Born Lean:
Generative AI and the Making of New Firms

Philip G. Berger
Booth School of Business, University of Chicago
philip.berger@chicagobooth.edu

Rebecca N. Hann
Robert H. Smith School of Business, University of Maryland
rhann@umd.edu

Kristen Valentine
Terry College of Business, The University of Georgia
kristen.valentine@uga.edu

Jingyu Zhang
Robert H. Smith School of Business, University of Maryland
jingyuz@umd.edu

This version: July 2026

Born Lean:
Generative AI and the Making of New Firms

Abstract

We examine how the introduction of generative AI reshapes the organizations new firms build. Using the ChatGPT's public release as a shock, we compare venture-backed startups by their GenAI exposure and find the more-exposed ones scale headcount more slowly after the release. They also begin to carry more capital and market value per employee, consistent with capital-labor substitution: the increase comes from lower headcount, not from a contraction in financing. The slowed hiring rate is most pronounced where a firm's ability to substitute GenAI for labor is greatest: among startups building AI products, those led by AI-technical executives, and those already hiring for AI roles before the ChatGPT release. Within firms, the effect of ChatGPT's release on hiring growth falls on early-career, remote-suitable, and specialized workers. The lean scaling response is specific to startup firms, which react to the release with little legacy employment. Newly public firms with established workforces increase hiring rates, even for early-career, remote-suitable, and specialized workers, consistent with GenAI exposure prompting increased labor investment at scaling firms that face labor force adjustment costs. Overall, the results indicate that generative AI is changing how new firms are built, altering their mix of labor and capital rather than only the exposure of individual occupations.

JEL Classification: G24; J23; L26; O33; D24

Keywords: generative AI; firm formation; GenAI exposure; startup hiring; capital-labor substitution; venture capital

"In the coming years, some of the most valuable companies in the world will have 'sub-100 employees…I think the world's not prepared for that… the world economy—and certainly the tech economy—is going to look unrecognizable.'"

-Daniel Nadler, founder of OpenEvidence, as quoted in Fortune

1. Introduction

Generative artificial intelligence (GenAI) is reshaping the nature of work, with shifts in task-level technology use potentially prompting ripple effects on employment, business strategy, and economy-wide growth. Even as GenAI is changing the landscape of existing firms, understanding how startups scale in this new environment offers an important window into the future of work. Established firms face frictions in altering their existing pool of labor given existing policies, systems, and culture and thus, the choices established firms make about labor in the GenAI era are likely path dependent. In contrast, startups in the scaling phase are determining their relative mix of labor and technology from an almost clean slate and may reflect the long-term equilibrium labor-technology composition more rapidly than established firms in this adjustment period. We investigate whether startup firms scale with fewer employees when exposed to the widespread availability of GenAI, using the introduction of ChatGPT in November 2022 as a shock that had rapid impact and was largely unanticipated by employers and employees.

New firms have historically been a principal engine of job growth, accounting for about 20% of U.S. gross job creation (Decker, Haltiwanger, Jarmin, and Miranda 2014). However, if firms that grow natively with GenAI scale with fewer employees, startups may not continue to be as dominant a source of job creation as they have been. At the same time, GenAI-related business applications surged in 2023 (Dinlersoz, Dogan, and Zolas 2024), suggesting that future venture-backed startups will be disproportionately GenAI enabled. Thus, the historical link between startups and job creation may weaken even as GenAI-related startup activity intensifies. Therefore,

understanding how firms scale in the age of GenAI has important implications for both individual employees and the economy at large.

Whether generative AI in fact makes new firms leaner is ex ante ambiguous. On one hand, if GenAI substitutes for labor in exposed tasks, exposed firms can meet a given level of output or growth with fewer people, operating more capital-intensively and scaling with a smaller workforce. This potential is the displacement force emphasized in the task-based view of automation (Acemoglu and Restrepo 2019). On the other hand, AI may instead complement labor by raising the productivity and value of GenAI-exposed work, thereby inducing exposed firms to hire *more* and scale *faster* to exploit opportunities, as the adoption literature on complementary investment suggests (Babina et al. 2024). Which force dominates can be different in the startup setting than for more mature firms. For startups, we predict displacement prevails, as these new firms build their workforces from scratch under tight cash constraints, thus making the opportunity to build leaner highly attractive.

To study GenAI's impact on startup hiring, we construct a quarterly panel of venture-backed startups spanning the period around the ChatGPT release and measure each firm's workforce GenAI exposure from the detailed occupational composition of its employees following Brynjolfsson, Mitchell, and Rock (2018). We classify firms by whether their GenAI exposure is high or low and estimate how the hiring and capital deployment of high-exposure firms change relative to low-exposure firms after November 2022. The sample comprises approximately 6,400 startups from 2020 to 2024.

We first examine whether GenAI-exposed startups scale headcount more slowly after the release of ChatGPT. High-exposure startups reduce their hiring growth rate by about 4 percentage points a year relative to low-exposure peers, a roughly 32% decrease from their pre-release hiring

pace. The hiring slowdown is increasing in firms' GenAI exposure: each standard deviation of GenAI exposure lowers hiring growth by about 2 percentage points. This dose response is consistent with the results being attributable to the technology rather than to broader trends among exposed firms. The inferences are robust to matching high-exposure firms to comparable less-exposed startups on size and industry, and to using finer fixed effects that compare firms in the same city and quarter or the same Series A financing year.

We next examine whether exposed startups scale lean due to capital substituting for labor after the release of ChatGPT. We find that high-exposure startups have about 5 percent more funding per employee and 7 percent more market value per employee than low-exposure peers. As with hiring, the effect is increasing in GenAI exposure. The decrease in hiring growth paired with higher funding or value per employee is most consistent with startup firms scaling lean.

A concern is that these patterns reflect the 2022 funding winter rather than GenAI exposure, since the venture-capital tightening coincided with the release of ChatGPT and could have fallen harder on GenAI-exposed startups. The evidence is inconsistent with this alternative interpretation. Exposed startups raise no less capital after the release than their peers and the slowdown in hiring persists when we control for a firm's pre-shock funding. The effect we observe of scaling with fewer employees is thus more likely a response to the technology, not the funding winter.

To further establish that GenAI exposure prompts the scaling lean effect, we exploit variation in firms' ex ante propensity to adopt GenAI. The reduced hiring rate is concentrated among firms with AI-related products, whose leadership has technical AI ability, and firms that were hiring for AI positions prior to the public release of ChatGPT. We also separate hiring rates by worker types and find that the hiring slowdown is more pronounced for workers whose tasks are most substitutable with GenAI. Specifically, employees aged 25 or younger, workers in

remote-suitable positions, and specialist employees experience greater declines in hiring rates than more senior, onsite and generalist employees. However, the hiring rates are no different based on worker education or gender, pointing to the substitutability of human-performed tasks by GenAI, rather than demographics or credentials of workers, driving the observed change. Together, the firm-level and worker-level evidence are most consistent with substitution as the mechanism.

Finally, we show that the leaner configuration is specific to startup firms in the scaling phase. We compare hiring growth rates between startups and newly public firms in similar sectors and size brackets. The analysis compares firms in similar industries and size to exploit variation in firms' life cycle stages – specifically, startups in the scaling phase relative to newly public firms with more established workforces. We find that for startup firms, hiring rates decrease in general and for entry-level workers, remote workers and specialists in particular, while the exact opposite is true for newly public firms. The results show that greater GenAI exposure predicts faster hiring growth among newly public firms. Newly public firms have an established workforce and thus may be more prone to expand their workforce to facilitate augmentation or productivity gains from GenAI rather than using GenAI to substitute for labor and incur the adjustment and reputation costs from firing existing employees. We find that the observed increase in hiring rates for newly public firms is bigger in cash-rich firms and extends to investments in capital expenditures, though the effect is statistically indistinguishable from that of cash-lean firms. Taken together, the results are consistent with startup firms relying more on capital relative to labor when they can use GenAI to scale, whereas newly public firms use their post-IPO enhanced financing capacity to hire faster, perhaps to gain GenAI capabilities to augment their existing workforce.

Our paper makes three contributions. First, we contribute to the literature on artificial intelligence and labor, which measures the exposure of occupations and tasks to AI (Brynjolfsson

and Mitchell 2017; Felten, Raj, and Seamans 2021; Eloundou et al. 2024; Pizzinelli et al. 2023) and traces its consequences for employment, wages, and job postings (Acemoglu and Restrepo 2019; Hampole, Papanikolaou, Schmidt, and Seegmiller 2025; Teutloff et al. 2025; Humlum and Vestergaard 2025; Johnston and Makridis 2026; Berger et al. 2026). This evidence is generally established at the level of the worker or occupation and is often identified from firms' realized AI adoption. We move the analysis to the firm, where hiring decisions are actually made, and identify the effect from predetermined workforce exposure interacted with a single common shock rather than from endogenous adoption. The work closest to ours documents that the economy-wide effect of generative AI falls hardest on young, entry-level, and junior workers (Brynjolfsson, Chandar, and Chen 2025; Hosseini Maasoum and Lichtinger 2025). We differ in two ways: we tie the slowdown in hiring to the firm's capacity to substitute GenAI for labor, and we study a setting that cleanly isolates the hiring margin. Startup firms build their workforces from scratch, so the reduced labor growth we document is a decision not to hire uncontaminated by the retention and restructuring frictions in the presence of an existing workforce.

Second, we add to the literature on technology, factor demand, and firm value, which treats labor as a central, partially intangible input whose configuration shapes firm value (Becker 1962; Eisfeldt and Papanikolaou 2013; Belo, Lin, and Bazdresch 2014) and studies how readily that human capital can be substituted or reconfigured (Jaravel, Petkova, and Bell 2018; Jäger and Heining 2022; Custódio, Ferreira, and Matos 2013; Gathmann and Schönberg 2010; Becker and Murphy 1992). We show that a generative-AI shock shifts the labor-capital configuration toward capital: exposed startups grow with a smaller workforce, as corroborated by more capital and market value per employee. This provides a firm-level micro-foundation for the macroeconomic evidence that AI raises capital intensity and lowers the labor share (Acemoglu and Restrepo 2019;

Johnston and Makridis 2026), and distinguishes us from work on GenAI adoption in established firms, whose factor proportions are less variable (Babina et al. 2024; Eisfeldt et al. 2026).

Third, we extend research on entrepreneurship, firm scaling, and business dynamism, which studies how new firms raise capital and grow (Gompers 1995; Hellmann and Puri 2002; Davila, Foster, and Gupta 2003; Ewens, Nanda, and Rhodes-Kropf 2018; Janeway, Nanda, and Rhodes-Kropf 2021) and how scaling firms build their workforces (DeSantola and Gulati 2017; Core, Ryduchowska, and Wang 2024; Lee and Kim 2024; Ouimet and Zarutskie 2014; Kim 2018; Sorenson, Dahl, Canales, and Burton 2021; Choi, Goldschlag, Haltiwanger, and Kim 2023; Bessen, Poege, and Röttger 2024). Yet we know little about how a technology shift reshapes the organization that startup firms build. We find that generative AI changes the blueprint of the firm being built. Exposed startups scale with fewer employees and greater capital intensity, resulting in a different mix of labor and capital, not just fewer employees. The study closest to ours examines how generative AI affects the performance of entrepreneurs (Otis et al. 2024); we differ in studying the organizations firms build in the GenAI era, with implications for workforce composition (Decker, Haltiwanger, Jarmin, and Miranda 2014; Karahan, Pugsley, and Şahin 2024).

2. Institutional Background and Hypothesis Development

2.1 Scaling-Stage Venture-Backed Startups

The firms we study are venture-backed startups in their scaling stage. Such firms raise capital in a sequence of priced rounds, seed and then Series A, B, C, and beyond, each financing a phase of development and staged on the milestones reached in the previous one (Gompers 1995). The earliest rounds fund the search for product-market fit by a small founding team. The Series A round marks the onset of scaling: having established that fit, the firm raises its first substantial institutional round and begins building the organization to grow a proven model (Hellmann and

Puri 2002). Series B and C rounds finance continued growth, until the firm matures, exits through acquisition or a public offering, or stalls.

The defining activity of the scaling stage is hiring, the process through which a venture organizes itself and grows (DeSantola and Gulati 2017; Lee and Kim 2024). After its first institutional round, a startup adds employees steadily, building out its organization from the bottom up by expanding junior and support roles (Core, Ryduchowska, and Wang 2024). The hiring is also a startup's largest financial commitment, as payroll is the dominant use of its funds and the majority of its expenditures (Davila, Foster, and Gupta 2003; Ewens, Nanda, and Rhodes-Kropf 2018), so the central operating decision a scaling startup makes is how many workers to add and in which roles, the margin of job creation (Mortensen and Pissarides 1994).

Two conditions shape this workforce. First, the startup assembles its workforce from a small base, choosing each hire to fit the technology of production it operates as it grows. Second, startups fund hiring from raised capital under binding cash constraints, lacking the debt capacity and retained earnings that let established firms buffer their payrolls (Robb and Robinson 2014), and startups compete for scarce talent with larger employers drawing on the same labor pool (Bessen, Poege, and Röttger 2024). The hiring decision is thus tightly coupled with the firm's production technology and to the financing it has on hand.

These features stand in contrast to established firms, which carry a full workforce built under an earlier technology and embedded in fixed roles and reporting lines. An incumbent that desires a leaner configuration could achieve it only by restructuring its organization, against the labor adjustment costs, internal labor markets, and organizational inertia that make incumbent structures slow to change (Hamermesh and Pfann 1996; Hannan and Freeman 1984). Thus, the marginal cost of adopting a leaner organization is low for the startup firm and high for incumbents.

2.2 Generative AI and Workforce Exposure

The release of ChatGPT on November 30, 2022 made generative AI broadly usable for the first time, putting a general-purpose tool for open-ended cognitive work in the hands of non-specialists (Eloundou et al. 2024). This was a step-change from earlier machine learning, which performed narrow tasks on structured data. Large language models instead produce contextually relevant text, code, and analysis, and so can perform or support the writing, programming, customer service, research, and other knowledge work that had previously required employees. Early causal evidence confirms that ChatGPT's release shifted labor demand, with demand for the most substitutable online freelance tasks falling sharply afterward (Teutloff et al. 2025) and GenAI-exposed employers cutting back hiring (Berger et al. 2026). By lowering the amount labor needed to carry out these tasks, the technology can act as a shock to the relative price of labor and capital in producing a firm's output.

How much this shock changes a given startup's input choice depends on how readily the technology substitutes for the tasks its particular workforce performs. Generative AI is more capable in some occupations than others, so startups assembled with varying occupational mixes differ in workforce exposure. The occupational exposure is a feature of the tasks a firm's employees do, and it reflects the technology's potential to substitute for that work rather than its realized adoption, which is partial and uneven (Humlum and Vestergaard 2025). We exploit the variation in exposure across otherwise-similar startups in our hypotheses and tests.

2.3 Hypothesis Development

2.3.1 Hiring

We view generative AI as a shock to the relative price of labor and develop its consequences through the task-based view of automation (Acemoglu and Restrepo 2019; Acemoglu 2025). The

technology exerts three forces on a startup's demand for labor. The first is displacement: AI substitutes for workers in the tasks it can perform, so a more-exposed startup can meet a given level of output or growth with fewer people, a substitution margin that task-level estimates find to be large (Hampole, Papanikolaou, Schmidt, and Seegmiller 2025). The second is a productivity, or scale, effect: by lowering the cost of producing output, AI can let an exposed startup grow faster and pursue a larger market, raising its demand for labor in the tasks that are not automated, the pattern emphasized in evidence that AI adoption accompanies firm growth (Babina et al. 2024) and that AI-exposed sectors have expanded output and employment (Johnston and Makridis 2026). The third is augmentation: AI creates new tasks in which labor retains an advantage, from overseeing the technology to building it, adding demand for AI-complementary workers. The first force lowers headcount, while the second and third raise it. Which of these forces dominates therefore varies with a startup's workforce exposure, with displacement weighing more heavily the more of its work the technology can perform.

For the scaling startups we study, we expect displacement to dominate. Startups assemble their workforces fresh and under cash constraints, so an increase in the relative price of labor prompts building leaner from the outset rather than adjusting staffing around the new capability. A startup whose workforce is more exposed to GenAI should therefore require fewer workers to scale after the release of ChatGPT and add headcount more slowly than an otherwise-comparable, less-exposed startup.

> **H1.** After the release of ChatGPT, startups with greater GenAI exposure scale headcount more slowly than less-exposed startups.

2.3.2 Capital-Labor Substitution

How a startup combines labor and capital is central to venture-backed growth, because capital comes in discrete rounds that the firm spends down as it builds its organization (Gompers 1995; Davila, Foster, and Gupta 2003), and the configuration of this labor input shapes firm value (Eisfeldt and Papanikolaou 2013; Belo, Lin, and Bazdresch 2014). If generative AI substitutes for labor in exposed tasks, a more-exposed startup can reach a given scale with a smaller workforce, shifting its mix of labor and capital toward capital.

Overall, we expect the substitution toward capital to dominate, since we predict in the previous hypothesis that the rate of headcount increase will slow for the growing young private firms that are more exposed to GenAI. We therefore predict that, following the release of ChatGPT, more-exposed startups substitute capital for labor, operating more capital-intensively than less-exposed startups.

> **H2.** After the release of ChatGPT, more-exposed startups substitute capital for labor, operating with a more capital-intensive mix of inputs than less-exposed startups.

However, this hypothesis is not without tension. Concurrent work suggests that GenAI adoption could increase hiring, perhaps even more than the attendant increased use of capital, thereby prompting a less capital-intensive mix. For example, Kharazian, Simon, and Stevens (2026) combine observed GenAI spending with employment records to study how GenAI adoption (rather than exposure) affects firm-level employment. They conclude that GenAI adoption is associated with higher employment, especially among high-intensity adopters, with firms that make sustained and intensive GenAI investments expanding employment faster than "otherwise

comparable firms."[1] If this expansion in labor exceeds capital use from GenAI adoption, the mix of labor and capital could shift to be more labor intensive.

3. Data, Sample, and Research Design

3.1 Data Sources

We combine four datasets to study how venture-backed startups' hiring and capital intensity respond to the generative-AI shock according to their predetermined exposure.

Revelio Labs. Our workforce data come from Revelio Labs, which compiles individual-level information on workers' careers and characteristics from public professional profiles. The career records give each worker's sequence of positions, with the employer, job title, start and end dates, an occupation code (the O*NET-SOC classification), and a seniority level on a seven-point scale. The characteristic records give demographic attributes such as gender and ethnicity, along with educational background. We restrict to full-time United States positions, dropping internships and temporary, seasonal, part-time, and contract roles. We then build a firm-quarter panel, counting a worker as employed at a firm in a quarter when a position at that firm is active at quarter-end, and measure quarter-end headcount, new hires, and the workforce's composition by seniority, age, education, and occupation.[2]

[1] We note that their approach is quite different from ours as their definition of otherwise comparable firms is a firm that eventually adopts GenAI within the same intensity tercile as the focal firm, but that has not yet adopted by the relevant focal firm's adoption time. Moreover, they examine whether employment expands, not whether the *rate* of expansion is increasing or declining. In a Financial Times article (Murray 2026) accompanying the release of the Kharazian et al. study, an anonymous labor economist cautions that "the groups using AI most in the sample tend to be smaller" and that thus it is hard for the results to separate between the paper's explanation of "intense AI adopters grow faster" versus an alternative explanation that "small fast-growing start-ups buy a lot of AI quite early."

[2] Revelio builds worker-level employment histories from public professional profiles, so coverage varies across workers, occupations, and geographies, tilting toward higher-skill roles. Such measures are widely used to study firm-level employment (Benson and Shaw 2025; Kharazian, Simon, and Stevens 2026) and track external benchmarks: Li, Lourie, Nekrasov, and Shevlin (2022) find that Revelio-based turnover correlates with Bureau of Labor Statistics turnover. In our sample, our raw quarter-end headcount is strongly rank-correlated with the employee count PitchBook reports independently (Spearman rank correlation 0.69, untabulated); the level gap reflects that we count United States full-time workers, a subset of PitchBook's total. Such coverage does not threaten identification: with firm and sector-

PitchBook. Our venture-financing and firm-characteristic data come from PitchBook, which records both each firm's financing history and its company characteristics. The deal records give the date, stage, and amount of every financing round, together with its pre- and post-money valuation. The company records give the founding year, primary industry sector, and ownership status. We link PitchBook firms to their Revelio records through a firm-level crosswalk.[3]

Other data sources. Compustat provides public firms' financial characteristics, including their employee counts and IPO dates, which we link to Revelio by issuer CUSIP and use to benchmark the startup response against comparable public firms. LightCast provides broad coverage of online job postings, with the posting date, the hiring firm, and the full posting text.[4] We identify artificial-intelligence postings using the classifier of Shi et al. (2025), which labels a posting as AI-related from its text, and use them as a hiring-based measure of a firm's engagement with AI.

3.2 Startup Sample

We construct a firm-quarter panel that follows each venture-backed startup across its growth trajectory. Beginning from PitchBook, we take firms that raised a Series A round and link them to their Revelio employment records. A firm enters the panel in the quarter of its Series A financing, the early growth-stage round at which it builds out its core team, and remains until two years after

by-quarter fixed effects, coverage stable within a firm or common to a sector-quarter is absorbed, so identification requires only that residual, time-varying coverage not correlate with GenAI exposure, consistent with Hosseini Maasoum and Lichtinger (2025), who find the composition of a comparable Revelio panel stable across the release of ChatGPT.

[3] PitchBook is a leading commercial database of private capital markets and a standard source for research on venture-backed firms, described as offering the most comprehensive venture universe available (Howell, Parker, and Xu 2024). It covers several million companies and financing events, assembled through automated collection and a research team that verifies records directly with market participants (PitchBook 2024).

[4] LightCast (formerly Burning Glass Technologies) aggregates online job postings scraped from roughly 40,000 company websites and job boards, covering close to the universe of vacancies posted online in the United States, with an occupational and industry composition that tracks BLS Job Openings and Labor Turnover Survey vacancies and Occupational Employment Statistics employment shares reasonably closely (Hershbein and Kahn 2018; Acemoglu, Autor, Hazell, and Restrepo 2022).

its Series C round, or until the end of the sample period if it never raises a Series C, so that the window spans the firm's active scaling years.

From this panel we draw the estimation sample by imposing three restrictions. First, we retain quarters from 2020 through 2024 and require a firm's Series A to have closed before 2022; this keeps firms' GenAI exposure predetermined, measured from the pre-ChatGPT workforce rather than as a response to the technology, and leaves a long enough window after Series A to observe the firm through its scaling phase. Second, we require quarter-end headcount of at least five, both contemporaneously and four quarters earlier, to limit noise in the growth rate from very small firms. Third, we drop firm-quarters with missing GenAI exposure or primary sector. The resulting sample comprises approximately 6,400 startups, young and small (median quarter-end headcount about 30 employees), spanning the seven PitchBook sectors and concentrated in information technology and healthcare.

3.3 Measuring GenAI Exposure

The firm's GenAI exposure (*GenAIExposure*) is its predetermined exposure to artificial intelligence through the composition of its workforce. We measure GenAI exposure at the occupation level using the suitability-for-machine-learning index (SML) of Brynjolfsson, Mitchell, and Rock (2018), which scores each occupation by how amenable its tasks are to machine learning, the technology underlying generative AI.[5] We map the index to each worker's job via a standard occupational crosswalk and aggregate to an employee-weighted firm average, so that the

[5] Acemoglu, Autor, Hazell, and Restrepo (2022) find SML less predictive of firm AI adoption in their pre-generative-AI setting than the measures of Felten, Raj, and Seamans (2021) and Webb (2020), because SML is more forward-looking, scoring tasks AI could perform in the near term rather than those within reach at the time. That property is what makes this measure appropriate for a generative-AI shock: the release of ChatGPT realizes exactly those near-term language and cognition tasks. Our results are robust to alternative exposure measures, including the language-model occupational exposure from Felten, Raj, and Seamans (2023), which Acemoglu, Autor, Hazell, and Restrepo (2022) find performs best in their setting (Table 8).

measure is increasing the more the tasks a firm's employees undertake are exposed to GenAI. We measure GenAI exposure at a single point in time based on the workforce in the firm's Series A year, so that exposure is set well before the release of ChatGPT and is neither contaminated by anticipation nor itself an outcome of the technology. We assess robustness to alternative exposure measures, including the language-model occupational-exposure index of Felten, Raj, and Seamans (2023), the generative-AI task-exposure ratings of Eloundou et al. (2024), and the Anthropic Economic Index of Handa et al. (2025) as further discussed in Section 4.5.

3.4 Research Design

3.4.1 Identification Strategy

We identify the effect of GenAI from the November 30, 2022 public release of ChatGPT, a discrete and largely unanticipated advance in generative-AI capability. The release was common to all firms, so on its own it does not identify a causal effect; identification comes instead from firms' exposure to the technology ex ante. ChatGPT raised the value of automating the tasks captured in our GenAI exposure measure, raising the relative price of labor for firms whose workforces are more exposed to GenAI. We thus compare the change in outcomes after the release of ChatGPT across firms that differ in predetermined exposure, treating more-exposed firms as treated and less-exposed firms as controls, in a differential-exposure difference-in-differences design.

The identifying assumption is parallel trends: absent the release of ChatGPT, the outcomes of high- and low-exposure firms would have evolved in parallel. The design further requires that exposure is predetermined with respect to the release and affects post-release outcomes only through a firm's responsiveness to the AI shock, not through correlation with other firm characteristics that affect outcomes after 2022. Two threats to identification deserve attention. The first is that exposure is correlated with differential trends unrelated to AI, most plausibly those

accompanying the concurrent 2022 venture funding contraction. A contraction in venture financing began in 2022, after the boom of 2021. This swing, although unusually severe, was characteristic of the venture cycle and its effects on startup financing (Janeway, Nanda, and Rhodes-Kropf 2021). The broad contraction reshaped the funding environment facing every venture-backed startup during the period we study.

The second threat to identification is that exposure is correlated with other firm attributes, such as size, valuation, or capitalization, that independently shape the post-release response. For example, Eisfeldt et al. (2026) and Berger et al. (2026) find that publicly traded firms with high GenAI exposure experienced positive abnormal stock returns in the short window around ChatGPT's release. The design addresses the first threat to identification directly, by comparing exposed and unexposed firms within the same sector and quarter and by examining pre-release trends relative to the quarter before the release (2022Q3); we probe both concerns further in Section 4. Even low-exposure firms are AI-exposed to a positive, if smaller, degree, so our estimate is a relative effect and therefore a conservative lower bound on the absolute effect of the technology.

3.4.2 Estimation

Our baseline specification is

$$y_{it} = \beta\,(\text{High GenAIExposure}_i \times \text{Post}_t) + \theta \log(\text{Headcount}_{i,t-4}) + \alpha_i + \delta_{st} + \varepsilon_{it},$$

where y_{it} is the outcome for firm i in quarter t, namely hiring growth in our test of leaner hiring (H1) and the per-employee funding ratios in our test of capital-labor substitution (H2). *High GenAIExposure*$_i$ is an indicator equal to one for firms in the top tercile of *GenAIExposure* (the treatment group) and zero for those in the bottom tercile (the control group), with the middle tercile dropped. We also report a continuous specification that replaces the indicator measure with

standardized *GenAIExposure* (per one standard deviation) as an intensity complement, in the spirit of the continuous difference-in-differences design of Callaway, Goodman-Bacon, and Sant'Anna (2024). $Post_t$ equals one for quarters from the fourth quarter of 2022 onward, and $log(Headcount_{i,t-4})$ is the log of the firm's quarter-end headcount four quarters earlier (*Log Lagged Headcount*), which controls for the firm's initial size, since larger firms tend to grow more slowly (mean reversion).[6] The firm fixed effect α_i absorbs all time-invariant differences across firms, including the level of GenAI exposure. The sector-by-quarter fixed effect δ_{st} absorbs the contemporaneous macroeconomic environment common to a sector, including the 2022 venture funding contraction. The coefficient of interest, β, is the differential change in the outcome for high-GenAI-exposure firms relative to low-GenAI-exposure firms after the release of ChatGPT. We cluster standard errors by firm.

4. Empirical Results

4.1 Descriptive Statistics

Table 1, Panel A reports summary statistics for the variables used in the analysis, separately for the two estimation samples that enter the firm-quarter panel of venture-backed startups over 2020 through 2024. The binary sample, which enters the baseline tercile specifications, retains the top and bottom *GenAIExposure* terciles (4,252 firms and 56,745 firm-quarters) and drops the middle third; the continuous sample, which enters the intensity specifications, is the full estimation sample (6,377 firms and 85,666 firm-quarters) and contains the binary sample as a subset. The startups are small and young: median quarter-end headcount is about 30 employees and median Series A

[6] We also present our results without any controls and find the results are robust (Table 2 columns 1 and 5). Nickell (1981) recommends against including both the lagged dependent variable and firm fixed effects in the same model to avoid bias. We note that the Nickell bias is primarily concerned with the coefficient on the lagged term, not our coefficient of interest. Furthermore, the bias is decreasing in the number of time periods in the panel, which is 14 (quarters) at the median in our sample. Thus, any bias is unlikely to significantly impact our primary inferences.

capital is about 11 million dollars, albeit with a long upper tail of larger, better-funded firms. The startups deploy substantial capital per worker, a median of roughly 1.1 million dollars of cumulative funding per employee. *Hiring Growth*, the twelve-month growth rate of quarter-end headcount, averages -0.035 in the binary sample and -0.044 in the continuous sample with a median of zero, reflecting the broad scaling slowdown over the sample period. These are scaling-stage ventures that are still being built, not mature firms adapting already existing organizations.

Table 1, Panel B reports the industry composition of the binary sample at the firm level. The startups concentrate in information technology and healthcare, which together account for about 70 percent of firms, with the remainder spread across consumer and business products, financial services, materials, and energy. Table 1, Panel C reports Pearson correlations below the diagonal and Spearman rank correlations above, with the two broadly consistent. Hiring growth is mildly negatively correlated with firm size (Pearson -0.14 with *log(lagged headcount)*), the mean reversion that the lagged-size control absorbs. *GenAIExposure* is only weakly correlated with the hiring and funding or valuation per employee outcomes (Pearson -0.04 with *Hiring Growth*), consistent with high- and low-exposure startups being comparable on the outcomes before the ChatGPT release.

4.2 Leaner Hiring and Capital-Labor Substitution

4.2.1 GenAI Exposure and Hiring Growth

We begin by testing whether generative AI changes the pace of hiring at exposed startups. Generative AI substitutes for human labor in the tasks our exposure measure captures, lowering the value of hiring workers for those tasks. After the release of ChatGPT, more-exposed startups should therefore add workers at a lower rate than otherwise-comparable, less-exposed startups, with no differential trend beforehand.

The outcome is *Hiring Growth*, the symmetric growth rate of quarter-end headcount over four quarters: the change in headcount between a quarter and the same quarter one year earlier, divided by the average of the two (Davis and Haltiwanger 1992). The measure is bounded between -2 and +2 and treats entry and exit symmetrically, which makes it well suited to young firms whose lagged headcount can be small, where an ordinary percentage change would be volatile or undefined. We measure growth over a twelve-month horizon.[7] Our primary specifications control for the log of lagged headcount, taken four quarters back to match that window, since larger firms tend to grow more slowly (mean reversion).[8]

Table 2, Panel A reports the binary design, in which *High GenAIExposure* × *Post* compares the top and bottom *GenAIExposure* terciles before and after the release of ChatGPT. Column 1 omits the lagged-headcount control and includes firm and sector-by-quarter fixed effects: high-exposure startups reduce the hiring growth rate by 3.4 percentage points (t = -2.07) relative to low-exposure peers. Column 2 adds the headcount control and the results remain statistically significant with a point estimate of 4.1 percentage points (t = -2.76). The effect is sizable relative to the pre-ChatGPT hiring pace that it slows: with these startups hiring at roughly 13 percent a year before the release, a 4.1-point reduction slows their hiring pace by 32%.

Table 2 column 3 addresses the leading concurrent event confound: the 2022 venture funding contraction that coincided with the release of ChatGPT. High-exposure startups are somewhat smaller and less funded at baseline (Series A capital of 14.8 versus 20.5 million dollars; untabulated), so the contraction could have slowed their hiring through thinner capitalization rather

[7] The nine-month horizon yields the same pattern: a reduction of 0.032 in the binary specification (t = -2.54) and of 0.015 per standard deviation of exposure in the continuous design (t = -3.00), both with clean pre-trends (joint p = 0.28 and 0.31).

[8] As a robustness check, we also use the change in log headcount over the same horizon, defined as the log of quarter-end headcount minus its log four quarters earlier. This conventional log growth rate is unbounded and so provides a check that the result does not depend on the bounded functional form of the symmetric measure.

than through AI. We therefore add the firm's initial capitalization interacted with the post period, *Log(Series-A funding)* × *Post*, to absorb any post-shock trend running through pre-shock funding. The estimate remains similar, but marginally larger in magnitude (4.9 points, t = -3.28), so the decreased hiring growth does not reflect the funding contraction. Column 4 ensures that the comparison is between comparable firms, matching high- to low-exposure startups on pre-shock size (their mean log headcount before the release) and industry by propensity score, with the effect essentially unchanged (4.1 points). Columns 5 through 8 repeat these four specifications with the change in log headcount as the dependent variable, and the effect is robust throughout.

Table 2, Panel B reports the continuous design, which estimates the response across the full exposure gradient rather than comparing the top and bottom terciles alone. The specification replaces the tercile indicator with standardized exposure, *GenAIExposure (1 sd)* × *Post*, so the coefficient measures the change in hiring growth per standard deviation of exposure. Column 1 omits the lagged-headcount control (1.6 points per standard deviation, t = -2.53); column 2 adds it, the baseline, and each standard deviation of exposure lowers the hiring growth rate by 1.9 percentage points (t = -3.24). The estimate is robust to the same funding control (column 3), to entropy balancing on pre-shock size and industry (column 4), and to using the log-headcount outcome (columns 5 through 8).[9] Collectively, the continuous estimates corroborate the binary ones, with a sharper hiring rate slowdown the more exposed a firm's workforce is to AI.

Figure 1 plots the event-study estimates, the exposure effect on hiring growth in each quarter relative to the ChatGPT release, using the 3rd quarter of 2022 as the reference quarter, for the binary design (Panel A) and the continuous design (Panel B). Before the release of ChatGPT,

[9] The two designs balance covariates by different methods, each suited to its treatment. Panel A uses propensity-score matching, which is built for the binary treated-versus-control contrast of the tercile design, whereas the continuous treatment in Panel B has no such split, so Panel B uses entropy balancing, which reweights observations to balance the same covariates, pre-shock size and industry, across the full exposure gradient.

the two groups are on the same hiring trajectory: the pre-release coefficients are flat and jointly insignificant, and high- and low-exposure startups had statistically indistinguishable hiring growth in the pre-period. The effect emerges only in the quarters after the 4th quarter of 2022 and then persists through the end of the sample period, albeit with the top of the 95% confidence intervals getting to just above zero in the 3rd and 4th quarters of 2024.[10]

4.2.2 Capital and Value per Employee

Beyond hiring, we examine whether exposed startups appear to substitute capital for labor. Generative AI performs work that would otherwise require employees, letting exposed startups reach a given scale with a smaller workforce. More-exposed startups should therefore carry higher funding and market value per employee than otherwise-comparable less-exposed startups.

Table 3, Panel A reports the results from the binary design. After the release of ChatGPT, high-exposure startups carry 0.049 log points more *Funding per Employee* than otherwise-comparable low-exposure peers (column 1, t = 2.03), so each worker is backed by roughly 5 percent more capital. Their *Valuation per Employee* is 0.070 log points higher (column 2, t = 2.51), about 7 percent more market value per worker. Columns 3 and 4 scale by net-new hires, the headcount a firm has added since its Series A round. This captures the workforce the firm used to scale this funding round rather than its founding team. If the introduction of ChatGPT prompted firms to scale with a leaner workforce, we expect the ratios of capital and value per net new hire to be higher since series A funding. We indeed find both ratios rise using this alternative scalar.[11]

[10] The slower hiring accumulates rather than reversing: in a level difference-in-differences with log headcount as the outcome (firm and sector-by-quarter fixed effects), high-exposure startups carry about 5 percent fewer employees than comparable low-exposure peers by the end of the sample period (untabulated), an economically meaningful gap that mirrors the roughly one-third slowdown in their pre-release hiring pace.

[11] The results are essentially unchanged controlling for the firm's time since Series A, interacted with Post. High- and low-exposure firms are in any case similar in Series-A timing, with a median Series-A year of 2018 in both groups and a firm-level correlation of 0.06 between exposure and the Series-A year.

Table 3, Panel B reports the continuous design and reproduces the pattern, with all four ratios rising per standard deviation of exposure (0.026 to 0.049 log points).

Figure 2 plots the event-study estimates for the two per-employee ratios, *Funding per Employee* (Panel A) and *Valuation per Employee* (Panel B), using the 3rd quarter of 2022 as the reference quarter. Before the release, the exposure coefficients are flat and jointly insignificant. Both ratios rise only after the 4th quarter of 2022 and remain elevated through the end of the sample, although for *Valuation per Employee* the bottom of the 95% confidence interval just crosses zero for the 3rd and 4th quarters of 2024.

4.2.3 Disentangling the Funding Contraction

An alternative to the GenAI interpretation is the 2022 venture funding contraction, which began at the very start of 2022. Perhaps more importantly for our design, the funding contraction continued throughout our post-period. If the post-period funding contraction fell more heavily on GenAI-exposed startups, it could have produced both slower hiring and higher capital per employee with no role for the technology. The design already limits this concern: the sector-by-quarter fixed effects absorb the funding environment common to firms in the same sector and quarter, and the hiring estimate is robust to controlling for the firm's initial capitalization (Table 2, column 2). The remaining concern is that the post-period funding contraction fell disproportionately on more GenAI-exposed firms. Table 4 provides evidence inconsistent with this scarcer-capital explanation in two ways.

First, a capital-side placebo confirms that the higher capital per employee reflects leanness rather than a movement in funding. The ratio of funding to employees used in Table 3 measures leanness only if exposed firms' funding held steady through the contraction, and a control for funding cannot verify this, since funding is part of the ratio. Therefore, we test whether exposed

firms raised less after the ChatGPT release, and columns 1 and 2 of Table 4 show they did not. High-exposure startups are no less likely to close a financing round (column 1, -0.002, t = -0.49) and raise no less when they do (column 2, -0.008, t = -0.56). Thus, GenAI-exposed firms' funding held up through the contraction, which cannot account for the higher funding per employee we observe.

Second, funding controls in the hiring regressions add comfort that the reduced hiring growth in the post-period for high exposure firms is not driven by such firms suffering a greater post-period funding contraction. The funding contraction could have slowed exposed firms' hiring through a weaker funding position rather than through GenAI, in which case controlling for the funding position would absorb the effect. In place of the initial-capitalization control already used in Table 2, column 4 of Table 4 uses the firm's pre-shock cumulative funding and the time since its last financing round, each interacted with Post, which together capture its funding position entering the contraction.[12] Adding these controls does not attenuate the coefficient from column 3 of Table 4 on *GenAIExposure* × *Post*. If anything, it grows in absolute value, to -0.056 (t = -3.75). The reduced growth in hiring therefore does not arise due to the funding winter.[13]

4.3 The Substitution Channel: Firms and Workers

4.3.1 GenAI-Adoption Propensity

Having established that exposed startups operate leaner and that the pattern is not a financing artifact, we now ask whether the effect operates through firms' capacity to substitute GenAI for labor. An exposed firm that builds or deploys generative AI in its operations can use AI in place

[12] Total Funding and Months Since Round are each fixed at the firm's last pre-shock quarter, so their levels do not vary within firm and are absorbed by the firm fixed effects.

[13] The hiring effect remains negative and statistically significant when we restrict the sample to firms that raised a financing round within the two years preceding the ChatGPT release, and when the funding controls are entered in levels rather than logarithms.

of employees, whereas a firm only nominally exposed to AI cannot. The leaner hiring should therefore concentrate among the GenAI-exposed firms best positioned to adopt the technology.

Table 5 estimates the exposure response separately within firms with higher and lower propensity to adopt GenAI and tests whether the responses differ between the two groups. We gauge a firm's propensity to adopt GenAI in three ways, each capturing a different facet of the firm. The first is the firm's product. Higher-propensity firms are those classified in an artificial-intelligence industry category whose business description also contains AI or machine-learning terms, identified with the keyword dictionaries of Gofman and Jin (2024), Lou and Wu (2021), and Shi et al. (2025). The second is the firm's leadership. We flag each of the firm's board members and senior managers as AI-technical when their title denotes an AI or technical-leadership role, such as chief technology, data, or AI officer, and measure the firm by the share of its leadership in such roles, taking firms above the median as higher-propensity. The third is the firm's hiring. We classify each of the firm's job postings using the classifier of Shi et al. (2025), which labels a posting as AI-related from its text. Higher-propensity firms are those that posted any AI role before the ChatGPT release.[14]

The leaner-hiring effect concentrates in higher-propensity firms under every measure. By the product measure (columns 1 and 2), the exposure response is -0.103 for higher-propensity firms ($t = -3.02$) but an insignificant -0.027 for lower-propensity firms ($t = -1.60$), a significant difference (difference $p = 0.046$). The leadership measure (columns 3 and 4) shows the cleanest split: -0.079 for AI-led firms ($t = -3.80$) against essentially zero for the rest, the most significant difference of the three (difference $p = 0.011$). The hiring measure (columns 5 and 6) reproduces it, with the

[14] We note that for this test, we require startup firms to have job postings covered by Lightcast, which reduces the sample size used in this specification.

higher-propensity response again near -0.10 (difference p = 0.062). An untabulated continuous intensity specification holds for each measure.

Figure 3 plots the event-study estimates by subgroup under each of the three measures, using the 3rd quarter of 2022 as the reference quarter. In every panel the higher- and lower-propensity coefficients are flat and overlapping before the release and separate only after 2022Q4. The gap thus opens at the ChatGPT release rather than continuing a pre-existing trend.

These results are most consistent with the slowed hiring growth we observe being attributable to the substitution channel. The reduced hiring growth concentrates where the firm is best positioned to deploy GenAI in place of labor and holds whether adoption propensity is measured based on the firm's product, its leadership, or its hiring.

4.3.2 Worker Composition

If the mechanism through which GenAI exposure leads to slowed hiring is task substitution, the reduction in labor demand should fall on the workers whose tasks are most readily performed by the technology and should not associate with worker attributes that are unrelated to task content. We decompose the overall hiring effect into category-specific hiring-growth regressions and test formally whether the exposure response differs across paired worker categories.

Table 6 reports the decomposition, which reveals three margins of task substitutability along which the reduced labor demand concentrates. First, it intensifies toward labor-market entry: hiring of workers aged 25 and under falls by 0.119 (column 1) compared to 0.050 for those 26 and older. Second, it concentrates in remote-suitable roles: classifying the top tercile of roles by remote-suitability as remote-suitable knowledge work, remote-suitable hiring falls by 0.068 (column 3) compared to 0.033 for onsite roles. Third, it falls more heavily on specialists than on generalists: among senior hires, workers whose prior careers are concentrated in a few roles

experience a hiring-growth decline of 0.040 (column 5), whereas the most generalist employees experience no change in hiring (column 6, 0.017).[15] These margins share a common theme: entry-age, remote-suitable, and narrowly specialized workers are those whose tasks are most readily performed by generative AI.

We also examine two further dimensions along which generative AI is widely expected to fall disproportionately, and fail to find a significant differential. AI is often expected to weigh more heavily on women than on men, since clerical and administrative work is more exposed to AI and the associated workers are disproportionately female (Pizzinelli et al. 2023; Gmyrek, Berg, and Bescond 2023). Additionally, the most educated employees undertake work that is more exposed to large language models (Eloundou et al. 2024). In our startups sample, however, the exposure response does not differ significantly by gender (columns 9 and 10) or by education (columns 7 and 8).

We examine the age margin more closely, asking whether the negative impact of GenAI on labor demand declines steadily with a potential worker's age or instead breaks sharply around the typical age for labor-market entry. We partition workers into six age bands, following Brynjolfsson, Chandar, and Chen (2025), and estimate the exposure response within each. Figure 4 shows the restraint is concentrated in the youngest, entry band (age 25 and under), where the decline is roughly twice that of any older band, while from early career through senior bands the effect is similar in magnitude. The reduction in hiring therefore falls most heavily on workers at the point of labor-market entry rather than declining steadily with age.

[15] Role breadth is the Blau (1977) index of heterogeneity, one minus the Herfindahl concentration of a worker's prior positions across role clusters (Appendix A), computed for senior hires (seniority of 3 and above), the population for which a career-trajectory measure is defined. The construct of distinguishing generalists from specialists by the breadth of prior career follows Custódio, Ferreira, and Matos (2013), Crossland, Zyung, Hiller, and Hambrick (2014), Lazear (2004, 2009), and Gathmann and Schönberg (2010). The differential survives orthogonalizing role breadth on age and seniority and strengthens when the 2021 funding-boom quarters are dropped.

Whereas existing work documents the economy-wide impact of generative AI on workers (Hosseini Maasoum and Lichtinger 2025; Brynjolfsson, Chandar, and Chen 2025), we establish the effect of GenAI in the distinct setting of scaling-stage startups, which are the earliest indicators of what a post-GenAI organizational structure will look like. Furthermore, we link the leaner-hiring effect to task substitution. The worker-level and firm-level evidence thus far are most consistent with the interpretation that GenAI-exposed startups add fewer workers because GenAI substitutes for labor.

4.4 Additional Analysis: Startups versus New Public Firms

We have shown that GenAI-exposed startups scale headcount more slowly after the release of ChatGPT, building leaner organizations consistent with AI substituting for hiring additional workers. In contrast, newly public firms have a more established workforce and thus achieving a similar leanness could occur only by firing existing workers, which entails adjustment costs and has implications for internal labor markets. Thus, newly public firms may be more likely than startup firms to experience the productivity or augmentation benefits of AI adoption, with AI use increasing the pace of hiring.

To separate the AI-exposure effect in firm formation versus incumbent adjustment, we compare the hiring response of GenAI-exposed startups with that of otherwise-comparable public firms. We identify newly public firms that completed an initial public offering between 2010 and 2021 from Compustat, operating in the same broad sectors as our startups and within a comparable-size range.[16] We link public firms to Revelio by issuer CUSIP and retain firms whose Revelio

[16] We harmonize Compustat GICS sectors and PitchBook industry sectors into seven common groups and compare firms only within a group through the sector-by-quarter fixed effects. The mapping assigns GICS Information Technology (45) to IT, Health Care (35) to Health, Consumer Discretionary and Consumer Staples (25 and 30) to Consumer, Industrials and Communication Services (20 and 50) to Business-to-Business, Financials (40) to Financial Services, Energy (10) to Energy, and Materials (15) to Materials; the PitchBook startup sectors map to the same seven groups. The size screen retains public firms with Compustat employment between 25 and 5,000. The upper bound is

headcount tracks their reported employment.[17] Since public firms have no Series A funding, we fix their GenAI exposure at their 2019 workforce, a predetermined, pre-pandemic baseline well before the release of ChatGPT.

Table 7 Panel A estimates the hiring response separately for venture-backed startups and new public firms, and tests the startup-minus-public difference in the exposure response.[18] We find that startups reduce hiring growth overall and for entry-level, remote-suitable, and specialist roles specifically, whereas newly public firms increase hiring growth on all dimensions. Specifically, among startups, *GenAIExposure* × *Post* is negative (-0.041, column 1, t = -2.76): exposed startups hire more slowly after the release of ChatGPT. Among otherwise-comparable public firms the corresponding coefficient is positive (0.069, column 2, t = 2.11): exposed incumbents hire *faster*.[19] The startup-minus-public difference is -0.110 and highly significant ($p < 0.01$).[20] Figure 5 plots the two within-group exposure paths, both roughly flat before the release and separating sharply only after the 4th quarter of 2022.

approximately the largest startup in our sample, keeping public firms within the size range startups occupy, while the lower bound drops the few micro-cap public firms for which Revelio covers too little of the workforce to yield a reliable hiring measure. The startup-minus-public reversal is not sensitive to these bounds, remaining negative and significant across floors from 25 to 50 and ceilings from 2,000 to 10,000 employees, with the reported window yielding the most conservative estimate (untabulated).

[17] Revelio headcount tracks these firms' reported employment, with a median Revelio-to-Compustat employee ratio near one in the well-covered information-technology and healthcare sectors, so the Revelio-based hiring measure is reliable for the public-firm panel. Of the 1,570 screened Compustat firms, 1,082 match to Revelio by issuer CUSIP and clear a minimum coverage ratio of 0.10, and 880 remain after the size screen.

[18] The comparison is confined to the hiring margins, the outcomes Revelio constructs for both populations. The capital-efficiency ratios of Table 3 are built from venture-financing rounds, which have no counterpart for established firms, and the adoption-propensity measures of Table 5 are constructed from venture-data classifications of a firm's product and leadership, so neither analysis carries over to the public-firm panel.

[19] Other research finds that more-exposed S&P 1500 firms post fewer jobs after the release of ChatGPT (Berger et al. 2026). A potential reason from the opposite direction of the two results is that recently listed public firms tend to be growing much more aggressively than typical S&P 1500 firms.

[20] The reversal is robust to restricting the new public firms to the most lifecycle-comparable subsample, those that completed their initial public offering in 2013 or later: the startup-minus-public difference stays negative and significant across all four margins (overall hiring -0.113, entry-level -0.270, remote-suitable -0.139, specialist -0.113).

The same reversal appears in the composition of the restraint documented in Section 4.3.2. Estimating the same within-population comparison for each worker-composition margin, the startup-minus-public difference is again negative and significant: entry-level hiring of workers age 25 and under (-0.254, columns 3-4), remote-suitable hiring (-0.138, columns 5-6), and specialist hiring (-0.111, columns 7-8).[21] In every case, GenAI-exposed incumbents increase hiring growth while GenAI-exposed startups decrease it.

In section 4.2.3, we find that the contraction in startup hiring is not attributable to reduced funding and in section 4.3 the evidence for startups is most consistent with a substitution channel. Newly public firms have an established workforce and thus may be more prone to expand their workforce to facilitate productivity or augmentation hiring growth when exposed to AI rather than grow leaner. We explore this possibility in Table 7 Panel B. Specifically, we examine whether firms with financial slack are more likely to increase hiring after GenAI exposure, consistent with an investment in labor to develop AI capabilities. We split the new public firms at the median of pre-shock cash holdings scaled by assets and re-estimate the hiring response within each group. The increased hiring for newly public firms is concentrated in cash-rich firms: *High GenAIExposure* × *Post* is 0.117 among cash-rich incumbents (column 1, t = 2.15) and an insignificant 0.026 among cash-lean ones (column 2). These same cash-rich public firms also increase capital spending, not only labor: *Capex per Employee* rises by 0.278 log points (column 3, t = 2.51) and is flat among cash-lean firms (column 4). This pattern of results is consistent with expansion, not leanness: where the startups' capital per worker rises because they hire fewer people, newly public firms add capital and headcount simultaneously. However, we caution that

[21] The remote-suitability and specialist margins are top-tercile indicators, defined within each population separately (each population's own top third), consistent with the within-population treatment terciles. The reversal on both margins is similar under a pooled combined-population cutoff (untabulated).

this evidence is suggestive in nature as the differences between cash-rich and cash-lean public firms is not statistically significant (p-values in excess of 0.10). In sum, newly public firms do not appear to change how they combine labor and capital after GenAI exposure and instead use their relative financing capacity to increase hiring in response to the changing AI landscape. The lean reorganization of labor and capital is thus a property of startup firms that scale with GenAI.

A growing literature documents that generative AI reduces hiring across the economy, concentrated among junior and entry-level workers. Our comparison shows that this aggregate pattern is not applicable to every organization: GenAI-exposed startups and comparable incumbents react to the same exposure in opposite directions. The results suggest that firms in the scaling phase find it optimal to reduce the rate of hiring for junior roles, whereas newly public firms instead increase hiring, perhaps to build AI capability as an addition to an existing workforce.

4.5 Robustness

4.5.1 Alternative Exposure Measures

Our baseline measure of GenAI exposure is well suited to our setting: constructed from the occupational task content of a firm's workforce (Brynjolfsson, Mitchell, and Rock 2018) and fixed at the firm's Series-A year, well before the release of ChatGPT. It is also not an exposure score rated by an AI model. This matters because such AI-generated ratings are sensitive to the model used: Yin, Vu, and Persico (2026) have three frontier models apply a single AI-exposure rubric and find that the resulting difference-in-differences estimates can switch sign depending on which AI model generated the AI-exposure measure.

Table 8, Panel A re-estimates the baseline specification under four alternatives: the language-model occupational measure of Felten, Raj, and Seamans (2023, LM-AIOE; column 1), the *human* ratings of the field's dominant exposure rubric (Eloundou et al. 2024; column 2) and

its GPT-4-rated counterpart (column 3), and the realized-usage Anthropic Economic Index (Handa et al. 2025; column 4).[22] The effect is negative and significant under the language-model, human-rated, and realized-usage measures (columns 1, 2, and 4). The lone exception is the GPT-4-rated rubric (column 3), which is null. Model-rated exposure scores are sensitive to the rating model (Yin, Vu, and Persico 2026), but the GPT-4-rated exposure result nevertheless indicates that our result is not perfectly robust to all possible ways of measuring GenAI exposure. Nevertheless, finding that the result holds across measures grounded in occupational task content, human ratings, and realized usage indicates that it does not depend on using a particular operationalization of exposure.

4.5.2 Distinguishing AI-exposure from the 2022 Monetary Tightening

The results in Table 4 are inconsistent with the 2022 contraction of capital being a primary driver of the leaner hiring we observe among startup firms. The funding winter that coincided with the release of ChatGPT did not differentially cut exposed firms' financing, and the reduced growth in hiring is unchanged controlling for the funding winter. However, the end of the zero-rate era that triggered the funding contraction could also operate through two channels that a capital-side analysis leaves open. The first is a labor-market channel, specifically the unwinding of 2021 over-hiring: even firms whose capital was untouched by the funding contraction may cut hiring if they expanded rapidly in the 2021 boom and then reverted.

The second is a valuation channel. Our GenAI-exposure measure may proxy for high-growth software firms, whose long-duration equity is the most sensitive to rising discount rates

[22] The Anthropic Economic Index is a revealed-usage measure, constructed by classifying millions of Claude conversations into occupational tasks (Handa et al. 2025). It captures where AI is observed in use rather than where tasks are rated as exposed, providing an independent, behavior-based complement to the predicted-exposure measures. One caveat is that, drawn from realized 2025 usage after the release of ChatGPT, it is not predetermined, so we report it as corroborating rather than identifying evidence.

(Dechow, Sloan, and Soliman 2004; Weber 2018). Near-zero interest rates had supported rich valuations for scalable, not-yet-profitable business models such as business-to-business software, and these valuations fell sharply as the Federal Reserve raised rates from March 2022, pressuring such firms to conserve cash and slow hiring. If GenAI exposure merely captures software intensity, the slowdown could reflect this repricing rather than the effect of GenAI.

Our design partly addresses both concerns by construction: the sector-by-quarter fixed effects absorb any macroeconomic event operating at the sector level, comparing exposed and less-exposed startups within the same sector and quarter. Thus, an alternative interpretation of the results needs to occur within a sector and affect GenAI-exposed firms more than less exposed firms or requires GenAI exposure to proxy for software intensity even within an industry. Table 8, Panel B addresses these residual concerns directly.

The labor-market concern predicts that high-exposure firms over-hired during the 2021 boom and then corrected, which requires GenAI exposure to be positively related to 2021 hiring growth. It is not: GenAI exposure is essentially uncorrelated with a firm's 2021 hiring growth (a firm-level correlation of 0.01, untabulated), so the slowdown cannot be the boom's reversion. Consistent with this, adding the firm's predetermined 2021 hiring growth interacted with Post leaves the effect negative and significant (-0.029, t = -2.03, Panel B, column 2), even though the control is itself strongly significant. The slowdown is thus not the reversal of boom-era over-hiring.

The second concern predicts that GenAI exposure proxies for software intensity, so a measure of software intensity should be positively associated with the slowdown and should absorb the exposure effect. Neither implication holds; if anything, the association is of the opposite sign. GenAI exposure is negatively correlated with the share of a firm's workforce in computer and mathematical occupations (a correlation of -0.08, untabulated), and the most software-

intensive startups hired more after the release of ChatGPT, not less (Table 8, Panel B, column 3). Controlling for tech-intensity, the AI-exposure effect remains negative and significant (Panel B, column 4),[23] while tech-intensity enters with the opposite sign. The leaner hiring is thus a response to workforce GenAI-task exposure, not to software intensity that made firms vulnerable to the 2022 repricing of high-growth equity.

4.5.3 Placebo, Anticipation, and Estimation Window

If the leaner-hiring effect reflects the release of ChatGPT, it should fail to appear using pseudo-event dates. We confirm this in untabulated specifications, for both the binary and continuous designs of Table 2. First, a placebo test assigns a pseudo-release date at the 4th quarter of 2020 or 2021, and finds no decline at either, inconsistent with the effects we document being attributable to diverging trends underway before the release of ChatGPT. Second, the effect does not depend on the quarters closest to the release of ChatGPT. Our results are essentially unchanged when we drop Q2 and Q3 of 2022; thus GenAI-exposed firms to not appear to anticipate the technology. The results are also robust to dropping the three quarters surrounding the release, from Q3 2022 through Q1 2023. Third, our inferences are stable across alternative estimation windows: symmetric windows of six to nine quarters on each side of ChatGPT's release, and a tighter 2021-2024 sample that drops the 2020 pandemic period, indicating the slowdown is not an artifact of COVID-era hiring disruptions.

4.5.4 Alternative Fixed Effects

Beyond the timing of the effect, a separate concern is that leaner hiring reflects not GenAI but a factor that happens to correlate with GenAI exposure, such as the local labor markets where

[23] This is a conservative control: the share of computer and mathematical occupations overlaps the workforce composition from which the exposure measure is built, so it absorbs part of the treatment's own variation.

exposed firms concentrate or the funding vintages to which they belong. Results from untabulated specifications are inconsistent with both concerns. First, we use metropolitan-area-by-quarter fixed effects to confine the comparison to exposed and less-exposed firms in the same city and quarter, so the effect does not appear to reflect a cooling local labor market in the hubs where exposed firms cluster. Second, Series-A-cohort-by-quarter fixed effects make comparisons within firms that raised capital at the same time, inconsistent with the funding contraction hitting firms at a particular stage differentially. The leaner-hiring effect survives both alternative fixed effect structures: absorbing the finer time-varying variation attenuates it modestly, but the effect remains negative and significant in both the binary and continuous designs, with clean pre-trends.

5. Conclusion

The configuration of labor and capital is one of the most consequential firm choices, and a technology that substitutes for labor can change this mix. A large literature measures the exposure of occupations and tasks to artificial intelligence and traces its consequences for workers and aggregate employment, but little is known about how startup firms choose their mix of labor and capital as they scale with GenAI. In this paper, we examine how the introduction of generative AI reshapes the combination of labor and capital in the organizations that new firms build, comparing firms by their predetermined GenAI exposure around the public release of ChatGPT.

We find that GenAI-exposed startups build leaner organizations. After the release of ChatGPT, more exposed startups scale headcount significantly more slowly than comparable, less-exposed peers, and carry more capital and market value per employee. The leanness operates through the firm's capacity to substitute GenAI for labor and is most pronounced among the startups best positioned to do so, including those building AI products, those led by AI-technical executives, and those already hiring for AI roles. Within firms, the reduced demand for labor falls

on the workers whose tasks the technology most readily performs, the early-career, remote-suitable, and narrowly specialized tasks, rather than on workers distinguished by gender or education. The leaner response appears specific to young private firms: among newly public firms, greater GenAI exposure is associated with *more* hiring rather than less.

Taken together, these findings reveal that generative AI reshapes labor demand at the level of the firm, not only the worker: startups exposed to the technology scale with fewer employees and appear to substitute capital for labor, building a leaner organization from the start. Whereas extant literature has largely examined GenAI's influence on workers and occupations, our paper contributes by examining firm-wide responses to GenAI. Furthermore, we extend the finance literature on human capital and firm boundaries by showing that the introduction of ChatGPT shifts the firm's mix of labor and capital, raising capital and market value per employee.

As the economy shapes itself around generative AI as a disruptive technology, our evidence in the startup setting provides an important insight into how firms that grow organically with GenAI choose their mix of labor and capital. Our results suggest that GenAI acts as a substitute for labor for firms in the scaling phase. Although we cannot yet observe whether firms that scale with GenAI later ramp up their relative use of labor as they mature, the leaner scaling effect we observe has implications for the entry-level hiring pipeline, remote work, and specialist skills.

References

ACEMOGLU, D. "The Simple Macroeconomics of AI." *Economic Policy* 40 (2025): 13–58.

ACEMOGLU, D.; D. AUTOR; J. HAZELL; AND P. RESTREPO. "Artificial Intelligence and Jobs: Evidence from Online Vacancies." *Journal of Labor Economics* 40 (2022): S293–S340.

ACEMOGLU, D., AND P. RESTREPO. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33 (2019): 3–30.

BABINA, T.; A. FEDYK; A. HE; AND J. HODSON. "Artificial Intelligence, Firm Growth, and Product Innovation." *Journal of Financial Economics* 151 (2024): 103745.

BECKER, G. S. "Investment in Human Capital: A Theoretical Analysis." *Journal of Political Economy* 70 (1962): 9–49.

BECKER, G. S., AND K. M. MURPHY. "The Division of Labor, Coordination Costs, and Knowledge." *The Quarterly Journal of Economics* 107 (1992): 1137–1160.

BELO, F.; X. LIN; AND S. BAZDRESCH. "Labor Hiring, Investment, and Stock Return Predictability in the Cross Section." *Journal of Political Economy* 122 (2014): 129–177.

BENSON, A., AND K. SHAW. "What Do Managers Do? An Economist's Perspective." *Annual Review of Economics* 17 (2025): 635–664.

BERGER, P. G.; W. CAI; L. QIU; AND C. X. SHEN. "Employer and Employee Responses to Generative AI: Early Evidence." Working paper, University of Chicago, 2026.

BESSEN, J. E.; F. POEGE; AND R. RÖTTGER. "Competing for Talent: Large Firms and Startup Growth." Working paper, Boston University, 2024.

BLAU, P. M. *Inequality and Heterogeneity: A Primitive Theory of Social Structure*. New York: Free Press, 1977.

BRYNJOLFSSON, E.; B. CHANDAR; AND R. CHEN. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence." Working paper, Stanford Digital Economy Lab, 2025.

BRYNJOLFSSON, E., AND T. MITCHELL. "What Can Machine Learning Do? Workforce Implications." *Science* 358 (2017): 1530–1534.

BRYNJOLFSSON, E.; T. MITCHELL; AND D. ROCK. "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings* 108 (2018): 43–47.

CALLAWAY, B.; A. GOODMAN-BACON; AND P. H. C. SANT'ANNA. "Difference-in-Differences with a Continuous Treatment." Working paper, National Bureau of Economic Research, 2024.

CHOI, J.; N. GOLDSCHLAG; J. HALTIWANGER; AND J. D. KIM. "Early Joiners and Startup Performance." *The Review of Economics and Statistics* (2023). Forthcoming.

CORE, F.; G. RYDUCHOWSKA; AND S. WANG. "Human Capital Acquisition in Venture-Capital-Backed Companies." Working paper, 2024.

CROSSLAND, C.; J. ZYUNG; N. J. HILLER; AND D. C. HAMBRICK. "CEO Career Variety: Effects on Firm-Level Strategic and Social Novelty." *Academy of Management Journal* 57 (2014): 652–674.

CUSTÓDIO, C.; M. A. FERREIRA; AND P. MATOS. "Generalists versus Specialists: Lifetime Work Experience and Chief Executive Officer Pay." *Journal of Financial Economics* 108 (2013): 471–492.

DAVILA, A.; G. FOSTER; AND M. GUPTA. "Venture Capital Financing and the Growth of Startup Firms." *Journal of Business Venturing* 18 (2003): 689–708.

DAVIS, S. J., AND J. HALTIWANGER. "Gross Job Creation, Gross Job Destruction, and Employment Reallocation." *The Quarterly Journal of Economics* 107 (1992): 819–863.

DECHOW, P. M.; R. G. SLOAN; AND M. T. SOLIMAN. "Implied Equity Duration: A New Measure of Equity Risk." *Review of Accounting Studies* 9 (2004): 197–228.

DECKER, R.; J. HALTIWANGER; R. JARMIN; AND J. MIRANDA. "The Role of Entrepreneurship in US Job Creation and Economic Dynamism." *Journal of Economic Perspectives* 28 (2014): 3–24.

DESANTOLA, A., AND R. GULATI. "Scaling: Organizing and Growth in Entrepreneurial Ventures." *Academy of Management Annals* 11 (2017): 640–668.

DINLERSOZ, E.; T. DOGAN; AND N. ZOLAS. "Starting Up AI." Working paper, U.S. Census Bureau Center for Economic Studies, 2024.

EISFELDT, A. L., AND D. PAPANIKOLAOU. "Organization Capital and the Cross-Section of Expected Returns." *Journal of Finance* 68 (2013): 1365–1406.

EISFELDT, A. L.; G. SCHUBERT; M. B. ZHANG; AND B. TASKA. "Generative AI and Firm Values." *Journal of Finance* (2026).

ELOUNDOU, T.; S. MANNING; P. MISHKIN; AND D. ROCK. "GPTs Are GPTs: Labor Market Impact Potential of Large Language Models." *Science* 384 (2024): 1306–1308.

EWENS, M.; R. NANDA; AND M. RHODES-KROPF. "Cost of Experimentation and the Evolution of Venture Capital." *Journal of Financial Economics* 128 (2018): 422–442.

FELTEN, E.; M. RAJ; AND R. SEAMANS. "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses." *Strategic Management Journal* 42 (2021): 2195–2217.

FELTEN, E.; M. RAJ; AND R. SEAMANS. "How Will Language Modelers like ChatGPT Affect Occupations and Industries?" Working paper, 2023.

GATHMANN, C., AND U. SCHÖNBERG. "How General Is Human Capital? A Task-Based Approach." *Journal of Labor Economics* 28 (2010): 1–49.

GMYREK, P.; J. BERG; AND D. BESCOND. "Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality." Working paper, International Labour Organization, 2023.

GOFMAN, M., AND Z. JIN. "Artificial Intelligence, Education, and Entrepreneurship." *Journal of Finance* 79 (2024): 631–667.

GOMPERS, P. A. "Optimal Investment, Monitoring, and the Staging of Venture Capital." *Journal of Finance* 50 (1995): 1461–1489.

HAMERMESH, D. S., AND G. A. PFANN. "Adjustment Costs in Factor Demand." *Journal of Economic Literature* 34 (1996): 1264–1292.

HAMPOLE, M.; D. PAPANIKOLAOU; L. D. W. SCHMIDT; AND B. SEEGMILLER. "Artificial Intelligence and the Labor Market." Working paper, National Bureau of Economic Research, 2025.

HANDA, K.; A. TAMKIN; M. MCCAIN; S. HUANG; E. DURMUS; S. HECK; J. MUELLER; J. HONG; S. RITCHIE; T. BELONAX; K. K. TROY; D. AMODEI; J. KAPLAN; J. CLARK; AND D. GANGULI. "Which Economic Tasks Are Performed with AI? Evidence from Millions of Claude Conversations." Working paper, Anthropic, 2025.

HANNAN, M. T., AND J. FREEMAN. "Structural Inertia and Organizational Change." *American Sociological Review* 49 (1984): 149–164.

HELLMANN, T., AND M. PURI. "Venture Capital and the Professionalization of Start-Up Firms: Empirical Evidence." *Journal of Finance* 57 (2002): 169–197.

HERSHBEIN, B., AND L. B. KAHN. "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings." *American Economic Review* 108 (2018): 1737–1772.

HOSSEINI MAASOUM, S. M., AND G. LICHTINGER. "Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data." Working paper, 2025.

HOWELL, S. T.; D. PARKER; AND T. XU. "Tyranny of the Personal Network: The Limits of Arm's Length Fundraising in Venture Capital." Working paper, National Bureau of Economic Research, 2024.

HUMLUM, A., AND E. VESTERGAARD. "Large Language Models, Small Labor Market Effects." Working paper, National Bureau of Economic Research, 2025.

JÄGER, S., AND J. HEINING. "How Substitutable Are Workers? Evidence from Worker Deaths." Working paper, National Bureau of Economic Research, 2022.

JANEWAY, W. H.; R. NANDA; AND M. RHODES-KROPF. "Venture Capital Booms and Start-Up Financing." *Annual Review of Financial Economics* 13 (2021): 111–127.

JARAVEL, X.; N. PETKOVA; AND A. BELL. "Team-Specific Capital and Innovation." *American Economic Review* 108 (2018): 1034–1073.

JOHNSTON, A. C., AND C. A. MAKRIDIS. "AI, Output, and Employment." Working paper, CESifo, 2026.

KARAHAN, F.; B. PUGSLEY; AND A. ŞAHIN. "Demographic Origins of the Start-up Deficit." *American Economic Review* 114 (2024).

KHARAZIAN, A; L. SIMON; AND R. STEVENS. "A New Look at AI's Impact on Jobs: Firm-Level AI Spending and Workforce Adjustment." Working Paper, Ramp and Revelio Labs, 2026.

KIM, J. D. "Is There a Startup Wage Premium? Evidence from MIT Graduates." *Research Policy* 47 (2018): 637–649.

LAZEAR, E. P. "Balanced Skills and Entrepreneurship." *American Economic Review* 94 (2004): 208–211.

LAZEAR, E. P. “Firm-Specific Human Capital: A Skill-Weights Approach.” *Journal of Political Economy* 117 (2009): 914–940.

LEE, S., AND J. D. KIM. “When Do Startups Scale? Large-Scale Evidence from Job Postings.” *Strategic Management Journal* 45 (2024): 1633–1669.

LI, Q.; B. LOURIE; A. NEKRASOV; AND T. SHEVLIN. “Employee Turnover and Firm Performance: Large-Sample Archival Evidence.” *Management Science* 68 (2022): 5667–5683.

LOU, B., AND L. WU. “AI on Drugs: Can Artificial Intelligence Accelerate Drug Development? Evidence from a Large-Scale Examination of Bio-Pharma Firms.” *MIS Quarterly* 45 (2021): 1451–1482.

MORTENSEN, D. T., AND C. A. PISSARIDES. “Job Creation and Job Destruction in the Theory of Unemployment.” *The Review of Economic Studies* 61 (1994): 397–415.

MURRAY, C. “Heavy Corporate AI Spenders Add Staff Faster Than Peers.” *Financial Times* (June 29, 2026).

NICKELL, S. J. “Biases in Dynamic Models with Fixed Effects.” *Econometrica* 49 (1981): 1417–1426.

OTIS, N.; R. CLARKE; S. DELECOURT; D. HOLTZ; AND R. KONING. “The Uneven Impact of Generative AI on Entrepreneurial Performance.” Working paper, 2024.

OUIMET, P., AND R. ZARUTSKIE. “Who Works for Startups? The Relation between Firm Age, Employee Age, and Growth.” *Journal of Financial Economics* 112 (2014): 386–407.

PITCHBOOK DATA, INC. “How PitchBook Collects Data.” PitchBook Data, Inc., 2024. https://pitchbook.com/research-process (accessed July 2026).

PIZZINELLI, C.; A. PANTON; M. M. TAVARES; M. CAZZANIGA; AND L. LI. “Labor Market Exposure to Artificial Intelligence: Cross-country Differences and Distributional Implications.” Working paper, International Monetary Fund, 2023.

ROBB, A. M., AND D. T. ROBINSON. “The Capital Structure Decisions of New Firms.” *The Review of Financial Studies* 27 (2014): 153–179.

SHI, H.; A. GUPTA; W. W. DING; AND K. ZHANG. “Identifying AI Jobs in Online Job Postings: Existing versus Newly Developed LLM-Based Approaches.” Working paper, University of Maryland, 2025.

SORENSON, O.; M. S. DAHL; R. CANALES; AND M. D. BURTON. “Do Startup Employees Earn More in the Long Run?” *Organization Science* 32 (2021): 587–604.

TEUTLOFF, O.; J. EINSIEDLER; O. KÄSSI; F. BRAESEMANN; P. MISHKIN; AND R. M. DEL RIO-CHANONA. “Winners and Losers of Generative AI: Early Evidence of Shifts in Freelancer Demand.” *Journal of Economic Behavior & Organization* 235 (2025): 106845.

WEBB, M. “The Impact of Artificial Intelligence on the Labor Market.” Working paper, 2020.

WEBER, M. “Cash Flow Duration and the Term Structure of Equity Returns.” *Journal of Financial Economics* 128 (2018): 486–503.

YIN, M.; H. VU; AND C. PERSICO. “How (un)Stable Are LLM Occupational Exposure Scores? Evidence from Multi-Model Replication.” Working paper, National Bureau of Economic Research, 2026.

Appendix A: Variable Definitions

Variable	Definition
Dependent Variables	
Hiring Growth	The change in quarter-end Headcount (the count of unique employees active at the firm in the quarter) over four quarters, divided by the average of its current and four-quarters-prior values (the Davis-Haltiwanger symmetric growth rate; the twelve-month horizon): ($Headcount_t$ - $Headcount_{t-4}$) / (0.5 ($Headcount_t$ + $Headcount_{t-4}$)). Bounded between -2 and 2 by construction, with -2 denoting exit and +2 entry. Source: Revelio Labs.
ΔLog(headcount)	The log of quarter-end headcount minus the log of headcount four quarters earlier: log($Headcount_t$) - log($Headcount_{t-4}$). Source: Revelio Labs.
Funding / employee	The log of the firm's cumulative funding (the total capital it has raised across all rounds through quarter t, in millions of dollars) divided by quarter-end headcount: log(cumulative $funding_t$ / $Headcount_t$); winsorized 1/99 by year. Source: PitchBook, Revelio Labs.
Valuation / employee	The log of the firm's post-money valuation divided by quarter-end headcount: log(post-money $valuation_t$ / $Headcount_t$); winsorized 1/99 by year. Source: PitchBook (academic feed), Revelio Labs.
Funding / net-new emp	The log of cumulative funding divided by net-new employees, capital per scaling hire: log(cumulative $funding_t$ / ($Headcount_t$ - $Headcount^A$)); winsorized. Net-new employees $Headcount_t$ - $Headcount^A$ is the headcount the firm has added since its Series A round, where $Headcount^A$ is headcount in the Series A year; it counts the workers hired during the firm's scaling phase rather than its founding team. Source: PitchBook, Revelio Labs.
Valuation / net-new emp	The log of post-money valuation divided by net-new employees, market value per scaling hire: log(post-money $valuation_t$ / ($Headcount_t$ - $Headcount^A$)); winsorized (net-new employees defined as above). Source: PitchBook (academic feed), Revelio Labs.
Treatment, Split, and Control Variables	
High GenAIExposure	Indicator equal to one for firms in the top tercile of predetermined GenAI exposure and zero for those in the bottom tercile (the control group); the middle tercile is dropped. GenAI exposure is the employee-weighted average across the firm's workforce of the occupation-level suitability-for-machine-learning index (SML) of Brynjolfsson, Mitchell, and Rock (2018), mapped to workers through their ONET-SOC occupation and measured from the workforce in the firm's Series A year (startups) or the 2019 workforce (new public firms), so that it predates the ChatGPT release. Source: Revelio Labs; Brynjolfsson, Mitchell, and Rock (2018).
GenAIExposure (1 sd)	The firm's GenAI exposure (the SML-based measure defined above) standardized to mean zero and unit standard deviation within sample; the continuous treatment.
Post	Indicator equal to one for quarters from 2022Q4 onward, the quarter of the public release of ChatGPT.
Log(lagged headcount)	The log of quarter-end headcount four quarters earlier: log($Headcount_{t-4}$). Source: Revelio Labs.
Log(Series A funding)	Log of the amount raised in the firm's Series A round (in millions of dollars). Source: PitchBook.
Startup	Indicator equal to one for venture-backed startups and zero for new public firms. Source: PitchBook, Compustat.

Cash-Rich / Cash-Lean	Among new public firms, a median split on pre-shock cash slack, cash and short-term investments divided by total assets, averaged over 2019-2021; cash-rich denotes the upper half. Source: Compustat.
Capex per Employee	The log of one plus single-quarter capital expenditure per employee, $\log(1 + \text{capex}_t/\text{Headcount}_t)$, where quarterly capital expenditure is de-cumulated from the Compustat year-to-date series; winsorized 1/99 by quarter. Source: Compustat.
AI-Adoption Propensity	A firm's propensity to adopt AI, used in Table 5 to split firms into higher- and lower-propensity subsamples. We measure it three ways, each capturing a different facet of the firm: (i) product, an indicator that the firm is classified in an artificial-intelligence industry category and its description or keyword field contains AI or machine-learning terms; (ii) leadership, the share of the firm's board members and senior managers holding AI or technical-leadership titles, with firms split at the median; and (iii) hiring, an indicator that the firm posted any AI-related job before the release of ChatGPT. The product keyword list combines the AI/ML terms of Gofman and Jin (2024), Lou and Wu (2021), and Shi et al. (2025) with contemporary generative-AI terms (e.g., ChatGPT, generative AI); AI job postings are classified from their text by the classifier of Shi et al. (2025). Source: PitchBook; LightCast; Gofman and Jin (2024); Lou and Wu (2021); Shi et al. (2025).
LM-AIOE	Language-model occupational exposure, an alternative AI-exposure measure used for robustness in Table 8, aggregated to the firm on the same employee-weighted, Series-A-baseline basis as the primary SML measure. Source: Felten, Raj, and Seamans (2023); Revelio Labs.
Eloundou β (human)	Generative-AI task exposure, human-rated, under the Eloundou et al. (2024) rubric; aggregated to the firm on the same basis as above. Source: Eloundou et al. (2024); Revelio Labs.
Eloundou β (GPT-4)	Generative-AI task exposure under the same Eloundou et al. (2024) rubric but GPT-4-rated; aggregated to the firm on the same basis as above. Source: Eloundou et al. (2024); Revelio Labs.
Anthropic Index	Realized Claude-usage exposure from the Anthropic Economic Index built from post-ChatGPT (2025) Claude-conversation data, aggregated to the firm on the same basis as above. Source: Handa et al. (2025); Revelio Labs.
2021 Hiring Growth	The firm's average *Hiring Growth* over the four quarters of 2021, a predetermined measure of the firm's exposure to the 2021 venture-funding boom; standardized within sample. Source: Revelio Labs.
Tech-Occ Share	The share of the firm's workforce in Computer and Mathematical occupations (SOC major group 15), measured at the Series A year; standardized within sample and interacted with Post in Table 8. Source: Revelio Labs.
Entry-Level / Age 26+	Entry-level denotes workers aged 25 and under; the comparison group is 26 and over. Age is proxied as the calendar year minus the worker's bachelor-graduation year plus 22. Source: Revelio Labs.
Remote-Suitable / Onsite	Remote-suitable denotes the top tercile of the role-level remote-suitability score (0 to 1); onsite denotes the remainder. Source: Revelio Labs.
Specialist / Generalist	Worker role breadth, computed as one minus the Herfindahl concentration of the worker's predetermined (pre-hire) positions across Revelio role clusters (one minus the sum of squared role shares), from the worker's full-time, post-graduation, multi-employer career history before the hire date. Among senior hires (Revelio seniority of 3 or above), generalist denotes the top tercile of role breadth and specialist the remainder. Source: Revelio Labs.
Bachelor- / Graduate	Highest degree at or below a bachelor's versus a graduate degree (Master's, MBA, or PhD). Source: Revelio Labs.
Female / Male	Gender indicator imputed by Revelio. Source: Revelio Labs.

Total Funding	The log of one plus the firm's cumulative funding (total capital raised across all rounds), fixed at the firm's last pre-2022Q4 quarter. Source: PitchBook.
Log(1 + $M raised)	The log of one plus the millions of dollars raised in the quarter: log(1 + $M raised). Source: PitchBook.
Raised a Round	Indicator for any financing round closed in the quarter. Source: PitchBook.
Months Since Round	Number of months since the firm's most recent financing round, fixed at the firm's last pre-2022Q4 quarter (a proxy for remaining runway). Source: PitchBook.

Figure 1. Event Study: Hiring Growth by GenAI Exposure

Figure 1 displays event-study estimates of the leaner-hiring effect. Panel A plots the coefficients on *High GenAIExposure* (the top-versus-bottom tercile) interacted with each quarter relative to the ChatGPT release; Panel B plots the coefficients on *GenAIExposur (standardized)*, the continuous intensity design. Both are difference-in-differences event studies: each point is the exposure effect on hiring growth in a given quarter relative to the ChatGPT release, normalized to the reference quarter 2022Q3. *Hiring Growth* is the change in quarter-end headcount over four quarters, divided by the average of its current and four-quarters-prior values. *Change in Log Headcount* is the log of quarter-end headcount minus the log of headcount four quarters earlier. Vertical bars denote 95% confidence intervals from standard errors clustered at the firm level.

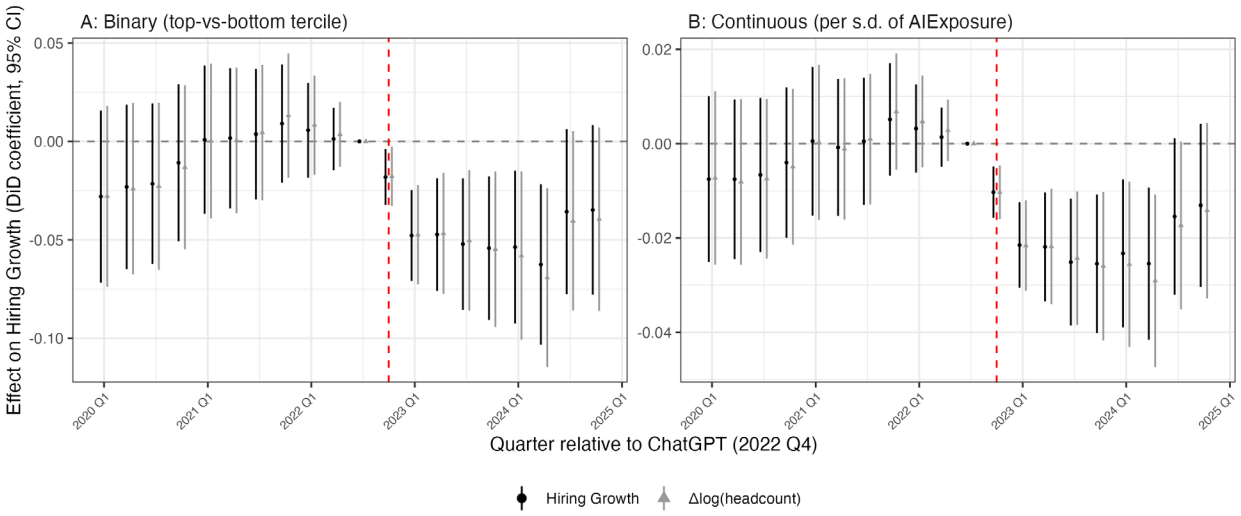


Figure 2. Event Study: Capital and Valuation per Employee

Figure 2 displays event-study estimates of the per-employee capital ratios. The figure plots the exposure coefficients at each quarter relative to the ChatGPT release, normalized to the reference quarter 2022Q3, for *Funding per Employee* and *Valuation per Employee*. *Funding per Employee* is the log of cumulative funding divided by quarter-end headcount. *Valuation per Employee* is the log of post-money valuation divided by quarter-end headcount. Vertical bars denote 95% confidence intervals from standard errors clustered at the firm level.

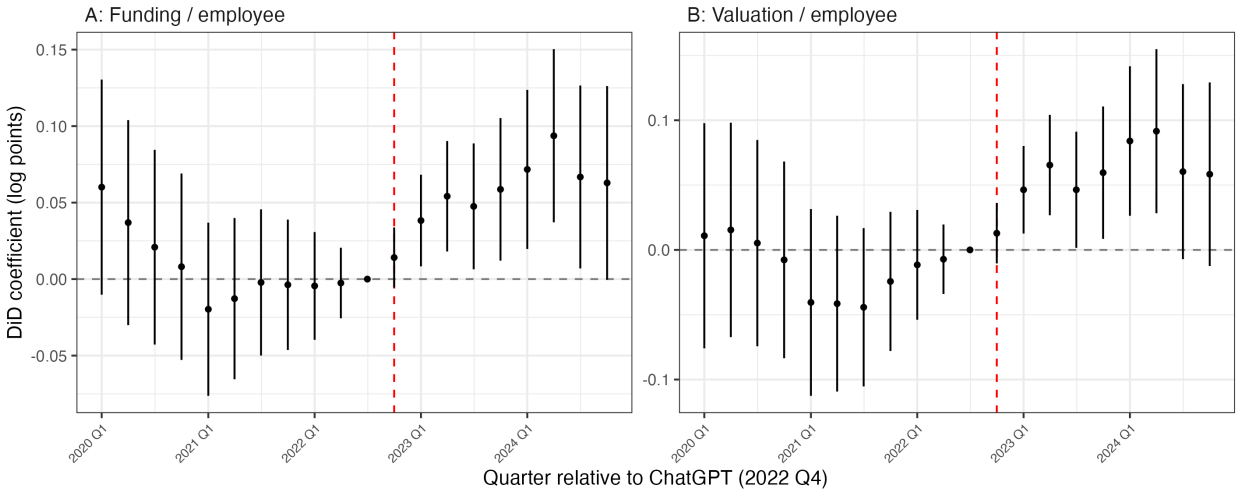


Figure 3. Event Study: Hiring by AI-Adoption Propensity

Figure 3 displays event-study estimates of the leaner-hiring effect separately for higher- and lower-AI-propensity startups, across the three measures of AI-adoption propensity in Table 5 (Panel A: AI product; Panel B: AI-technical leadership; Panel C: AI hiring). The figure plots the exposure coefficients at each quarter relative to the ChatGPT release, normalized to the reference quarter 2022Q3, for *Hiring Growth* within each subgroup. *Hiring Growth* is the change in quarter-end headcount over four quarters, divided by the average of its current and four-quarters-prior values. Vertical bars denote 95% confidence intervals from standard errors clustered at the firm level.

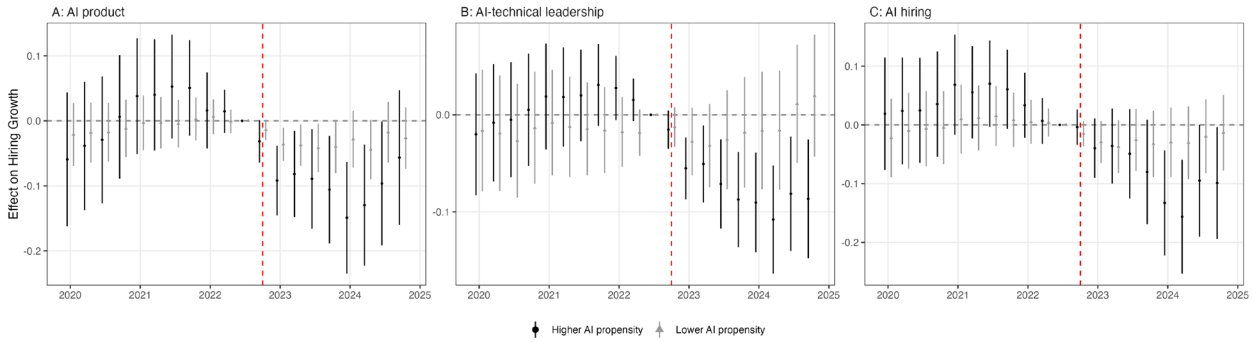


Figure 4. Hiring Response by Worker Age

Figure 4 displays the exposure effect on *Hiring Growth* across six worker age bands following Brynjolfsson, Chandar, and Chen (2025): entry (25 and under), early career (26 to 30), developing (31 to 34), mid-career 1 (35 to 40), mid-career 2 (41 to 49), and senior (50 and over). Each point is the difference-in-differences coefficient for the corresponding age band. *Hiring Growth* is the change over four quarters in the count of workers in a given age band, divided by the average of its current and four-quarters-prior values. Vertical bars denote 95% confidence intervals from standard errors clustered at the firm level.

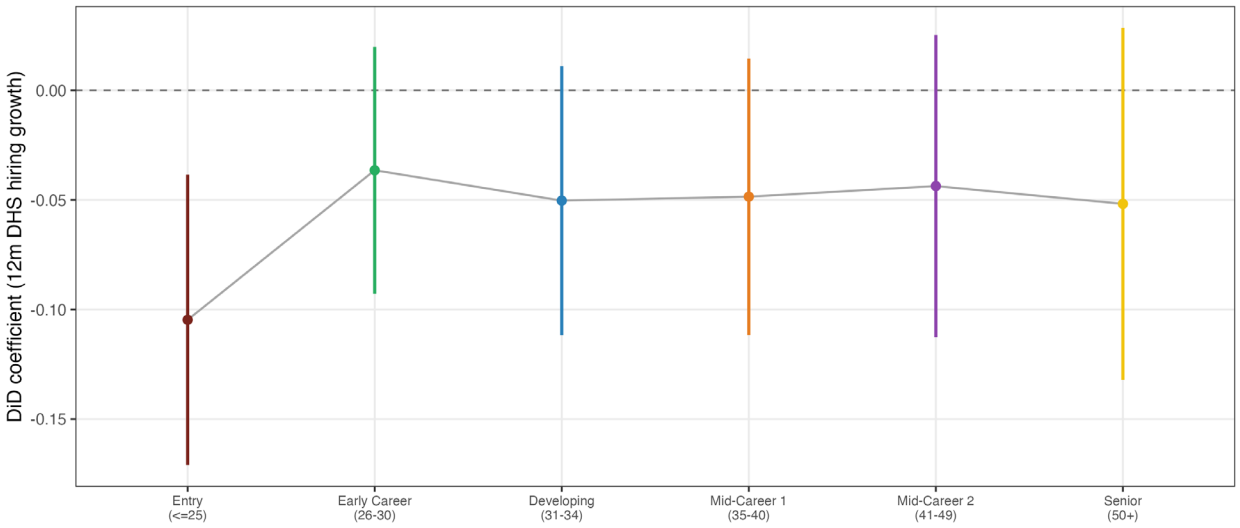


Figure 5. Event Study: Startups versus Public Firms

Figure 5 compares the AI-exposure hiring response of venture-backed startups and new public firms, asking whether the leaner-hiring effect is specific to new firms. The new public firms are Compustat firms that completed an initial public offering between 2010 and 2021, in the same sectors as the startups and overlapping their size range. The figure plots the coefficients on *High GenAIExposure* interacted with each quarter relative to the ChatGPT release, normalized to the reference quarter 2022Q3, estimated separately within each population, for *Hiring Growth* under the binary top-versus-bottom *GenAIExposure* tercile design. *High GenAIExposure* is fixed at each firm's pre-determined workforce, measured at the Series-A year for startups and at the 2019 workforce for new public firms. *Hiring Growth* is the change in quarter-end headcount over four quarters, divided by the average of its current and four-quarters-prior values. Vertical bars denote 95% confidence intervals from standard errors clustered at the firm level.

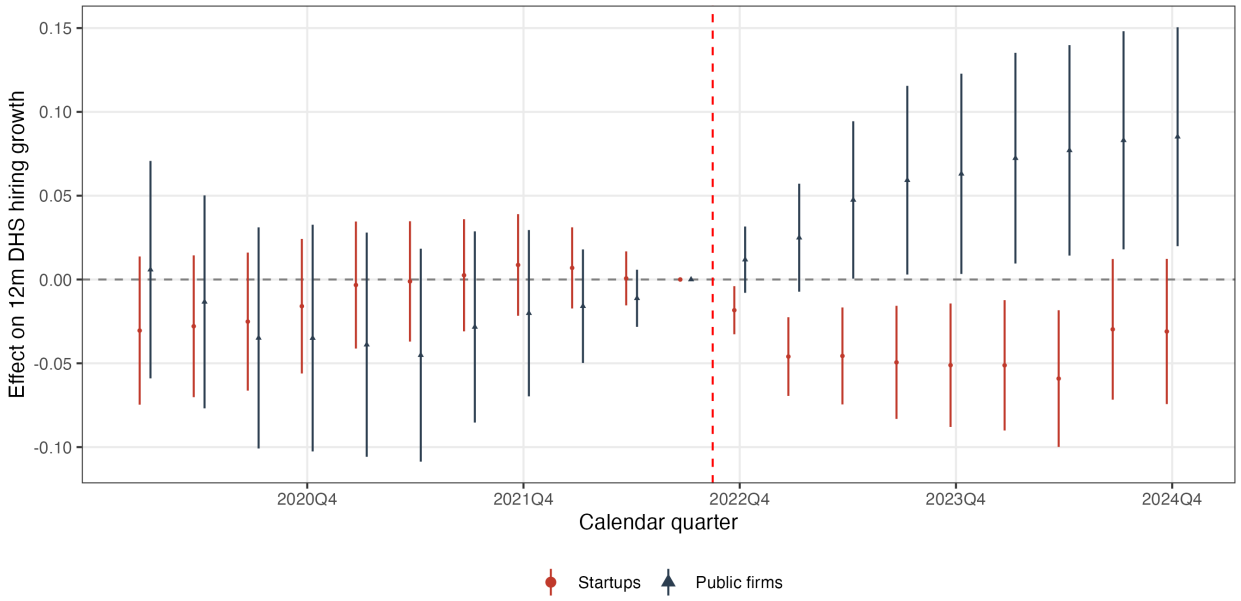


Table 1. Summary Statistics

Table 1 reports descriptive statistics for the main variables used in the analysis, on the firm-quarter panel of venture-backed startups covering 2020 through 2024. Panel A reports summary statistics for the two estimation samples: the binary sample (the top and bottom *GenAIExposure* terciles, with the middle tercile dropped) and the continuous sample (the full estimation sample). Panel B reports the industry-sector composition of the binary sample at the firm level. Panel C presents Pearson correlations (lower diagonal) and Spearman correlations (upper diagonal) among the main analysis variables, on the binary sample. All variables are defined in Appendix A.

Panel A: Descriptive Statistics

Variable	*N*	*Mean*	*SD*	*p25*	*p50*	*p75*
Binary sample						
Hiring Growth	56,745	-0.035	0.465	-0.314	0.000	0.267
ΔLog(headcount)	56,745	-0.038	0.488	-0.288	0.000	0.257
Funding / employee ($M)	56,745	2.041	2.752	0.528	1.149	2.364
Valuation / employee ($M)	50,208	5.774	9.390	1.323	2.851	6.250
Funding / net-new emp ($M)	46,804	4.056	7.228	0.791	1.778	4.088
Valuation / net-new emp ($M)	41,658	11.080	21.490	2.083	4.612	10.700
High GenAIExposure	56,745	0.498	0.500	0.000	0.000	1.000
Lagged headcount	56,745	57.854	131.726	14.000	30.000	62.000
Series-A funding ($M)	56,745	19.001	40.967	5.700	10.800	20.000
Post	56,745	0.447	0.497	0.000	0.000	1.000
Continuous sample						
Hiring Growth	85,666	-0.044	0.466	-0.321	0.000	0.257
ΔLog(headcount)	85,666	-0.047	0.490	-0.305	0.000	0.251
Funding / employee ($M)	85,666	1.978	2.689	0.528	1.118	2.254
Valuation / employee ($M)	76,070	5.625	9.292	1.321	2.771	6.000
Funding / net-new emp ($M)	69,594	3.923	7.053	0.781	1.716	3.897
Valuation / net-new emp ($M)	62,166	10.793	21.136	2.059	4.478	10.192
GenAIExposure (1 sd)	85,666	-0.002	0.996	-0.628	-0.021	0.591
Lagged headcount	85,666	60.157	127.108	15.000	31.000	65.000
Series-A funding ($M)	85,666	19.638	61.000	5.975	10.990	20.000
Post	85,666	0.444	0.497	0.000	0.000	1.000

Table 1, continued

Panel B: Industry Distribution

Sector	*Firms*	*% of firms*
Information Technology	1,762	41.40
Healthcare	1,199	28.20
Consumer Products and Services (B2C)	600	14.10
Business Products and Services (B2B)	379	8.90
Financial Services	196	4.60
Materials and Resources	81	1.90
Energy	35	0.80

Panel C: Correlation Matrix

Variable	*(1)*	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*	*(7)*	*(8)*	*(9)*
(1) Hiring Growth	1.000	1.000	-0.268	-0.254	-0.282	-0.295	-0.046	-0.122	0.056
(2) ΔLog(headcount)	0.994	1.000	-0.268	-0.254	-0.284	-0.296	-0.046	-0.123	0.055
(3) Funding / employee	-0.293	-0.297	1.000	0.768	0.891	0.708	-0.192	-0.176	0.438
(4) Valuation / employee	-0.278	-0.280	0.766	1.000	0.653	0.889	-0.145	-0.057	0.324
(5) Funding / net-new emp	-0.293	-0.294	0.898	0.658	1.000	0.805	-0.151	-0.253	0.362
(6) Valuation / net-new emp	-0.312	-0.311	0.712	0.896	0.812	1.000	-0.105	-0.136	0.273
(7) High GenAIExposure	-0.035	-0.031	-0.188	-0.145	-0.147	-0.103	1.000	0.050	-0.157
(8) Log(lagged headcount)	-0.143	-0.150	-0.197	-0.073	-0.265	-0.148	0.056	1.000	0.299
(9) Log(Series-A funding)	0.036	0.030	0.485	0.346	0.405	0.293	-0.141	0.308	1.000

Table 2. Leaner Hiring among GenAI-Exposed Startups

Table 2 examines whether higher-exposure startups slow headcount growth after the release of ChatGPT, using a differential-exposure difference-in-differences design. The dependent variable is *Hiring Growth*, the change in quarter-end headcount over four quarters (the twelve-month horizon) divided by the average of its current and four-quarters-prior values; the alternative outcome *Change in Log Headcount* is the log of quarter-end headcount minus the log four quarters earlier. Panel A reports the binary design, with key explanatory variable *High GenAIExposure* × *Post* (*High GenAIExposure* and *Low GenAIExposure* denote the top and bottom *GenAIExposure* terciles). Columns (1) to (4) report *Hiring Growth* without the lagged-headcount control, under the baseline that adds it, with *Log(Series A Funding)* × *Post* further added, and on a propensity-score-matched sample; Columns (5) to (8) repeat these for *Change in Log Headcount*. Panel B reports the continuous design, with key explanatory variable *GenAIExposure (1 sd)* × *Post*, under the same four specifications: without the lagged-headcount control, the baseline that adds it, the added funding control, and an entropy-balanced specification. All variables are defined in Appendix A. Statistical significance is denoted by *, **, and *** at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered at the firm level, with t-statistics reported in parentheses.

Panel A: Binary Exposure

	Hiring Growth				*ΔLog(headcount)*			
	No control	Baseline	Funding	PSM	No control	Baseline	Funding	PSM
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
High GenAIExposure × Post	-0.034**	-0.041***	-0.049***	-0.041***	-0.036**	-0.043***	-0.053***	-0.043***
	(-2.074)	(-2.761)	(-3.280)	(-2.742)	(-2.019)	(-2.679)	(-3.280)	(-2.657)
Log(lagged headcount)		-0.426***	-0.422***	-0.422***		-0.442***	-0.437***	-0.439***
		(-46.894)	(-46.558)	(-45.971)		(-45.440)	(-45.202)	(-44.503)
Log(Series-A funding) × Post			-0.031***				-0.039***	
			(-5.522)				(-6.456)	
Observations	56,658	56,658	56,658	54,947	56,658	56,658	56,658	54,947
Adjusted R-squared	0.497	0.603	0.605	0.601	0.479	0.583	0.585	0.581
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sector × quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cluster	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm

Table 2, continued

Panel B: Continuous Exposure Intensity

	Hiring Growth				*ΔLog(headcount)*			
	No control	Baseline	Funding	Entropy-bal.	No control	Baseline	Funding	Entropy-bal.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
GenAIExposure (1 sd) × Post	-0.016**	-0.019***	-0.023***	-0.016**	-0.017**	-0.021***	-0.025***	-0.018**
	(-2.533)	(-3.235)	(-3.807)	(-2.231)	(-2.477)	(-3.160)	(-3.817)	(-2.230)
Log(lagged headcount)		-0.416***	-0.412***	-0.412***		-0.435***	-0.431***	-0.433***
		(-56.521)	(-56.343)	(-53.296)		(-55.032)	(-54.927)	(-50.911)
Log(Series-A funding) × Post			-0.031***				-0.038***	
			(-6.797)				(-7.812)	
Observations	85,547	85,547	85,547	83,381	85,547	85,547	85,547	83,381
Adjusted R-squared	0.500	0.604	0.605	0.600	0.480	0.583	0.585	0.578
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sector × quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cluster	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm

Table 3. Capital-Labor Substitution among GenAI-Exposed Startups

Table 3 examines whether GenAI-exposed startups substitute capital for labor, operating more capital-intensively, after the release of ChatGPT. The dependent variables are four per-employee efficiency ratios, each in logs. *Funding (Valuation) per Employee* is cumulative funding (post-money valuation) divided by headcount. *Funding (Valuation) per Net-New Employee* is cumulative funding (post-money valuation) divided by net-new employees, the headcount added since the firm's Series A year. These specifications include firm and sector-by-quarter fixed effects but do not control for lagged headcount, which is the denominator of each ratio. Panel A reports the binary design, with key explanatory variable *High GenAIExposure* × *Post*; Panel B the continuous design, with key explanatory variable *GenAIExposure (1 sd)* × *Post*. All variables are defined in Appendix A. Statistical significance is denoted by *, **, and *** at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered at the firm level, with t-statistics reported in parentheses.

Panel A: Binary Exposure

	Funding / employee	*Valuation / employee*	*Funding / net-new emp*	*Valuation / net-new emp*
	(1)	(2)	(3)	(4)
High GenAIExposure × Post	0.049**	0.070**	0.085**	0.123***
	(2.031)	(2.510)	(2.507)	(3.358)
Observations	56,658	50,130	46,644	41,527
Adjusted R-squared	0.850	0.826	0.791	0.771
Firm FE	Yes	Yes	Yes	Yes
Sector × quarter FE	Yes	Yes	Yes	Yes
Cluster	Firm	Firm	Firm	Firm

Panel B: Continuous Exposure Intensity

	Funding / employee	*Valuation / employee*	*Funding / net-new emp*	*Valuation / net-new emp*
	(1)	(2)	(3)	(4)
GenAIExposure (1 sd) × Post	0.026**	0.037***	0.034**	0.049***
	(2.465)	(3.056)	(2.488)	(3.206)
Observations	85,547	75,962	69,362	61,969
Adjusted R-squared	0.846	0.822	0.786	0.767
Firm FE	Yes	Yes	Yes	Yes
Sector × quarter FE	Yes	Yes	Yes	Yes
Cluster	Firm	Firm	Firm	Firm

Table 4. Leanness versus the 2022 Funding Contraction

Table 4 examines whether the slower hiring and higher capital intensity of high-exposure startups reflect the concurrent 2022 venture funding contraction rather than a genuine shift to leaner operations. Columns (1) and (2) are a capital-side placebo, with fundraising as the outcome. In Column (1) the dependent variable is *Raised a round*, an indicator equal to one if the firm closed any financing round in the quarter; in Column (2) it is *Log(1 + $M raised)*, the log of one plus the millions of dollars raised in the quarter. In Columns (3) and (4), the dependent variable is *Hiring Growth*, the change in quarter-end headcount over four quarters divided by the average of its current and four-quarters-prior values. Column (3) is the baseline; Column (4) adds two pre-shock funding-environment controls interacted with *Post*: *Total Funding*, the log of one plus total capital raised as of the firm's last pre-2022Q4 quarter, and *Months Since Round*, months since the firm's last financing round at that quarter. The key explanatory variable throughout is *High GenAIExposure* × *Post*, where *High GenAIExposure* indicates the top *GenAIExposure* tercile and the bottom tercile is the control. All columns control for *Log Lagged Headcount* and include firm and sector-by-quarter fixed effects. All variables are defined in Appendix A. Statistical significance is denoted by *, **, and *** at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered at the firm level, with t-statistics reported in parentheses.

	Fundraising (capital-side placebo)		*Hiring Growth*	
	Raised a round	Log(1 + $M raised)	Baseline	Funding-winter controls
	(1)	(2)	(3)	(4)
High GenAIExposure × Post	-0.002	-0.008	-0.041***	-0.056***
	(-0.494)	(-0.561)	(-2.763)	(-3.749)
Total Funding × Post				-0.051***
				(-9.048)
Months Since Round × Post				0.000*
				(1.670)
Log(lagged headcount)	-0.029***	-0.112***	-0.422***	-0.406***
	(-11.573)	(-11.954)	(-46.065)	(-43.070)
Observations	56,658	56,658	55,178	55,178
Adjusted R-squared	-0.005	0.009	0.601	0.605
Firm FE	Yes	Yes	Yes	Yes
Sector × quarter FE	Yes	Yes	Yes	Yes
Cluster	Firm	Firm	Firm	Firm

Table 5. AI-Adoption Propensity and Leaner Hiring

Table 5 examines whether the leaner-hiring effect concentrates in firms where generative AI most plausibly substitutes for labor. The dependent variable is *Hiring Growth*, the change in quarter-end headcount over four quarters divided by the average of its current and four-quarters-prior values. The key explanatory variable is *High GenAIExposure* × *Post*, where *High GenAIExposure* indicates the top *GenAIExposure* tercile and the bottom tercile is the control. The table reports six columns, two per measure across three measures of a firm's AI-adoption propensity, each capturing a different facet of the firm: whether its product is AI, classified in an artificial-intelligence industry category with AI or machine-learning terms in its description (Columns 1 and 2); whether its senior leadership is AI-technical, the share of board and team members in AI or technical-leadership titles in the PitchBook academic records, split at the median (Columns 3 and 4); and whether it hires AI talent, having posted AI-related jobs before the release of ChatGPT (Columns 5 and 6). Within each pair, the first column is the lower-propensity subsample and the second the higher-propensity subsample. All variables are defined in Appendix A. Statistical significance is denoted by *, **, and *** at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered at the firm level, with t-statistics reported in parentheses.

	Dependent variable: Hiring Growth					
	AI product		AI-technical leadership		AI hiring	
	Lower	Higher	Lower	Higher	Lower	Higher
	(1)	(2)	(3)	(4)	(5)	(6)
High GenAIExposure × Post	-0.027	-0.103***	-0.001	-0.079***	-0.028	-0.105***
	(-1.600)	(-3.018)	(-0.033)	(-3.802)	(-1.263)	(-3.067)
Log(lagged headcount)	-0.422***	-0.444***	-0.402***	-0.446***	-0.432***	-0.368***
	(-41.348)	(-22.896)	(-30.678)	(-35.094)	(-32.182)	(-18.781)
Observations	46,372	10,266	27,972	26,918	24,851	10,527
Adjusted R-squared	0.607	0.589	0.600	0.611	0.616	0.655
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Sector × quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Cluster	Firm	Firm	Firm	Firm	Firm	Firm
p(Diff)		0.046		0.011		0.062

Table 6. Leaner Hiring across Worker Types

Table 6 examines which workers bear the hiring restraint, decomposing the headline effect by worker category. The dependent variable is a category-specific *Hiring Growth*: for each worker category, the change over four quarters in the count of workers in that category divided by the average of its current and four-quarters-prior values. Each specification controls for *Log Lagged Headcount* and includes firm and sector-by-quarter fixed effects. The table reports ten columns, two per dimension, across five worker dimensions: *Entry-Level* (age 25 and under) versus 26-and-over; *Remote-Suitable* (top tercile of role-level remote-suitability) versus *Onsite*; *Specialist* versus *Generalist* (role breadth, one minus the Herfindahl concentration of a worker's prior positions across role clusters, computed for senior hires only; *Generalist* is the top tercile of role breadth and *Specialist* the remainder), *Bachelor-* versus *Graduate*; and *Female* versus *Male*. The key explanatory variable is *High GenAIExposure* × *Post*, where *High GenAIExposure* indicates the top *GenAIExposure* tercile and the bottom tercile is the control. All variables are defined in Appendix A. Statistical significance is denoted by *, **, and *** at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered at the firm level, with t-statistics reported in parentheses.

	Age		*Remote*		*Generality*		*Education*		*Gender*	
	Entry (<=25)	Age 26+	Remote-suitable	Onsite	Specialist	Generalist	Bachelor-	Graduate	Female	Male
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
High GenAIExposure × Post	-0.119***	-0.050***	-0.068***	-0.033*	-0.040**	0.017	-0.042**	-0.023	-0.049**	-0.042***
	(-3.600)	(-2.871)	(-3.693)	(-1.954)	(-2.238)	(0.673)	(-2.272)	(-1.108)	(-2.521)	(-2.705)
Log(lagged headcount)	-0.512***	-0.323***	-0.320***	-0.384***	-0.320***	-0.334***	-0.398***	-0.313***	-0.419***	-0.332***
	(-24.385)	(-31.301)	(-28.768)	(-38.732)	(-28.862)	(-21.688)	(-35.473)	(-25.312)	(-36.020)	(-35.705)
Observations	50,909	56,601	56,253	56,611	56,418	54,358	56,175	55,965	56,230	56,640
Adjusted R-squared	0.290	0.485	0.420	0.539	0.441	0.339	0.471	0.387	0.459	0.522
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sector × quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cluster	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm
p(Diff)		0.017		0.038		0.031		0.288		0.655

Table 7. Firm Formation: Startups versus New Public Firms

Table 7 examines whether the leaner-hiring response is specific to new firms and what drives the differing response of established firms. The comparison group is new public firms, Compustat firms that completed an initial public offering between 2010 and 2021 in the same sectors as the startups and overlapping their size range. Exposure is fixed at a pre-determined workforce, the Series A year for startups and the 2019 workforce for new public firms. Panel A reports the cross-population comparison, with eight columns (two per outcome, startups then new public firms) and key explanatory variable *High GenAIExposure* × *Post*, where *High GenAIExposure* indicates the top *GenAIExposure* tercile and the bottom tercile is the control. *Hiring Growth* is the change in quarter-end headcount over four quarters divided by the average of its current and four-quarters-prior values; the composition margins apply the same growth rate to the count of entry-level (age 25 and under), remote-suitable (top tercile of role-level remote-suitability), and specialist (bottom two-thirds of senior hires by career role-breadth) workers. Panel B restricts the sample to new public firms, splits them at the median of pre-shock cash-to-assets into cash-rich and cash-lean firms. *Capex per Employee* is the log of one plus single-quarter capital expenditure per employee, de-cumulated from the year-to-date Compustat series. All variables are defined in Appendix A. Statistical significance is denoted by *, **, and *** at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered at the firm level, with t-statistics reported in parentheses.

Panel A: Hiring Response

	Hiring Growth		*Entry-level*		*Remote-suitable*		*Specialist*	
	Startups	Public	Startups	Public	Startups	Public	Startups	Public
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
High GenAIExposure × Post	-0.041***	0.069**	-0.119***	0.135**	-0.068***	0.070**	-0.040**	0.071**
	(-2.761)	(2.113)	(-3.600)	(2.229)	(-3.693)	(2.204)	(-2.238)	(2.021)
Log(lagged headcount)	-0.426***	-0.288***	-0.512***	-0.303***	-0.320***	-0.263***	-0.320***	-0.270***
	(-46.894)	(-16.528)	(-24.385)	(-6.513)	(-28.768)	(-14.410)	(-28.862)	(-13.980)
Observations	56,658	11,714	50,909	11,034	56,253	11,711	56,418	11,706
Adjusted R-squared	0.603	0.545	0.290	0.311	0.420	0.480	0.441	0.502
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sector × quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cluster	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm
p(Diff)		0.002		0.000		0.000		0.005

Table 7, continued

Panel B: Capital Access

	Hiring Growth		*Capex / Employee*	
	Cash-rich	Cash-lean	Cash-rich	Cash-lean
	(1)	(2)	(3)	(4)
High GenAIExposure × Post	0.117**	0.026	0.278**	0.085
	(2.145)	(0.459)	(2.509)	(1.001)
Log(lagged headcount)	-0.278***	-0.369***		
	(-11.125)	(-10.230)		
Observations	5,316	4,278	4,146	3,218
Adjusted R-squared	0.555	0.540	0.539	0.656
Firm FE	Yes	Yes	Yes	Yes
Sector × Quarter FE	Yes	Yes	Yes	Yes
Cluster	Firm	Firm	Firm	Firm
p(Diff)		0.242		0.166

Table 8. Robustness to Alternative Exposure Measures

Table 8 reports robustness of the leaner-hiring effect. The dependent variable is *Hiring Growth*, the change in quarter-end headcount over four quarters divided by the average of its current and four-quarters-prior values. Panel A re-estimates the baseline specification with alternative occupation-level measures, each aggregated to the firm on an employee-weighted basis and fixed at the firm's Series A year. The four columns are alternatives to the paper's primary suitability-for-machine-learning measure (Brynjolfsson, Mitchell, and Rock 2018): *LM-AIOE* (language-model occupational exposure; Felten, Raj, and Seamans 2023), *Eloundou β (human)* and *Eloundou β (GPT-4)* (generative-AI task exposure under the same rubric, rated by humans and by GPT-4, respectively; Eloundou et al. 2024), and the *Anthropic Index* (Handa et al. 2025). Panel B examines whether the effect reflects the 2022 monetary tightening, the end of the zero-rate era, rather than generative AI. *2021 Hiring Growth* is the firm's average twelve-month hiring growth over the four quarters of 2021. *Tech-Occ Share* is the share of the firm's workforce in computer and mathematical occupations (SOC major group 15) at the Series A year. All variables are defined in Appendix A. Statistical significance is denoted by *, **, and *** at the 10%, 5%, and 1% levels, respectively. Standard errors are clustered at the firm level, with t-statistics reported in parentheses.

Panel A: Alternative Exposure Measures

	Hiring Growth			
	LM-AIOE (2023)	Eloundou β, human	Eloundou β, GPT-4	Anthropic Index
	(1)	(2)	(3)	(4)
High GenAIExposure × Post	-0.036***	-0.048***	0.005	-0.042***
	(-2.578)	(-3.017)	(0.311)	(-2.754)
Log(lagged headcount)	-0.415***	-0.419***	-0.419***	-0.429***
	(-45.791)	(-47.516)	(-46.794)	(-48.007)
Observations	56,897	56,981	56,431	56,740
Adjusted R-squared	0.603	0.609	0.602	0.605
Firm FE	Yes	Yes	Yes	Yes
Sector × quarter FE	Yes	Yes	Yes	Yes
Cluster	Firm	Firm	Firm	Firm

Table 8, continued

Panel B: Ruling out the 2022 Monetary Tightening

	Hiring Growth			
	Baseline	2021 boom	Tech-share only	Tech-share
	(1)	(2)	(3)	(4)
High GenAIExposure × Post	-0.041***	-0.029**		-0.029*
	(-2.761)	(-2.027)		(-1.853)
2021 hiring growth × Post		-0.158***		
		(-14.667)		
Tech-occ share × Post			0.030***	0.027***
			(4.203)	(3.599)
Log(lagged headcount)	-0.426***	-0.329***	-0.423***	-0.423***
	(-46.894)	(-31.466)	(-46.078)	(-46.134)
Observations	56,658	50,368	54,949	54,949
Adjusted R-squared	0.603	0.599	0.606	0.606
Firm FE	Yes	Yes	Yes	Yes
Sector × Quarter FE	Yes	Yes	Yes	Yes
Cluster	Firm	Firm	Firm	Firm