. A natural question is whether these firm-level adjustments also reshaped the broader entrepreneurial ecosystem. If Gen AI lowered entry barriers or shifted investor expectations toward Gen AI-intensive sectors, it may have influenced both the rate and composition of new firm formation. To assess this extensive-margin response, we analyze how investment patterns and new startup entry evolved across VC firms with differing exposure to Gen AI. We extend our sample to include VC-backed startups that raised their first funding round between 2020 and 2024.

²⁷Figure B4 plots the dynamics for the cumulative amount raised, corresponding to Column (6) of Table 8. The event-study coefficients are flat before the shock, then trend upward over the next six to eight quarters, indicating no pre-trend and a sustained post-shock increase in cumulative funding amounts.

Our sample contains 15,830 first-round funding events in PitchBook between 2020 and 2024, associated with 8,208 unique VC firms. The unit of observation is a VC-year. We estimate:

$$Y_{jt} = \beta(HighExposure_j \times Post_t) + \theta_j + \theta_t + \epsilon_{jt}, \quad (4)$$

where Y_j denotes the outcome for venture capital firm j . *HighExposure_j* is an indicator equal to one if VC j 's Gen AI exposure is above median. We measure a VC's Gen AI exposure as the dollar-weighted average exposure of the markets (industry-state pairs) it invested in before the shock. For each market, exposure is computed using occupation-level Gen AI measures from Eisefeldt et al. (2023) combined with the occupational composition of startups in that market in the quarter before ChatGPT's introduction. $Post_t$ equals one for financing rounds occurring after ChatGPT's release (2023 onward). Our baseline specification include VC-firm fixed effects (θ_j), and year (θ_t) fixed effects. We also control for VC founding year $\times$ year and portfolio size $\times$ year fixed effects in our most robust specification. We estimate the regressions using the Poisson pseudo-maximum likelihood estimator, consistent with our baseline approach.

5.1 Initial Funding Size of New Startups

We begin by examining how the introduction of ChatGPT affected the initial funding size of new VC-backed startups. The outcome variable is the average funding size for first-round investments (in millions of dollars) for a VC in a given year. Table 9 reports the results. Column (1) presents the results with just the VC firm and year fixed effects. We add VC founding year $\times$ year and portfolio size $\times$ year fixed effects in Columns (2) and (3), to make sure that differential trends in funding by older or larger VCs do not drive our results.

We find a large and statistically significant decline in initial funding size for more exposed VCs following the introduction of Gen AI, consistent across all specifications. In the most robust specification (Column (3)), the average size of initial investments by highly exposed VCs fell by roughly 12% after ChatGPT's release. These results

are consistent with the notion that lower labor costs have decreased the cost of experimentation, thereby increasing the number of investments a VC fund can make.

5.2 Number and Total Value of New Investments

Next, we study how the number and total value of new investments changed across VC firms more exposed to Gen AI, after the introduction of ChatGPT. The outcome variable is the number and value of first-round investments (in millions of dollars) for a VC in a given year. We estimate the regressions using the Poisson pseudo-maximum likelihood estimator, consistent with our baseline approach. Table 10 presents the results. Our preferred specifications in Columns (2) and (4) include VC, VC founding-year-by-year, and portfolio-size-by-year fixed effects.

The results show that venture activity expanded significantly in highly exposed industries after ChatGPT’s release. Both Columns (1) and (2) show that the number of new investments for Gen AI-exposed VCs increased by roughly 8% following the ChatGPT release. In contrast, Columns (3) and (4) show moderately negative but statistically insignificant effects. Consequently, the total dollar value of investments by more exposed VCs either declined slightly or remained unchanged following the introduction of ChatGPT.

Taken together, these results indicate that exposed VCs reduced initial funding size but compensated by financing more new firms, resulting in minimal overall impact on capital deployment. This pattern reflects increased experimentation, with VCs making more frequent but smaller investments. Similar behavior was documented during the launch of Amazon Web Services (AWS). Ewens et al. (2018) show that AWS lowered fixed startup costs and encouraged a more “spray-and-pray” VC strategy. Gen AI instead cuts variable costs and boosts labor productivity. We find that this shift similarly increases VC experimentation, yielding smaller but more frequent investments. But unlike AWS, Gen AI’s cost savings persist throughout the firm life cycle, aligning with the stronger outcomes we observe for mature startups.

6 Aggregate Impact

We next assess how these startup- and investor-level adjustments translated into broader entrepreneurial dynamics. If Gen AI reshaped startup entry costs, labor demand, or capital allocation, its impact should be visible in the evolution of new firm creation and startup employment across markets. Aggregating to the market level also provides a clearer view of Gen AI’s overall impact on startup employment. While this approach limits our ability to include granular controls, it offers valuable insights into the market-level consequences of Gen AI adoption.

We conduct our analysis on a panel of 8,665 industry–state pairs observed quarterly between July 2021 and February 2025, where a “market” is defined as an industry–state pair.²⁸ The regressions are estimated at the market–quarter level and take the form:

$$Y_{mt} = \beta(HighExposure_m \times Post_t) + \theta_m + \theta_t + \varepsilon_{mt}, \quad (5)$$

where Y_{mt} is the outcome of market m in quarter t , $HighExposure_m$ is an indicator equal to one if market m ’s Gen AI exposure is above median, $Post_t$ is an indicator for quarters after ChatGPT’s introduction (December 2022– February 2025). All specifications include market (θ_m) and year-quarter (θ_t) fixed effect, and our more robust models further add market-characteristics–by–time fixed effects.

Table 11 Panel A presents the market-level effects of Gen AI exposure on startup formation and employment. We find that markets fully exposed to Gen AI experienced substantial and statistically significant increases in startup activity following the introduction of ChatGPT. In Column (2), the number of new startups founded per quarter increased by roughly 6% relative to the pre-ChatGPT mean in a market fully exposed to Gen AI. Column (4) puts this effect in context by showing the net impact on the total stock of startups in a market. We find that the total stock of active startups increased by approximately 1% for more exposed markets after ChatGPT’s release.

²⁸Industries are defined using 686 Crunchbase categories.

We next examine overall employment outcomes. With the increase in new startup formation, we expect employment at these firms to rise, and we find this to be the case. Column (6) shows that employment in newly founded startups increased by approximately 5% in markets more exposed to Gen AI following the ChatGPT introduction.

Figure 4 plots these dynamics. Panels (A) and (B) display flat pre-trends before the ChatGPT release and a persistent, sizable post-shock increase in both the number of new startups and the total stock of active firms.

What remains less clear is the effect on overall startup employment. Previous results indicate that existing startups reduced employment due to Gen AI. Table 11 Panel A Column (8) examines, therefore, total startup employment in more exposed markets following the shock. We find a modest, statistically insignificant 0.5% decline in total employment for fully exposed markets. These findings suggest that reductions in employment at existing startups were largely offset by the formation of new firms, resulting in a negligible net impact on market-level employment.

Table 11, Panel B, examines changes in the composition of startups. Column (2) shows that the average new startup formed in high-exposure markets was 7% smaller after the shock. A similar pattern holds for incumbent firms: existing startups in more exposed markets also contracted, becoming approximately 6% smaller in the post period at the market level. These findings are consistent with our earlier VC- and firm-level results, which show that more exposed VCs funded smaller startups and that more exposed incumbent firms reduced headcount following the Gen AI shock. Column (6) further indicates that, in aggregate, startups operating in exposed markets became about 4% smaller overall. In addition to shrinking in size, the composition of startups shifted. The share of managerial employees increased by 0.5 percentage points at the market level, suggesting a modest tilt toward relatively more managerial-intensive employment following the shock.

Overall, these results indicate that Gen AI spurred a substantial increase in

entrepreneurial entry and early-stage job creation, particularly in more exposed markets. At the same time, startups became smaller on average, driven in part by employment declines among incumbent firms. The reduction in employment at existing startups was largely offset by job creation from the surge in new firm entry, resulting in a negligible net effect on total market-level employment. However, this aggregate null effect masks meaningful compositional changes. In particular, the share of managerial workers increased even at the market level, indicating a shift in the structure of employment rather than a simple change in its level.

7 Conclusion

This paper examines how Gen AI has transformed startup employment, productivity, and financing following the introduction of ChatGPT. Using variation in task exposure across firms, industries, and markets, we show that startups with higher exposure to Gen AI adjusted rapidly after the shock. These firms replaced workers in younger, implementation-intensive roles while retaining and rehiring for positions that complement AI capabilities. The reorganization of startup teams was concentrated among technically proficient founders, who adopted AI tools more aggressively and restructured their workforces around higher-skill, more flexible roles.

Displaced workers, however, faced significant frictions. Employees exiting more exposed startups were slower to find reemployment and often transitioned into lower-paying or less AI-exposed occupations, consistent with involuntary displacement rather than voluntary mobility. Despite these labor adjustments, startups became markedly more productive and financially successful. More exposed firms achieved higher sales per employee and achieved more fundraising milestones, suggesting that Gen AI enhanced efficiency.

Venture capital investors also adapted their strategies in response to the AI shock. VCs in Gen AI-intensive sectors made smaller initial investments per firm but expanded the number of deals, broadening their exposure across a larger set of AI-related startups.

This shift toward a more diversified, exploratory investment approach fueled a wave of new startup formation, particularly in markets with higher exposure to Gen AI.

On aggregate, these forces largely offset one another. Employment among incumbent startups declined, but the surge in new firm entry muted the net effect on total startup employment. However, the share of managerial workers increased, indicating a shift in the structure of employment rather than a simple change in its level. These findings suggest that while Gen AI adoption may accelerate business dynamism, it also reshapes how startups scale, how they are financed, and how they structure their workforce.

References

- Acemoglu, Daron, Gary W. Anderson, David N. Beede, Catherine Buffington, Eric E. Childress, Emin Dinlersoz, Lucia S. Foster, Nathan Goldschlag, John Haltiwanger, Zachary Kroff, Pascual Restrepo, and Nikolas Zolas**, *Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey*, University of Chicago Press, May 2024.
- Adelino, Manuel, Song Ma, and David Robinson**, “Firm Age, Investment Opportunities, and Job Creation,” *The Journal of Finance*, 2017, 72 (3), 999–1038.
- Autor, David H and David Dorn**, “The growth of low-skill service jobs and the polarization of the US labor market,” *American economic review*, 2013, 103 (5), 1553–1597.
- Beckman, Christine M. and M. Diane Burton**, “Founding the Future: Path Dependence in the Evolution of Top Management Teams from Founding to IPO,” *Organization Science*, 2008, 19 (1), 3–24.
- Bena, J. and E. Simintzi**, “Machines Could Not Compete with Chinese Labor: Evidence from U.S. Firms’ Innovation,” *Review of Finance*, Forthcoming.
- , **H. Ortiz-Molina, and E. Simintzi**, “Shielding Firm Value: Employment Protection and Process Innovation,” *Journal of Financial Economics*, 2022, 146, 637–664.
- Bessen, James**, “Industry Concentration and Information Technology,” *The Journal of Law and Economics*, 2020, 63 (3), 531–555.
- Brynjolfsson, Erik, BHarat Chandar, and Ruyu chen**, “Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence,” 2025. Working Paper.
- , **Danielle Li, and Lindsey Raymond**, “Generative AI at Work*,” *The Quarterly Journal of Economics*, 02 2025, 140 (2), 889–942.
- Cajner, Tomaz, Leland D. Crane, Ryan A. Decker, Adrian Hamins-Puertolas, and Christopher J. Kurz**, “Using Payroll Processor Microdata to Measure Aggregate Labor Market Activity,” Finance and Economics Discussion Series 2018-005, Board of Governors of the Federal Reserve System (U.S.) 2018.
- Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant’Anna**, “Difference-in-Differences with a Continuous Treatment,” NBER Working Paper 32117, National Bureau of Economic Research Feb 2024.
- Chen, Jun and Micheal Ewens**, “Venture Capital and Startup Agglomeration,” *The Journal of Finance*, 2025, 80 (4), 2153–2198.
- Cohn, Jonathan B., Zack Liu, and Malcolm I. Wardlaw**, “Count (and count-like) data in finance,” *Journal of Financial Economics*, 2022, 146 (2), 529–551.
- Cook, Lisa D, Matt Marx, and Emmanuel Yimfor**, “Funding Black High-Growth Startups,” 2023. *Journal of Finance*, Forthcoming.
- Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda**, “Declining Business Dynamism: What We Know and the Way Forward,” *American Economic Review*, May 2016, 106 (5), 203–07.
- Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda**, “The Role of Entrepreneurship in US Job Creation and Economic Dynamism,” *Journal of Economic*

Perspectives, September 2014, 28 (3), 3–24.

Deshpande, Vijeta, Dan Pechi, Shree Thatte, Vladislav Lialin, and Anna Rumshisky, “Honey, I Shrunk the Language: Language Model Behavior at Reduced Scale,” *arXiv*, 2023.

Dingel, Jonathan I and Brent Neiman, “How many jobs can be done at home?,” *Journal of public economics*, 2020, 189, 104235.

Eisfeldt, Andrea L., Gregor Schubert, Miao Ben Zhang, and Bledi Taska, “Generative AI and Firm Values,” 2023. Working Paper.

Ewens, Michael, Ramana Nanda, and Matthew Rhodes-Kropf, “Cost of experimentation and the evolution of venture capital,” *Journal of Financial Economics*, 2018, 128 (3), 422–442.

Firpo, Sergio, Nicole M Fortin, and Thomas Lemieux, “Occupational tasks and changes in the wage structure,” Technical Report, IZA Discussion Papers 2011.

Ford, Jeffrey D. and John W. Slocum, “Size, Technology, Environment and the Structure of Organizations,” *The Academy of Management Review*, 1977, 2 (4), 561–575.

Gambacorta, Leonardo, Han Qiu, Shuo Shan, and Daniel M. Rees, “Generative AI and Labour Productivity: A Field Experiment on Coding,” *BIS Working Paper No. 1208*, 2024.

García-Trujillo, Gonzalo, Nathalie González-Prieto, and Álvaro Silva, “Startup Employment and Career Trajectories,” Working Paper 978, Banco Central de Chile April 2023.

Geng, H. G., Y. Huang, C. Lin, and S. Liu, “Minimum Wage and Corporate Investment: Evidence from Manufacturing Firms in China,” *Journal of Financial and Quantitative Analysis*, 2022, 57, 94–126.

Gupta, Abhinav, Arpit Gupta, and Elena Simintzi, “Who benefits from Remote Work?,” 2025. Working paper.

—, **Naman Nishesh, and Elena Simintzi**, “Big Data and Bigger Firms: A Labor Market Channel,” *Working Paper*, 2024.

Haltiwanger, John, Ron S. Jarmin, and Javier Miranda, “Who Creates Jobs? Small versus Large versus Young,” *The Review of Economics and Statistics*, 05 2013, 95 (2), 347–361.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence Schmidt, and Bryan Seegmiller, “Artificial Intelligence and the Labor Market,” *Working Paper*, 2025.

Hecker, Daniel E., “High-technology employment: a NAICS-based update,” *Monthly Labor Review*, 2005, 128 (7), 57–72.

Kerr, William R., Josh Lerner, and Antoinette Schoar, “The Consequences of Entrepreneurial Finance: Evidence from Angel Financings,” *The Review of Financial Studies*, 2014, 27 (1), 20–55.

Kim, J. Daniel, Joonkyu Choi, Nathan Goldschlag, and John Haltiwanger, “High-Growth Firms in the United States: Key Trends and New Data Opportunities,” Finance and Economics Discussion Series (FEDS) 2024-074, Board of Governors of the Federal Reserve System 2024.

- Ma, Wenting, Paige Ouimet, and Elena Simintzi**, “Mergers and Acquisitions, Technological Change and Inequality,” 2025. Working paper.
- Ouimet, P., E. Simintzi, and K. Ye**, “The Impact of the Opioid Crisis on Firm Value and Investment,” *The Review of Financial Studies*, 2025.
- Ouimet, Paige and Rebecca Zarutskie**, “Who works for startups? The relation between firm age, employee age, and growth,” *Journal of Financial Economics*, 2014, *112* (3), 386–407.
- Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer**, “The Impact of AI on Developer Productivity: Evidence from GitHub Copilot,” 2023.
- Sauermann, Henry**, “Fire in the belly? Employee motives and innovative performance in start-ups versus established firms,” *Strategic Entrepreneurship Journal*, 2018, *12* (4).
- Schumpeter, Joseph A.**, *Capitalism, Socialism and Democracy*, 1st ed., New York: Harper & Brothers, 1942.
- Sorenson, Olav, Michael S. Dahl, Rodrigo Canales, and M. Diane Burton**, “Do Startup Employees Earn More in the Long Run?,” *Organization Science*, 2021, *32* (3), 587–604.
- Tomlinson, Kiran, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri**, “Working with AI: Measuring the Occupational Implications of Generative AI,” July 2025.
- Tuzel, S. and M. B. Zhang**, “Economic Stimulus at the Expense of Routine-Task Jobs,” *The Journal of Finance*, 2021, *76*, 3347–3399.
- Zolas, Nikolas, Zachary Kroff, Erik Brynjolfsson, Kristina McElheran, David N. Beede, Cathy Buffington, Nathan Goldschlag, Lucia Foster, and Emin Dinlersoz**, “Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey,” 2020.

Figures






Figure 1: The Effects of Gen AI Exposure on Employment


Notes: This figure presents an event study on the employment size of the startup around ChatGPT introduction from Poisson pseudo-maximum likelihood regressions. A unit of observation is a startup-quarter. The dashed line indicates December 2022 to February 2023. The independent variable is the startup's Gen AI exposure, interacted with the time indicator for each quarter. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisefeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. The dependent variable is the startup's employment size. We control for firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, founding-year-by-year-quarter FEs, size-bucket-by-year-quarter FEs, and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Standard errors are clustered by startup.




(a) Share of Senior



(b) Med. Tenure




(c) Avg. Execution Score




(d) Avg. Managerial Score

Figure 2: The Effects of Gen AI Exposure on Team Composition

Notes: This figure presents an event study on the team composition of the startup around ChatGPT introduction from OLS regressions. A unit of observation is a startup-quarter. The dashed line indicates December 2022 to February 2023. The independent variable is the startup's Gen AI exposure interacted with the time indicator for each quarter. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisefeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. We present effects for the share of employees with seniority at the manager level and above, the median tenure, and the average execution score of the startup's employees. The tenure is from the start of the job position to the observation. The execution score is predicted by LLM based on job titles, measuring the extent to which a role involves hands-on implementation tasks such as coding, analysis, design, or operating tools. We control for firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, founding-year-by-year-quarter FEs, size-bucket-by-year-quarter FEs, and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Standard errors are clustered by startup.




(a) by Founder's CS Degree




(b) by Founder's Coding Experience

Figure 3: Heterogeneous Effects of Gen AI Exposure on Employment

Notes: This figure presents an event study on the employment size of the startup by startup characteristics around ChatGPT introduction from Poisson pseudo-maximum likelihood regressions. A unit of observation is a startup-quarter. The dashed line indicates December 2022 to February 2023. The independent variable is a triple interaction of the startup's Gen AI exposure, the indicator for whether the startup founders have computer science degrees or previous experience in computer occupations, and the time indicator for each quarter. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. The dependent variable is the startup's employment size. We control for firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, founding-year-by-year-quarter FEs, size-bucket-by-year-quarter FEs, and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Standard errors are clustered by startup.



(a) Number of New Firms



(b) Number of Firms

Figure 4: The Effects of Gen AI Exposure on New Startup Formation

Notes: This figure presents an event study on new startup formation around ChatGPT introduction from Poisson pseudo-maximum likelihood regressions. A unit of observation is a market-quarter. A market is an industry-state. We use CrunchBase categories as industries. The dashed line indicates December 2022 to February 2023. The independent variable is the market's Gen AI exposure interacted with the time indicator for each quarter. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the occupational composition of the existing startups in the market in the quarter immediately preceding ChatGPT's introduction. We present effects for the number of startups and the number of new startups founded in the quarter of observation. We control for market FEs and growth-bucket-by-year-quarter FEs. Markets are classified into ten growth buckets based on their growth in the number of startups in the quarter immediately preceding ChatGPT's introduction. Standard errors are clustered by market.

Tables

	mean	sd	p10	p50	p90	count
GenAI Exposure	0.39	0.11	0.25	0.41	0.50	94,789
Share of Tech-oriented Occupations	0.16	0.24	0	0	0.50	94,789
Emp. Size	11.00	12.96	1	6	34	1,516,624
Hiring Rate	0.07	0.17	0	0	0.22	1,439,623
Exit Rate	0.06	0.15	0	0	0.17	1,439,623
Payroll (M\$)	1.52	1.61	0.13	0.90	4.40	1,516,624
Share of Senior Employees	0.57	0.30	0.20	0.54	1.00	1,436,325
Med. Tenure (yrs)	2.54	2.12	0.50	2.08	5.00	1,436,325
Avg. Execution Score	0.37	0.21	0.10	0.39	0.64	1,356,692
Avg. Managerial Score	0.66	0.21	0.38	0.66	0.93	1,356,692
Share of Gen AI Job Postings	0.01	0.09	0	0	0	294,127
GitHub Contributions	232.50	513.35	2	21	775	104,864
GitHub Contributions per capita	23.14	64.30	1.40	8.00	52.82	104,864
Sales (M\$)	0.69	2.52	0	0	1.20	209,349
Sales per capita (M\$)	0.19	0.97	0	0	0.33	209,349
Cumulative Amount Raised (M\$)	3.10	29.10	0	0	1.50	1,327,046

Table 1: Summary Statistics

Notes: This table presents descriptive statistics for key variables used in the analysis. We obtain employment, salary, team composition, and employment dynamics from Revelio. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the startup’s occupational composition from Revelio. We follow the definition of tech-oriented occupations of Hecker (2005). We define senior employees as manager-level and above. The tenure is from the start of the job position to the observation. The execution score is predicted by LLM based on job titles, measuring the extent to which a role involves hands-on implementation tasks such as coding, analysis, design, or operating tools. We obtain fundraising data from CrunchBase and PitchBook. We obtain sales data from Data Axle.

Dependent Variable:	Emp. (Poisson)				Hire Rate	Exit Rate
	(1)	(2)	(3)	(4)	(5)	(6)
1(Post) $\times$ GenAI Exposure	-0.0957*** (0.0168)	-0.0749*** (0.0187)	-0.0931*** (0.0183)	-0.0784*** (0.0178)	0.00374 (0.00504)	0.0143*** (0.00451)
Firm FEs	Y	Y	Y	Y	Y	Y
Time FEs	Y					
Industry $\times$ Time FEs		Y	Y	Y	Y	Y
State $\times$ Time FEs		Y	Y	Y	Y	Y
Founding-Year $\times$ Time FEs			Y	Y	Y	Y
Firm Char. $\times$ Time FEs				Y	Y	Y
Y-Mean	10.13	10.13	10.13	10.13	0.112	0.0575
Magnitude (%)	-9.126	-7.217	-8.894	-7.536	3.346	24.81
Observations	1,464,016	1,408,016	1,408,016	1,400,208	610,020	610,020

Table 2: The Effects of Gen AI Exposure on Employment

Notes: This table reports estimates on the impact of Gen AI shock on the employment of the startup from Poisson pseudo-maximum likelihood and OLS regressions. A unit of observation is a startup-quarter. The independent variable is the startup's Gen AI exposure, interacted with an indicator for the post-ChatGPT period. We define the post-ChatGPT period as from December 2022. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. The dependent variable is the startup's employment size, hire rate, and exit rate. The hire rate is the ratio of newly hired employees in a given firm-quarter to the firm's average employment over the current and previous quarter. The exit rate is the ratio of employees leaving the firm in a given firm-quarter to the firm's average employment over the current and previous quarters. Column (1) controls for firm FEs and year-quarter FEs. Column (2) controls for firm FEs, industry-by-year-quarter FEs, and state-by-year-quarter FEs. Column (3) adds controls for founding-year-by-year-quarter FEs. Columns (4)-(6) add controls for size-bucket-by-year-quarter FEs and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Y-Mean is the average of the outcomes before ChatGPT's introduction. Standard errors are clustered by startup. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Panel A: Exit Rate with Non-employment Duration				
	1 Month	3 Months	6 Months	12 Months
	(1)	(2)	(3)	(4)
1(Post) $\times$ GenAI Exposure	0.0103*** (0.00369)	0.00639** (0.00321)	0.00654** (0.00295)	0.00277 (0.00253)
Firm FEs	Y	Y	Y	Y
Industry $\times$ Time FEs	Y	Y	Y	Y
State $\times$ Time FEs	Y	Y	Y	Y
Founding-Year $\times$ Time FEs	Y	Y	Y	Y
Firm Char. $\times$ Time FEs	Y	Y	Y	Y
Y-Mean	0.0575	0.0575	0.0575	0.0575
Magnitude (%)	17.83	11.11	11.38	4.813
Observations	610,020	610,020	610,020	610,020

Table 3: The Effects of Gen AI Exposure on Exit Rates

Notes: This table reports estimates on the impact of Gen AI shock on exit rates from OLS regressions. A unit of observation is a startup-quarter. The independent variable is the startup's Gen AI exposure, interacted with an indicator for the post-ChatGPT period. We define the post-ChatGPT period as from December 2022. We define the post-ChatGPT period as from December 2022. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. Panel A presents effects for exit rates with non-employment durations of 1, 3, 6, and 12 months. Panel B presents effects for exit rates with follow-up positions of lower or higher occupation-level Gen AI exposure and of lower or higher salary. All regressions include firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, founding-year-by-year-quarter FEs, size-bucket-by-year-quarter FEs, and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Y-Mean is the average of the exit rates before ChatGPT's introduction. Standard errors are clustered by startup. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Panel B: Exit Rate with Follow-up Position					
	GenAI Exposure		Salary		As Founder
	Lower	Higher	Lower	Higher	
	(1)	(2)	(3)	(4)	(5)
1(Post) $\times$ GenAI Exposure	0.0211*** (0.00223)	-0.0152*** (0.00171)	0.0123*** (0.00282)	0.00205 (0.00275)	0.000845 (0.000861)
t-Test	0.000***		0.003***		
Firm FEs	Y	Y	Y	Y	Y
Industry $\times$ Time FEs	Y	Y	Y	Y	Y
State $\times$ Time FEs	Y	Y	Y	Y	Y
Founding-Year $\times$ Time FEs	Y	Y	Y	Y	Y
Firm Char. $\times$ Time FEs	Y	Y	Y	Y	Y
Y-Mean	0.0575	0.0575	0.0575	0.0575	0.0575
Magnitude (%)	36.73	-26.35	21.45	3.571	1.468
Observations	610,020	610,020	610,020	610,020	610,020

Table 3: The Effects of Gen AI Exposure on Exit Rates

Notes: This table reports estimates on the impact of Gen AI shock on exit rates from OLS regressions. A unit of observation is a startup-quarter. The independent variable is the startup's Gen AI exposure, interacted with an indicator for the post-ChatGPT period. We define the post-ChatGPT period as from December 2022. We define the post-ChatGPT period as from December 2022. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. Panel A presents effects for exit rates with non-employment durations of 1, 3, 6, and 12 months. Panel B presents effects for exit rates with follow-up positions of lower or higher occupation-level Gen AI exposure and of lower or higher salary. All regressions include firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, founding-year-by-year-quarter FEs, size-bucket-by-year-quarter FEs, and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Y-Mean is the average of the exit rates before ChatGPT's introduction. Standard errors are clustered by startup. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Dependent Variable:	Share of Senior		Med. Tenure		Avg. Execution Score		Avg. Managerial Score	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1(Post) $\times$ GenAI Exposure	0.0178*** (0.00503)	0.0182*** (0.00499)	0.225*** (0.0400)	0.144*** (0.0390)	-0.00784** (0.00361)	-0.00796** (0.00358)	0.0110*** (0.00386)	0.0109*** (0.00383)
Firm FEs	Y	Y	Y	Y	Y	Y	Y	Y
Industry $\times$ Time FEs	Y	Y	Y	Y	Y	Y	Y	Y
State $\times$ Time FEs	Y	Y	Y	Y	Y	Y	Y	Y
Founding-Year $\times$ Time FEs	Y	Y	Y	Y	Y	Y	Y	Y
Firm Char. $\times$ Time FEs		Y		Y		Y		Y
Y-Mean	0.588	0.588	1.664	1.664	0.362	0.362	0.674	0.674
Magnitude (%)	3.022	3.097	13.51	8.644	-2.164	-2.196	1.633	1.621
Observations	1,381,731	1,379,469	1,381,731	1,379,469	1,305,551	1,304,002	1,305,551	1,304,002

Table 4: The Effects of Gen AI Exposure on Team Composition

Notes: This table reports estimates on the impact of Gen AI shock on the team composition of the startup from OLS regressions. A unit of observation is a startup-quarter. The independent variable is the startup's Gen AI exposure interacted with an indicator for the post-ChatGPT period. We define the post-ChatGPT period as from December 2022. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisefeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. We present effects for the share of employees with seniority at the manager level and above, the median tenure, and the average execution or managerial score of the startup's employees. The tenure is from the start of the job position to the observation. The execution or managerial score is predicted by LLM based on job titles following the prompt in Appendix A.1. Columns (1), (3), (5), and (7) control for firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, and founding-year-by-year-quarter FEs. Columns (2), (4), (6), and (8) add controls for size-bucket-by-year-quarter FEs and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Y-Mean is the average of the outcomes before ChatGPT's introduction. Standard errors are clustered by startup. ***p < 0.01, **p < 0.05, *p < 0.1.

Panel A				
Dependent Variable:	Emp. (Poisson)			
Heterogeneity Indicator:	1(Founder's CS Degree)	1(Founder's Coding Exp.)		
	(1)	(2)	(3)	(4)
1(Post) $\times$ GenAI Exposure	-0.0823*** (0.0270)	-0.0528** (0.0263)	-0.0548* (0.0293)	-0.0348 (0.0285)
1(Post) $\times$ GenAI Exposure $\times$ Founder Heterogeneity	-0.201*** (0.0772)	-0.158** (0.0750)	-0.161*** (0.0555)	-0.113** (0.0538)
Firm FEs	Y	Y	Y	Y
Industry $\times$ Time FEs	Y	Y	Y	Y
State $\times$ Time FEs	Y	Y	Y	Y
Founding-Year $\times$ Time FEs	Y	Y	Y	Y
Firm Char. $\times$ Time FEs		Y		Y
Observations	804,192	800,208	804,192	800,208
Panel B				
Heterogeneity Indicator:	1(Serial Entrepreneur)	1(Pre-shock VC Funding)		
	(1)	(2)	(3)	(4)
1(Post) $\times$ GenAI Exposure	-0.0922*** (0.0335)	-0.0652** (0.0324)	-0.0806*** (0.0187)	-0.0730*** (0.0182)
1(Post) $\times$ GenAI Exposure $\times$ Heterogeneity Indicator	-0.00195 (0.0480)	0.0150 (0.0466)	-0.0824 (0.0689)	0.044 (0.0671)
Firm FEs	Y	Y	Y	Y
Industry $\times$ Time FEs	Y	Y	Y	Y
State $\times$ Time FEs	Y	Y	Y	Y
Founding-Year $\times$ Time FEs	Y	Y	Y	Y
Firm Char. $\times$ Time FEs		Y		Y
Observations	804,192	800,208	1,408,016	1,400,208

Table 5: Heterogeneous Effects of Gen AI Exposure on Employment

Notes: This table reports estimates on the heterogeneous effects of Gen AI shock on the employment of the startup by startup characteristics from Poisson pseudo-maximum likelihood regressions. A unit of observation is a startup-quarter. The independent variable is a triple interaction of the firm's Gen AI exposure, an indicator for heterogeneity, and an indicator for the post-ChatGPT period. We define the post-ChatGPT period as from December 2022. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. The dependent variable is the startup's employment size. We present how the effect depends on whether founders have computer science degrees or previous experience in computer occupations, whether founders have previous entrepreneurial experience, and whether the startup received VC funding before ChatGPT's introduction. Columns (1) and (3) control for firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, and founding-year-by-year-quarter FEs. Columns (2) and (4) add controls for size-bucket-by-year-quarter FEs and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Y-Mean is the average of the outcomes before ChatGPT's introduction. Standard errors are clustered by startup. ***p < 0.01, **p < 0.05, *p < 0.1.

Dependent Variable:	Share of Gen AI Job Postings			
	(1)	(2)	(3)	(4)
1(Post) $\times$ GenAI Exposure	0.0453*** (0.00486)	0.0140** (0.00565)	0.0136** (0.00564)	0.0138** (0.00569)
Firm FEs	Y	Y	Y	Y
Time FEs	Y			
Industry $\times$ Time FEs		Y	Y	Y
State $\times$ Time FEs		Y	Y	Y
Founding-Year $\times$ Time FEs			Y	Y
Firm Char. $\times$ Time FEs				Y
Observations	285,297	274,962	274,962	273,615

Table 6: The Effects of Gen AI Exposure on Share of Gen AI Job Postings

Notes: This table reports estimates on the impact of Gen AI shock on the share of Gen AI job postings of the startup from OLS regressions. A unit of observation is a startup-year. The independent variable is the startup's Gen AI exposure, interacted with an indicator for the post-ChatGPT period. We define the post-ChatGPT period as from December 2022. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. Gen AI postings are identified through keywords listed in Appendix A.2 in raw titles and descriptions of job postings. Column (1) controls for firm FEs and year-quarter FEs. Column (2) controls for firm FEs, industry-by-year-quarter FEs, and state-by-year-quarter FEs. Column (3) adds controls for founding-year-by-year-quarter FEs. Column (4) adds controls for size-bucket-by-year-quarter FEs and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Standard errors are clustered by startup. ***p < 0.01, **p < 0.05, *p < 0.1.

Panel A: GitHub Contributions				
Dependent Variable:	Contributions (Poisson)		Contributions per capita (Poisson)	
	(1)	(2)	(3)	(4)
1(Post) $\times$ GenAI Exposure	0.437 (0.282)	0.494* (0.291)	0.492** (0.245)	0.558** (0.247)
Firm FEs	Y	Y	Y	Y
Industry $\times$ Time FEs	Y	Y	Y	Y
State $\times$ Time FEs	Y	Y	Y	Y
Founding-Year $\times$ Time FEs	Y	Y	Y	Y
Firm Char. $\times$ Time FEs		Y		Y
Y-Mean	218.9	218.9	22.17	22.17
Observations	69,009	66,944	69,009	66,944

Panel B: Sales				
Dependent Variable:	1(Sales in Top Quintile)		1(Sales per capita in Top Quintile)	
	(1)	(2)	(3)	(4)
1(Post) $\times$ GenAI Exposure	0.0405*** (0.0130)	0.0435*** (0.0133)	0.0383*** (0.0143)	0.0327** (0.0146)
Firm FEs	Y	Y	Y	Y
Industry $\times$ Time FEs	Y	Y	Y	Y
State $\times$ Time FEs	Y	Y	Y	Y
Founding-Year $\times$ Time FEs	Y	Y	Y	Y
Firm Char. $\times$ Time FEs		Y		Y
Y-Mean	0.177	0.177	0.160	0.160
Magnitude (%)	22.92	24.62	23.88	20.39
Observations	202,462	192,552	202,462	192,552

Table 7: The Effects of Gen AI Exposure on Productivity

Notes: This table reports estimates on the impact of Gen AI shock on the productivity of the startup from Poisson pseudo-maximum likelihood and OLS regressions. A unit of observation is a startup-quarter for Panel A and a startup-year for Panel B. The independent variable is the startup's Gen AI exposure interacted with an indicator for the post-ChatGPT period. We define the post-ChatGPT period as from December 2022. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisefeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. We present effects for the number of contributions to the GitHub repositories owned by the startup and the number of contributions per active contributor in Panel A, and whether the startup's sales or sales per employee are in the top quintile in Panel B. Columns (1) and (3) control for firm FEs, industry-by-time FEs, state-by-time FEs, and founding-year-by-time FEs. Columns (2) and (4) add controls for size-bucket-by-time FEs and growth-bucket-by-time FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Y-Mean is the average of the outcomes before ChatGPT's introduction. Standard errors are clustered by startup. ***p < 0.01, **p < 0.05, *p < 0.1.

Dependent Variable:	1(VC Round)		1(Late-stage VC Round)		Cumulative Raised Amount (Poisson)	
	(1)	(2)	(3)	(4)	(5)	(6)
1(Post) $\times$ GenAI Exposure	0.0117*** (0.00378)	0.0166*** (0.00391)	0.000782 (0.00114)	0.00326*** (0.00119)	0.355 (0.223)	0.406* (0.215)
Firm FEs	Y	Y	Y	Y	Y	Y
Industry $\times$ Time FEs	Y	Y	Y	Y	Y	Y
State $\times$ Time FEs	Y	Y	Y	Y	Y	Y
Founding-Year $\times$ Time FEs	Y	Y	Y	Y	Y	Y
Firm Char. $\times$ Time FEs		Y		Y		Y
Y-Mean	0.0794	0.0794	0.00550	0.00550	2.162	2.162
Magnitude (%)	14.68	20.91	14.23	59.25	42.65	50.09
Observations	1,276,646	1,220,884	1,276,646	1,220,884	1,276,646	1,220,884

Table 8: The Effects of Gen AI Exposure on VC Financing

Notes: This table reports estimates on the impact of Gen AI shock on VC financing of the startup from OLS regressions. A unit of observation is a startup-quarter. The independent variable is the startup's Gen AI exposure, interacted with an indicator for the post-ChatGPT period. We define the post-ChatGPT period as from December 2022. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisefeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. We present effects for whether the startup receives any VC round, whether the startup receives late-stage VC rounds (Series C onward), and the cumulative amount of funding raised. Columns (1), (3), and (5) control for firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, and founding-year-by-year-quarter FEs. Columns (2), (4), and (6) add controls for size-bucket-by-year-quarter FEs and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Y-Mean is the average of the outcomes before ChatGPT's introduction (in millions of dollars for the cumulative amount of funding raised). Standard errors are clustered by startup. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Dependent Variable:	Avg. Initial Funding Size (Poisson)		
	(1)	(2)	(3)
1(Post) $\times$ 1(High Exposure)	-0.114** (0.0526)	-0.122** (0.0522)	-0.122** (0.0512)
VC FEs	Y	Y	Y
Year FEs	Y		
Founding-Year $\times$ Year FEs		Y	Y
Portfolio-Size $\times$ Year FEs			Y
Y-Mean (M\$)	7.908	7.908	7.908
Observations	11,130	11,091	11,091

Table 9: The Effects of Gen AI Exposure on Initial Funding Size

Notes: This table reports estimates on the impact of Gen AI shock on initial funding size from Poisson pseudo-maximum likelihood regressions. A unit of observation is a VC-by-year. The independent variable is the VC's Gen AI exposure interacted with an indicator for the post-ChatGPT period. We define the post-ChatGPT period as from 2023. The Gen AI exposure of a VC is the average exposure of the markets (industry-state pairs) it invested in before the shock, weighted by the funding size. The Gen AI exposure of a market is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the occupational composition of the existing startups in the market in the quarter immediately preceding ChatGPT's introduction. Column (1) controls for VC FEs and year FEs. Column (2) controls for VC FEs and VC founding-year-by-year FEs. Column (3) adds controls for portfolio-size-by-year FEs. The portfolio size is the total number of new investments of a VC before the shock. Y-Mean is the average of the outcomes in million dollars before ChatGPT's introduction. Standard errors are clustered by VC firm. ***p < 0.01, **p < 0.05, *p < 0.1.

Dependent Variable:	Num. of New Investments (Poisson)		Value of New Investments (Poisson)	
	(1)	(2)	(3)	(4)
1(Post) $\times$ 1(High Exposure)	0.0721** (0.0366)	0.0745** (0.0373)	-0.0186 (0.0455)	-0.0177 (0.0467)
VC FEs	Y	Y	Y	Y
Year FEs	Y		Y	
Founding-Year $\times$ Year FEs		Y		Y
Portfolio-Size $\times$ Year FEs		Y		Y
Y-Mean	0.361	0.361	2.139	2.139
Observations	25,216	25,163	25,216	25,163

Table 10: The Effects of Gen AI Exposure on Number and Total Value of Investments

Notes: This table reports estimates on the impact of Gen AI shock on the number and value of new investments from Poisson pseudo-maximum likelihood regressions. A unit of observation is a VC-by-year. The independent variable is the VC's Gen AI exposure interacted with an indicator for the post-ChatGPT period. We define the post-ChatGPT period as from 2023. The Gen AI exposure of a VC is the average exposure of the markets (industry-state pairs) it invested in before the shock, weighted by the funding size. The Gen AI exposure of a market is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the occupational composition of the existing startups in the market in the quarter immediately preceding ChatGPT's introduction. Columns (1) and (3) control for VC FEs and year FEs. Columns (2) and (4) add controls for VC founding-year-by-year and portfolio-size-by-year FEs. The portfolio size is the total number of new investments of a VC before the shock. Y-Mean is the average of the outcomes (in million dollars for the value of investments) before ChatGPT's introduction. Standard errors are clustered by VC firm. ***p < 0.01, **p < 0.05, *p < 0.1.

Panel A: Num. and Total Emp. of Startups								
Dependent Variable:	Num. of New Startups (Poisson)		Num. of Startups (Poisson)		Total Emp. of New Startups (Poisson)		Total Emp. of Startups (Poisson)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1(Post) $\times$ 1(High Exposure)	0.0737*** (0.00314)	0.0612** (0.00298)	0.0160*** (0.00334)	0.00881*** (0.00313)	0.0763** (0.0327)	0.0547* (0.0328)	-0.000496 (0.00729)	-0.00572 (0.00731)
Market FEs	Y	Y	Y	Y	Y	Y	Y	Y
Time FEs	Y		Y		Y		Y	
Market Char. $\times$ Time FEs		Y		Y		Y		Y
Y-Mean	0.743	0.743	3.845	3.845	5.036	5.036	65.49	65.49
Observations	108,920	107,534	108,920	107,534	108,948	107,534	108,948	107,534
Panel B: Composition of Startups								
Dependent Variable:	Avg. Emp. of New Startups (Poisson)		Avg. Emp. of Existing Startups (Poisson)		Avg. Emp. of Startups (Poisson)		Share of Senior Emp.	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1(Post) $\times$ 1(High Exposure)	-0.0592 (0.0373)	-0.0665* (0.0372)	-0.0544*** (0.01147)	-0.0540*** (0.0148)	-0.0474*** (0.0169)	-0.0452*** (0.0168)	0.00566** (0.00265)	0.00571** (0.00264)
Market FEs	Y	Y	Y	Y	Y	Y	Y	Y
Time FEs	Y		Y		Y		Y	
Market Char. $\times$ Time FEs		Y		Y		Y		Y
Y-Mean	6.169	6.169	15.84	15.84	15.81	15.81	0.589	0.589
Observations	31,251	30,833	99,720	99,690	106,986	106,355	108,007	106,790

Table 11: The Effects of Gen AI Exposure on New Startup Formation

Notes: This table reports estimates on the impact of Gen AI shock on new startup formation from Poisson pseudo-maximum likelihood regressions. A unit of observation is a market-quarter. A market is an industry-state. We use CrunchBase categories as industries. The independent variable is the market's Gen AI exposure interacted with an indicator for the post-ChatGPT period. We define the post-ChatGPT period as from December 2022. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the occupational composition of the existing startups in the market in the quarter immediately preceding ChatGPT's introduction. We present the effects on the number of startups and their total employment, as well as the number of new startups founded within one year and their total employment in the quarter of observation in Panel A. We present the effects on the average employment of new, existing, and all startups, as well as the share of employees with seniority at the manager level and above in aggregate startup employment in Panel B. Columns (1), (3), (5), (7) control for market FEs and year-quarter FEs. Columns (2), (4), (6), (8) control for market FEs and growth-bucket-by-year-quarter FEs. Markets are classified into ten growth buckets based on their growth in the number of startups in the quarter immediately preceding ChatGPT's introduction. Y-Mean is the average of the outcomes before ChatGPT's introduction. Standard errors are clustered by market. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Appendix For Online Publication

A	Variable Construction	52
A.1	Managerial and Execution Score	52
A.2	Gen AI keywords	52
B	Appendix Figures	54
C	Appendix Tables	58

A Variable Construction

A.1 Managerial and Execution Score

The managerial and execution score is predicted from raw job titles using the `deepseek-chat` API with the following prompt:

You are an HR classifier. Given a job title, estimate two role characteristics:

1. **Managerial Score:** Reflects how much the job involves setting direction, supervising others, or making strategic or team-level decisions. This includes roles like project managers, product leads, team supervisors, executives, and founders—people who influence others’ work or make key decisions.
2. **Technical/Implementation Score:** Reflects how much the job involves hands-on work—such as coding, designing, analyzing data, operating tools, or directly producing output. This includes roles like engineers, analysts, interns, and technicians who typically execute specific tasks as part of a broader process.

Return only a JSON object with the following keys:

- `"managerial_score"`: float between 0 and 1
- `"technical_score"`: float between 0 and 1

A.2 Gen AI keywords

generative ai, gen ai, gen-ai, genai, large language model, large language models, llm, llms, transformer model, transformer models, prompt engineering, prompt engineer, prompt engineers, rag, retrieval augmented generation, fine-tuning, finetuning, fine tuned, model fine-tuning, chatbot, chatbots, conversational ai, text generation, chatgpt, gpt, gpt-3, gpt-4, gpts, claude, anthropic, llama, llama-2, llama-3, mistral, falcon, gemini, bard, copilot, codex, midjourney, stable diffusion, diffusion model, diffusion models, langchain, llamaindex, vector database, embedding model, embedding models,

openai api, huggingface, transformers library, transformer library, ai agent, ai agents,
autonomous agent, autonomous agents, ai assistant, ai assistants, ai copilot, ai copilots,
ai content generation, ai code generation, ai writing assistant

B Appendix Figures




Figure B1: The Effects of Gen AI Exposure by Firm Age

Notes: This figure plots heterogeneous effects of Gen AI shock on employment by firm age from Poisson pseudo-maximum likelihood regressions. A unit of observation is a firm-quarter. The independent variable is a triple interaction of the firm's Gen AI exposure, an indicator for the age bucket, and an indicator for the post-ChatGPT period. We define the post-ChatGPT period as from December 2022. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisefeldt et al. (2023) with the firm's occupational composition in the quarter immediately preceding ChatGPT's introduction. The dependent variable is the firm's employment size or total salary. The firms are divided into 1-5, 6-10, 11-25, and more than 25 years old in 2022 by row. We also plot the baseline estimate on our main sample of startups as the first row. We control for firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, founding-year-by-year-quarter FEs, and growth-bucket-by-year-quarter FEs. Firms are classified into twenty growth buckets based on their employment growth in the quarter immediately preceding ChatGPT's introduction. Standard errors are clustered by startup.




Figure B2: The Effects of Gen AI Exposure on Employment by Firm Age and Size

Notes: This figure plots heterogeneous effects of Gen AI shock on employment by firm age and size from Poisson pseudo-maximum likelihood regressions. A unit of observation is a firm-quarter. The independent variable is a triple interaction of the firm's Gen AI exposure, an indicator for the age-by-size bucket, and an indicator for the post-ChatGPT period. We define post-ChatGPT period as from December 2022. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the firm's occupational composition in the quarter immediately preceding ChatGPT's introduction. The dependent variable is the firm's employment size. The firms are divided into 1-5, 6-10, 11-25, and more than 25 years old by age, and small (≤ 50 employees), mid-sized (51 – 1000 employees), and large (> 1000 employees) by size in 2022 by row. We control for firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, founding-year-by-year-quarter FEs, and growth-bucket-by-year-quarter FEs. Firms are classified into twenty growth buckets based on their employment growth in the quarter immediately preceding ChatGPT's introduction. Standard errors are clustered by startup.




Figure B3: The Effects of Gen AI Exposure on Employment

Notes: This figure presents an event study on the employment size of the startup around ChatGPT introduction from Poisson pseudo-maximum likelihood regressions. A unit of observation is a startup-quarter. The dashed line indicates December 2022 to February 2023. The independent variable is whether the startup's Gen AI exposure is above median, interacted with the time indicator for each quarter. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. The dependent variable is the startup's employment size. We control for firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, founding-year-by-year-quarter FEs, size-bucket-by-year-quarter FEs, and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Standard errors are clustered by startup.




Figure B4: The Effects of Gen AI Exposure on VC Financing

Notes: This figure presents an event study on VC financing of the startup around ChatGPT introduction from Poisson pseudo-maximum likelihood regressions. A unit of observation is a startup-quarter. The dashed line indicates December 2022 to February 2023. The independent variable is the startup's Gen AI exposure, interacted with the time indicator for each quarter. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. The dependent variable is the cumulative amount of funding raised for the startup. We control for firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, founding-year-by-year-quarter FEs, size-bucket-by-year-quarter FEs, and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Standard errors are clustered by startup.

C Appendix Tables

Step	Number of Firms
CrunchBase firms founded 2018–2021	116,869
Matched to Revelio	87,549
PitchBook firms founded 2018–2021	40,794
Matched to Revelio	28,330
CrunchBase firms matched to Revelio	87,549
+ PitchBook firms matched to remaining Revelio firms	7,240
Final baseline sample	94,789
matched to Data Axle	43,250
matched to GitHub	18,947

Table C1: Sample Construction

Notes: This table presents the number of firms at each step of sample construction.

Panel A: Distribution of Industry Sectors	
Primary Industry Sector	Percent
Information Technology	36.79
Business Products and Services (B2B)	30.73
Consumer Products and Services (B2C)	16.34
Healthcare	11.41
Financial Services	2.92
Energy	0.95
Materials and Resources	0.87
Panel B: Distribution of Top 10 States	
State	Percent
California	21.86
New York	10.97
Texas	8.88
Florida	7.45
Massachusetts	3.61
Illinois	3.33
Georgia	2.76
Colorado	2.68
Pennsylvania	2.42
Washington	2.33

Table C2: Distribution of Industry Sectors and States

Notes: This table presents the distribution of industry sectors (PitchBook) and the top 10 states of the startups in our sample.

Dependent Variable:	Payroll (Poisson)			
	(1)	(2)	(3)	(4)
1(Post) $\times$ GenAI Exposure	-0.124*** (0.0147)	-0.109*** (0.0160)	-0.122*** (0.0156)	-0.0895*** (0.0151)
Firm FEs	Y	Y	Y	Y
Time FEs	Y			
Industry $\times$ Time FEs		Y	Y	Y
State $\times$ Time FEs		Y	Y	Y
Founding-Year $\times$ Time FEs			Y	Y
Firm Char. $\times$ Time FEs				Y
Y-Mean (M\$)	1.414	1.414	1.414	1.414
Observations	1,464,016	1,408,016	1,408,016	1,400,208

Table C3: The Effects of Gen AI Exposure on Payroll

Notes: This table reports estimates on the impact of Gen AI shock on the total payroll of the startup from Poisson pseudo-maximum likelihood regressions. A unit of observation is a startup-quarter. The independent variable is the startup's Gen AI exposure, interacted with an indicator for the post-ChatGPT period. We define the post-ChatGPT period as from December 2022. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. The dependent variable is the startup's total payroll. Column (1) controls for firm FEs and year-quarter FEs. Column (2) controls for firm FEs, industry-by-year-quarter FEs, and state-by-year-quarter FEs. Column (3) adds controls for founding-year-by-year-quarter FEs. Column (4) adds controls for size-bucket-by-year-quarter FEs and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Y-Mean is the average of the outcomes before ChatGPT's introduction in million dollars. Standard errors are clustered by startup. ***p < 0.01, **p < 0.05, *p < 0.1.

Dependent Variable:	Emp. (Poisson)		
	(1)	(2)	(3)
1(Post) $\times$ Exposure	-0.0784*** (0.0178)	-0.154*** (0.0346)	0.172* (0.102)
1(Post) $\times$ ShareSupp		-0.129*** (0.0473)	
1(Post) $\times$ ShareSupp $\times$ Exposure		0.326** (0.130)	
1(Post) $\times$ ShareCore			0.129*** (0.0473)
1(Post) $\times$ ShareCore $\times$ Exposure			-0.326** (0.130)
Firm FEs	Y	Y	Y
Industry $\times$ Time FEs	Y	Y	Y
State $\times$ Time FEs	Y	Y	Y
Founding-Year $\times$ Time FEs	Y	Y	Y
Firm Char. $\times$ Time FEs	Y	Y	Y
Observations	1,400,208	1,398,368	1,398,368

Table C4: Heterogeneous Effects of Gen AI Exposure by Share of Supplemental and Core Tasks

Notes: This table reports estimates on the heterogeneous effects of Gen AI shock on the employment of the startup by share of supplemental and core tasks from Poisson pseudo-maximum likelihood regressions. A unit of observation is a startup-quarter. The independent variable is the startup's Gen AI exposure, interacted with an indicator for the post-ChatGPT period. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. The dependent variable is the startup's employment size. The core and supplemental tasks are classified by O*Net. Column (1) presents the overall effect of exposure from all tasks. Column (2) presents how the effect depends on *ShareSupp*, i.e., the share of exposure deriving from supplemental tasks. Column (3) presents how the effect depends on *ShareCore*, i.e., the share of exposure deriving from core tasks. We control for firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, founding-year-by-year-quarter FEs, size-bucket-by-year-quarter FEs, and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Standard errors are clustered by startup. ***p < 0.01, **p < 0.05, *p < 0.1.

Dependent Variable:	Emp. (Poisson)					
	Drop Tech. Occupations	Drop Tech. Industries	Drop CA/NY	Drop Gen AI Descriptions	Drop VC-backed	Founding Years from 2016
	(1)	(2)	(3)	(4)	(5)	(6)
1(Post) $\times$ Exposure	-0.0837*** (0.0214)	-0.0718*** (0.0187)	-0.0833*** (0.0212)	-0.0589** (0.0230)	-0.0831*** (0.0183)	-0.0643*** (0.0126)
Firm FEs	Y	Y	Y	Y	Y	Y
Industry $\times$ Time FEs	Y	Y	Y	Y	Y	Y
State $\times$ Time FEs	Y	Y	Y	Y	Y	Y
Founding-Year $\times$ Time FEs	Y	Y	Y	Y	Y	Y
Firm Char. $\times$ Time FEs	Y	Y	Y	Y	Y	Y
Observations	1,232,957	1,208,432	942,288	862,362	1,262,672	2,291,088

Table C5: Robustness Tests with Alternative Sample

Notes: This table reports results of several robustness tests with alternative samples. A unit of observation is a startup-quarter. The independent variable is the startup's Gen AI exposure, interacted with an indicator for the post-ChatGPT period. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisfeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. In Column (1), we drop all the technology-oriented occupations (Hecker, 2005) in the sample. In Column (2), we drop startups in the high-technology industries, i.e., those with shares of technology-oriented occupations (Hecker, 2005) above 1.5 times of overall average. In Column (3), we drop startups in California or New York. In Column (4), we only keep startups with descriptions that don't contain Gen AI keywords listed in Appendix A.2. In Column (5), we drop all the startups with VC funding before the shock. In Column (6), we extend the startup sample to include firms founded from 2016 to 2021. All regressions include firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, founding-year-by-year-quarter FEs, size-bucket-by-year-quarter FEs, and growth-bucket-by-year-quarter FEs. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Standard errors are clustered by startup. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Dependent Variable:	Emp. (Poisson)			
	(1)	(2)	(3)	(4)
1(Post) $\times$ GenAI Exposure	-0.0624*** (0.0206)	-0.0608*** (0.0215)	-0.0777*** (0.0179)	-0.0506** (0.0227)
1(Post) $\times$ Remoteness	-0.00985 (0.00765)			-0.00542 (0.00852)
1(Post) $\times$ Offshorability		-0.00869 (0.00630)		-0.00868 (0.00715)
1(Post) $\times$ Share of Tech. Occ.			-0.00302 (0.00834)	-0.00439 (0.00849)
Firm FEs	Y	Y	Y	Y
Industry $\times$ Time FEs	Y	Y	Y	Y
State $\times$ Time FEs	Y	Y	Y	Y
Founding-Year $\times$ Time FEs	Y	Y	Y	Y
Firm Char. $\times$ Time FEs	Y	Y	Y	Y
Observations	1,353,360	1,394,656	1,400,208	1,352,560

Table C6: Robustness Tests with Additional Controls

Notes: This table reports results of several robustness tests with additional controls. A unit of observation is a startup-quarter. The independent variable is the startup's Gen AI exposure, remoteness, offshorability, or the technology-oriented occupations interacted with an indicator for the post-ChatGPT period. The firm-level measures are constructed by combining occupation-level measures with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. We use the occupation-level remoteness measure from Dingel and Neiman (2020), the occupation-level offshorability measure from Firpo et al. (2011) and Autor and Dorn (2013), and the definition of technology-oriented occupations Hecker (2005). We present effects on employment size. All regressions include firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, founding-year-by-year-quarter FEs, size-bucket-by-year-quarter FEs, and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Standard errors are clustered by startup. ***p < 0.01, **p < 0.05, *p < 0.1.

Dependent Variable:	Emp. (Poisson)		
	(1)	(2)	(3)
1(Post) $\times$ 1(High Exposure)	-0.00567* (0.00331)		
1(Post) $\times$ Exposure (Core Tasks)		-0.0865*** (0.0209)	
1(Post) $\times$ Exposure (Microsoft)			-0.237*** (0.0347)
Firm FEs	Y	Y	Y
Industry $\times$ Time FEs	Y	Y	Y
State $\times$ Time FEs	Y	Y	Y
Founding-Year $\times$ Time FEs	Y	Y	Y
Firm Char. $\times$ Time FEs	Y	Y	Y
Observations	1,400,208	1,400,208	1,399,696

Table C7: Robustness Tests with Alternative Exposure Measures

Notes: This table reports results of several robustness tests with alternative exposure measures. A unit of observation is a startup-quarter. We use a binary indicator for whether the Gen AI exposure of the startup is above median in Column (1). We use Gen AI exposure from Eisfeldt et al. (2023) based on all tasks in Column (2) and AI applicability score from Tomlinson et al. (2025) in Column (3). The firm-level measures are constructed by combining occupation-level measures with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. We present effects on employment size. All regressions include firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, founding-year-by-year-quarter FEs, size-bucket-by-year-quarter FEs, and growth-bucket-by-year-quarter FEs. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Standard errors are clustered by startup. ***p < 0.01, **p < 0.05, *p < 0.1.

Dependent Variable:	Log(Emp.)	Log(1+Emp.)	Log(1+Raised Amount)	Log(1+Sales)	Log(1+Sales per Employee)
	(1)	(2)	(3)	(4)	(5)
1(Post) $\times$ GenAI Exposure	-0.114*** (0.0143)	-0.102*** (0.0124)	0.0499*** (0.00751)	0.256** (0.111)	0.233*** (0.0887)
Firm FEs	Y	Y	Y	Y	Y
Industry $\times$ Time FEs	Y	Y	Y	Y	Y
State $\times$ Time FEs	Y	Y	Y	Y	Y
Firm Char. $\times$ Time FEs	Y	Y	Y	Y	Y
Observations	1,379,469	1,400,208	1,220,884	159,108	159,108

Table C8: Robustness Test with Log

Notes: This table reports the robustness of the results with log instead of Poisson regressions. A unit of observation is a startup-quarter. The independent variable is the startup's Gen AI exposure, interacted with an indicator for the post-ChatGPT period. The Gen AI exposure is constructed by combining occupation-level exposure estimates from Eisefeldt et al. (2023) with the startup's occupational composition in the quarter immediately preceding ChatGPT's introduction. We present effects for employment size and the cumulative amount raised by the startup. All regressions include firm FEs, industry-by-year-quarter FEs, state-by-year-quarter FEs, founding-year-by-year-quarter FEs, size-bucket-by-year-quarter FEs, and growth-bucket-by-year-quarter FEs. We use CrunchBase categories as industries. Startups are classified into three size buckets based on their employment in the quarter immediately preceding ChatGPT's introduction, and into two growth buckets based on their employment growth in the same quarter. Standard errors are clustered by startup. ***p < 0.01, **p < 0.05, *p < 0.1.