Firm-Specific Capital and Corporate Financing: Evidence from Generative AI

Iftekhar Hasan, Yunying Huang, David Matsa, and Buhui Qiu*

August 2026

Abstract

How does labor-substituting technology affect corporate debt capacity? Examining generative AI adoption, we show that firms' capital structure choices depend on the form of productive capital that replaces labor. We identify adoption from job postings for roles that build and integrate large-language-model systems at U.S. public firms and estimate staggered difference-in-differences models. Unlike conventional automation, which is associated with higher leverage, generative AI can shift production toward less redeployable capital. We find that following adoption, firms reduce leverage and raise interest coverage, with larger effects among firms more exposed to AI-substitutable labor and firms that already rely heavily on intangible and organizational capital. Deleveraging is strongest when the AI implementation is closely tied to the firm's own business, where the resulting productive capital less transferable and redeployable.

Keywords: Organizational Capital, Generative AI Adoption, Labor Restructuring, Corporate Leverage

JEL Codes: G32, O33, J23

*Iftekhar Hasan is affiliated with Gabelli School of Business, Fordham University; Email: ihasan@fordham.edu. Yunying Huang is affiliated with Gabelli School of Business, Fordham University; Email: yhuang374@fordham.edu. David Matsa is affiliated with Kellogg School of Management, Northwestern University; Email: dmatsa@kellogg.northwestern.edu. Buhui Qiu is affiliated with University of Sydney Business School, The University of Sydney; Email: buhui.qiu@sydney.edu.au. We thank Isaac Pan and seminar participants at the University of Sydney for helpful comments. All errors are our own.

1 Introduction

What happens to corporate debt capacity when technology replaces labor? A natural answer is that debt capacity should increase. Labor is difficult to collateralize, creates fixed operating commitments, and can make financial distress more costly, whereas many earlier forms of automation replace workers with machinery and other tangible assets that can be pledged to creditors. Consistent with this intuition, a large literature shows that labor-market frictions shape firms' financing choices (e.g., Titman, 1984; Bronars & Deere, 1991; Berk et al., 2010; Matsa, 2010; Bae et al., 2011; Agrawal & Matsa, 2013; Simintzi et al., 2015; Serfling, 2016; Campello et al., 2018; Matsa, 2018; Favilukis et al., 2020; Kim, 2020), and recent evidence suggests that conventional workplace automation can expand debt capacity (Bates et al., 2025; Cheng et al., 2025). This reasoning, however, leaves open an important question: does labor substitution still increase debt capacity when the technology replacing workers does not take the form of tangible, redeployable capital?

The answer depends critically on *what replaces the labor*. A central principle of capital structure is that productive assets support more borrowing when they can be pledged, separated from the firm, and redeployed in financial distress (Myers, 1977; Titman, 1984; Rampini & Viswanathan, 2013). Asset tangibility is accordingly among the most reliable empirical determinants of leverage (Rajan & Zingales, 1995; Frank & Goyal, 2009), and the secular rise of intangible capital has been linked to shrinking corporate debt capacity and greater reliance on internal liquidity (Falato et al., 2022; Li, 2025). Yet the conventional tangible-intangible distinction does not fully characterize the financing capacity of modern firms. Some intangible assets, including codified intellectual property, transferable software, and patents, may retain value outside the firm and can sometimes be pledged to creditors (Mann, 2018). Other productive assets derive their value from proprietary data, internal workflows, organizational routines, system integration, and tacit employee know-how. Such capital can be highly valuable in place but difficult to value, seize, sell, or redeploy when the organization is disrupted.

This distinction becomes especially important when technological change affects not only how much labor a firm uses, but also *where productive knowledge is embodied*. Generative artificial intelligence (Generative AI) provides a unique setting in which to study this transition. Firms generally do not create lasting value merely by obtaining access to a foundation model: the same models and application interfaces are often available to many firms on broadly comparable terms. Instead, firms create value by integrating Generative AI with proprietary data, internal systems, established procedures, and organizational know-how. In doing so, they can shift tasks previously performed by employees into AI-enabled workflows, data pipelines, decision processes, and internal routines. Generative AI thus can substitute for some forms of cognitive labor while reorganizing productive knowledge into capital whose value remains deeply tied to the adopting organization.

The financing consequences of this reorganization are not obvious. If Generative AI simply codifies employee knowledge into transferable software or broadly reusable AI agents, then moving productive knowledge away from human labor could make the firm's productive assets *more* separable from employees and potentially expand debt capacity. By contrast, if adoption reorganizes knowledge into customized workflows, proprietary data systems, internal routines, and specialized implementation know-how that cannot readily be transferred to another organization, then the resulting productive capital may be less redeployable than the labor it replaces. A creditor can repossess machinery, inventory, or real estate, but cannot readily operate a customized retrieval system, reconstruct an internal data pipeline, or reproduce a decision process built around firm-specific data and routines. In this case, replacing labor might not increase debt capacity and may instead reduce it.

The role of employees also changes rather than disappears. Organizational capital is not fully alienable from the individuals who create, adapt, and operate it (Eisfeldt & Papanikolaou, 2013). As some codifiable tasks migrate from workers into AI-enabled processes, productive knowledge becomes increasingly concentrated in the systems and specialized employees that connect the technology to the firm's proprietary data, workflows, customers,

and internal operations. Concerns about financial distress can make such employees more difficult to retain, discourage investment in firm-specific human capital, and weaken the firm's ability to preserve organizational know-how. Workers demand compensation for unemployment risk (Berk et al., 2010; Agrawal & Matsa, 2013), leverage can undermine the job security needed to sustain firm-specific human-capital investment (Jaggia & Thakor, 1994), and employee mobility can increase distress costs when organizational know-how leaves with key personnel (Donangelo, 2014; Klasa et al., 2018). The relevant financing question is thus not simply whether Generative AI reduces labor, but how adoption reorganizes productive knowledge between labor, technology, and firm-specific organizational capital.

Emerging evidence makes this question particularly timely. Generative AI can substitute for human labor along specific margins, including junior hiring, externally contracted digital work, and employment in highly exposed occupations, although evidence that adoption reduces firms' aggregate employment remains mixed (Demirci et al., 2025; Hosseini Maasoum & Lichtinger, 2025; Humlum & Vestergaard, 2026). The technology reorganizes labor demand by automating or augmenting particular cognitive tasks, shifting the form in which productive knowledge resides. We examine the how these changes affect firms' debt capacity.

Studying this relationship requires distinguishing organizational implementation from casual use of publicly available Generative AI tools. We measure firm-level Generative AI adoption using hiring for positions that build, fine-tune, integrate, deploy, or operate large-language-model systems within the organization. Using more than 350 million LinkedIn job postings from U.S. public firms in Revelio Labs, we first identify postings containing Generative AI terminology and then use a large language model to distinguish "integrator" positions from roles that merely list Generative AI as a desired skill or general-purpose tool. A firm is classified as an adopter in the quarter of its first integrator posting. The resulting measure identifies 642 adopting firms between 2020Q1 and 2025Q2 and captures a deliberate organizational decision to commit resources to implementation rather than generic exposure to the technology.

We combine this measure with a staggered difference-in-differences (DiD) design to examine how firms adjust capital structure following adoption. We find that Generative AI adoption reduces corporate leverage and improves debt-service capacity. Adopting firms reduce book leverage by 1.3 to 1.7 percentage points, approximately a 5% decline relative to the sample mean of 33%. Interest coverage increases by approximately 17% to 20%. The estimates remain stable when year-quarter fixed effects are replaced with industry-by-year-quarter fixed effects, indicating that the results reflect firm-level financing responses rather than broader industry trends in deleveraging.

Dynamic event-study estimates show no evidence of differential pre-adoption trends. The leverage reduction and interest-coverage improvement emerge gradually over subsequent quarters, consistent with firms reoptimizing their capital structures as implementation progresses and the organization of productive capital changes. The results are robust across alternative leverage definitions, including market leverage, long-term leverage, and net debt to assets. They also scale with adoption intensity. Moving from the 25th to the 75th percentile of cumulative Generative AI integrator hiring, corresponding to an increase from 2 to 36 postings, reduces leverage by an additional 2.0 percentage points and raises interest coverage by approximately 16.5%. For comparison, Kim (2020) finds that large manufacturing plant openings, typically involving thousands of workers and substantial tangible investment, increase leverage by 2.5 to 3.9 percentage points.

The findings survive an extensive set of identification and robustness tests. We obtain similar estimates using modern DiD estimators designed for staggered treatment timing and heterogeneous treatment effects, sensitivity analysis allowing meaningful departures from parallel trends, entropy balancing, stacked-cohort matching, synthetic difference-in-differences, pre-treatment placebo tests, and a treatment definition requiring that postings be followed by realized hires. Across these approaches, adopting firms are not on differential leverage trends before adoption, and the post-adoption effects remain economically and statistically significant.

We also exploit two sources of quasi-experimental variation in adoption. The first is a shift-share instrument that interacts industries' predetermined occupational composition with the evolving alignment between occupational tasks and the capabilities of newly released large language models. As the frontier of language-model capabilities expands, industries employing more workers in occupations whose tasks align with these new capabilities experience greater incentives to adopt. The second instrument uses firms' geographic proximity to university AI research, which affects access to relevant human capital following the diffusion of Generative AI. Both instruments produce strong first stages, and the resulting estimates again indicate that Generative AI adoption reduces leverage and improves interest coverage.

We next establish that Generative AI adoption is accompanied by meaningful workforce restructuring in our setting. We use Worker Adjustment and Retraining Notification (WARN) Act filings, which provide a concrete record of substantial workforce reductions. Generative AI adoption increases the probability that a firm files a WARN notice and raises the number of workers affected. Under specifications with firm and industry-by-year-quarter fixed effects, adopting firms are 1.7 percentage points more likely to file a WARN notice, relative to a sample mean of 2.1%. Adoption also increases the probability of a mass layoff affecting at least 500 workers by 0.7 percentage points, while having no significant effect on plant closings.

The workforce adjustment response is strongest precisely where Generative AI has greater potential to substitute for existing tasks. We construct an ex-ante firm-level exposure measure based on the share of the firm's pre-ChatGPT workforce employed in occupations whose tasks large language models can perform. Among high-exposure firms, adoption increases the probability of a WARN filing, the probability of a mass layoff, and the number of workers affected; low-exposure firms show no comparable response. Importantly, the financing response exhibits the same heterogeneity. High-exposure firms reduce leverage by approximately 2.2 percentage points after controls, whereas low-exposure firms show no significant leverage reduction once controls are included. Thus, the same firms in which Generative AI

has greater scope to reorganize labor also exhibit the strongest financing response.

The leverage reduction is also concentrated among firms that rely more heavily on labor before AI adoption. Firms with more employees relative to physical capital reduce leverage by 2.1 percentage points after adoption, while more capital-intensive firms show no significant response. Firms with lower sales per employee reduce leverage by 2.6 percentage points, whereas firms with higher sales per employee exhibit no leverage reduction. Firms with higher pre-adoption hiring rates reduce leverage by 2.8 percentage points, while low-hiring firms again show no significant effect. A separate measure based on the share of knowledge workers yields similar results. Taken together, the WARN, exposure, and labor-intensity evidence indicates that the financing response is strongest where Generative AI adoption represents a more consequential reorganization of cognitive labor and productive knowledge.

If firms reduce leverage to protect their organizational capital, then we expect larger responses when the Generative AI implementation is more closely tied to the firm's own business. When implementation is built around the firm's own products, operations, data, and internal processes, the resulting organizational capital is difficult to separate from the organization and redeploy elsewhere, and it would therefore be more likely to lose value in financial distress.

To measure this firm specificity, we compare the descriptions of integrator postings in the adoption quarter with the firm's pre-adoption 10-K Item 1 business description. A higher similarity score indicates that the firm is recruiting employees to develop applications closely connected to its own products, operations, or internal processes rather than broadly transferable AI capabilities. We construct the measure using both sentence-transformer embeddings and TF-IDF term vectors. Firms with more firm-specific adoption reduce leverage by approximately 2.2 percentage points and increase interest coverage by 26% to 31%. The effects among firms with more general adoption are substantially weaker and, under the sentence-embedding measure, statistically insignificant. The financing response is thus strongest when implementation is most closely tied to the firm's own business and least likely to retain its

value outside the organization.

We then examine broader measures of firms' reliance on intangible and organizational capital. Consistent with the same mechanism, the leverage reduction is concentrated among firms whose productive bases are already more intangible. High-R&D firms reduce leverage by 2.8 percentage points after adoption and increase interest coverage by approximately 24.5%, whereas low-R&D firms exhibit no significant response. Similar patterns arise when intangible intensity is measured using selling, general, and administrative expenditures (SGA), which capture organizational capital, brand investment, and employee training. High-SGA firms reduce leverage by 2.4 percentage points and increase interest coverage by approximately 34.0%, while low-SGA firms show no effect. Patenting firms also deleverage more than non-patenting firms. These results provide broader corroborating evidence that the financing consequences of Generative AI are greatest where productive activity relies heavily on intangible and organizational capital.

Firms' leverage response to Generative AI contrasts with their response to conventional workplace automation. Both conventional automation and Generative AI can substitute technology for labor, but they replace labor with different forms of productive capital. Conventional automation typically involves machinery and other physical assets that can be pledged to creditors and redeployed outside the firm. Using an occupation-based measure of the substitutability of labor with automated capital, we find that conventional automation exposure is positively associated with leverage and negatively associated with cash holdings, consistent with prior evidence that automation increases debt capacity (Bates et al., 2025; Cheng et al., 2025). When conventional automation exposure and Generative AI adoption are included in the same regression, their financing effects have opposite signs and the Generative AI estimate is essentially unchanged. This contrast shows that labor substitution itself does not determine the direction of the financing response. What matters is the nature of the productive capital that replaces the labor.

Our findings contribute to several strands of literature. First, we contribute to research

on how capital structure relates to the nature of productive assets. Existing work establishes that tangible assets support borrowing and that the rise of intangible investment has reduced corporate debt capacity (Rajan & Zingales, 1995; Frank & Goyal, 2009; Rampini & Viswanathan, 2013; Falato et al., 2022; Li, 2025). Our analysis points to an additional dimension of asset composition: the extent to which productive capital is separable from the organization, redeployable by creditors, and operable independently of the organizational knowledge that makes it productive. Existing research largely emphasizes how much intangible capital a firm owns. We emphasize how that capital is organized. Two firms can hold similar amounts of intangible capital yet support different amounts of debt if one firm's productive knowledge is more deeply embedded in proprietary workflows, internal systems, and organizational routines.

We show that firms reduce leverage the most when their Generative AI implementation is tied most closely to the firm's own business, indicating that the financing consequences of intangible investment depend on firm specificity and redeployability. In this sense, our results refine the usual tangible-intangible distinction: intangible capital that is codified and transferable need not have the same financing consequences as organizational capital whose value is embedded in firm-specific systems, routines, and know-how.

Second, we contribute to the literature on labor and capital structure. Prior work shows that unemployment risk, wage rigidity, employment protection, bargaining power, and local labor-market conditions affect financing choices (e.g., Titman, 1984; Bronars & Deere, 1991; Berk et al., 2010; Matsa, 2010; Bae et al., 2011; Agrawal & Matsa, 2013; Simintzi et al., 2015; Serfling, 2016; Campello et al., 2018; Matsa, 2018; Favilukis et al., 2020; Kim, 2020). We highlight a different interaction between labor and financing: the effect of labor substitution depends on the type of productive capital that replaces labor. Replacing labor with tangible capital can increase debt capacity, whereas replacing cognitive labor with firm-specific workflows, data systems, routines, and specialized know-how can have the opposite effect. Our evidence thereby connects labor-market frictions to the asset side of the balance sheet

and helps explain why apparently similar reductions in labor can generate opposite financing responses.

Third, we contribute to the rapidly growing literature on the economic effects of artificial intelligence. Existing research studies productivity, labor demand, innovation, seniority, market value, and firm growth (e.g., Eisfeldt et al., 2023; Noy & Zhang, 2023; Babina et al., 2024; Eloundou et al., 2024; Brynjolfsson et al., 2025; Hampole et al., 2025; Hosseini Maasoum & Lichtinger, 2025). We show that firm-level Generative AI adoption also changes financing policy–a core corporate-finance outcome. Unlike much of the existing research, we focus on adoption, not ex-ante exposure. Exposure captures whether a firm's existing task mix is susceptible to Generative AI, whereas our measure captures the firm's deliberate decision to commit resources to organizational implementation. A firm may be highly exposed but delay implementation, or adopt aggressively despite moderate exposure. By combining realized adoption with workforce restructuring and capital-structure outcomes, we show how the reorganization of production associated with Generative AI extends beyond productivity and labor demand to firms' financing choices.

More broadly, our results show that technological change is not a monolithic determinant of capital structure. The same substitution away from labor can increase or decrease debt capacity depending on what replaces it. Technologies that shift production toward tangible, transferable, and collateralizable assets may support more borrowing. Technologies that reorganize productive knowledge around firm-specific workflows, proprietary data, internal systems, and organizational know-how may support less. Generative AI thus illustrates a broader principle: to understand how technological change affects financing, one must consider not only whether technology replaces labor or raises productivity, but also where the firm's productive knowledge resides after adoption and whether its redeployable outside of the organization.

The rest of the paper proceeds as follows. Section 2 describes the data and measurement of Generative AI adoption. Section 3 presents the main capital-structure results and

identification tests. Section 4 examines workforce restructuring, the settings in which labor substitution is more consequential, firm-specific and intangible organizational capital, and the contrast with conventional automation. Section 5 concludes. Appendix Table A1 provides detailed variable definitions. The Internet Appendix reports additional analyses and supplementary results and provides details on the construction and validation of the Generative AI adoption measure, the instrumental variables, and the Generative AI exposure measure.

2 Data and Measurement

2.1 Data

We obtain workforce data from Revelio Labs, which compiles records from online professional profiles and job postings scraped from company career pages, job boards, and aggregators (Eisfeldt et al., 2023; Jiang et al., 2025; Kwan et al., 2026). We use two products. Job postings sourced from LinkedIn record the employer, job title, O*NET occupation code, location, salary, and full job description.[1] Individual position records compiled from online professional profiles capture employment spells including employer, job title, start and end dates, seniority level, and O*NET occupation code. Raw profile data overrepresent white-collar and urban workers; Revelio addresses this with sampling weights calibrated to Bureau of Labor Statistics employment totals, which we apply throughout. Revelio's company reference tables map subsidiaries to ultimate parents, so a firm's postings include both headquarters and subsidiary hiring activity.

Mass layoff data come from Worker Adjustment and Retraining Notification Act filings, which require U.S. employers with 100 or more employees to provide advance notice before plant closings or mass layoffs. Each filing records the employer name, facility loca-

[1] Where salary is not disclosed in the posting, Revelio predicts it using a model trained on visa application data, self-reported compensation, and posted salary ranges. See https://www.data-dictionary.reveliolabs.com/methodology.html for details.

tion, number of affected workers, and event type. Industry-occupation employment shares for the Bartik instrumental variable come from the BLS Occupational Employment and Wage Statistics survey, which reports employment counts by detailed occupation within each NAICS industry; we use annual releases lagged one year.

Quarterly financial data come from Compustat North America, daily stock returns from CRSP, linked through the CRSP/Compustat Merged database. Patent data come from Kogan et al. (2017), updated through 2024. Occupation-level Generative AI exposure scores are constructed from O*NET task descriptions classified by GPT-3.5 following Eloundou et al. (2024). Data on AI model capabilities come from Epoch AI's large-scale AI models database, which records each model's publication date, demonstrated capabilities, and application domain; we restrict to models in the Language or Multimodal domains. AI publication counts for the geographic proximity instrument come from the Dimensions database, restricted to education institutions. Firm headquarters locations come from SEC filing addresses.

2.2 Firm-level Generative AI Adoption

We use over 350 million LinkedIn job postings matched to U.S. public firms, out of 1.33 billion LinkedIn postings worldwide provided by Revelio Labs, to construct a firm-level measure of Generative AI adoption. Following Hampole et al. (2025) and Hosseini Maasoum and Lichtinger (2025), we identify adoption from job postings, which have become a standard source for measuring technology adoption at the firm level (Hershbein & Kahn, 2018; Babina et al., 2024). We first flag postings whose descriptions contain at least one of 27 Generative AI keywords, covering model names (ChatGPT, GPT, Claude, Gemini), technical infrastructure (LangChain, RAG, vector database), and implementation terms (prompt engineering, OpenAI API). We then classify each flagged posting using a large language model (OpenAI's GPT-4o-mini).[2] The classifier assigns each posting to one of four categories: *integrator* (the role involves developing, deploying, or operating Generative AI systems and applications

[2]Results are robust to classification using Meta's Llama-3.1-8b and Google's Gemini-2.5-flash. See Internet Appendix IA2.

within the firm), *user* (the role references Generative AI only among desired skills or as a general-purpose tool for routine tasks, without responsibilities for building such systems), *both*, or *neither*. The classification also reduces false positives where the matched keyword has a non-AI meaning in context. A posting that lists "skills with large language models" as a desired qualification tells us little about the firm's production choices. A posting that recruits someone to "build a retrieval-augmented generation pipeline" reveals that the firm is committing resources to embed Generative AI into its operations. Figure 1 shows a Generative AI integrator posting on LinkedIn, and Table IA16 in the Internet Appendix gives further examples of both categories. Only integrator postings enter our adoption measure. The full keyword list, classification prompt, model specifications, and robustness to alternative models and prompts are in Internet Appendix IA2.

[Insert Figure 1 about here]

Our adoption measure specifically includes firms whose first Generative AI integrator posting predates ChatGPT's November 2022 release, because earlier commercial generative models such as OpenAI's GPT-3 were already available and a small number of firms hired specialists to integrate them.,[3] We verify the release date of each Generative AI model and product in our keyword list, restricting each keyword to postings after the product's actual release date. Internet Appendix Table IA15 documents the release dates of each Generative AI technology in our keyword list. For robustness, we also construct an alternative adoption measure that considers only Generative AI job postings dated no earlier than ChatGPT's November 2022 release. Examples of pre-ChatGPT Generative AI job postings are displayed in Internet Appendix Table IA17.

A firm is a Generative AI adopter if it posts at least one integrator vacancy. The adoption

[3]OpenAI released GPT-3 in June 2020 as a generative language model accessible through its API. Subsequent Generative AI releases included GitHub Copilot (June 2022), built on OpenAI's Codex model. ChatGPT, launched in November 2022, was built on GPT-3.5 and provided the first public conversational interface, which made Generative AI accessible beyond software developers. Subsequent models include GPT-4 (March 2023), Anthropic's Claude (March 2023), Meta's Llama (February 2023), and Google's Gemini (December 2023).

date is the quarter of the firm's first integrator posting. The binary treatment indicator, *Generative AI Adoption*, equals one for all firm-quarters at and after the adoption quarter and zero otherwise. We also measure adoption intensity using two continuous variables: the natural logarithm of one plus cumulative annualized salary allocated to Generative AI postings since adoption, and the natural logarithm of one plus cumulative Generative AI postings since adoption.

Our firm-level Generative AI adoption measure differs from the exposure measure in Eisfeldt et al. (2023). Exposure quantifies the share of a firm's workforce tasks that large language models could automate, building on the task-based framework of Autor et al. (2003) and Eloundou et al. (2024). It is an ex-ante characteristic that depends on both technological capabilities and existing workforce composition before ChatGPT release in 2022Q4. Our adoption measure instead captures the firm's deliberate choice to hire Generative AI specialists. The two constructs can generate opposing coefficient signs without contradiction: exposure reflects market anticipation of AI's potential (Eisfeldt et al., 2023), while adoption reflects the realized costs of implementation, including workforce restructuring (Hosseini Maasoum & Lichtinger, 2025) and the intangible investment that our theoretical framework links to reductions in leverage.

2.3 Sample

Revelio Labs' LinkedIn posting data comprises over 1.33 billion job postings across all employers globally. We retain the 350 million postings from firms that Revelio Labs maps to the CRSP/Compustat Merged database through its company reference tables. Applying the Generative AI keyword filter to postings from these firms yields 541,331 matches, of which the language model classifier identifies 154,156 as integrator postings.

Our analysis panel spans 2020Q1 through 2025Q2, matching the coverage of Revelio Labs' job posting data.[4] We merge Revelio workforce data with Compustat quarterly funda-

[4]We use Revelio Labs' LinkedIn job posting records, which provide the largest single consistent source

mentals and exclude financial firms (SIC 60–67), utilities (SIC 49), and firms with zero total debt. The final sample contains 41,007 firm-quarter observations for 2,650 unique firms, of which 642 adopt Generative AI during the sample period. Generative AI adoption spans the economy broadly and is not confined to technology sectors. Figure IA2 in the Internet Appendix reports adoption rates by three-digit NAICS subsector: industries such as hospitals, accommodation, apparel manufacturing, educational services, and air transportation all exhibit adoption rates above 15%. Table 1 reports summary statistics.

[Insert Table 1 about here]

Figure 2 shows quarterly salary allocated to Generative AI roles and their share among technology postings. Apart from a small number of early Generative AI adopters for earlier models such as GPT-3, both series are close to zero before ChatGPT's (GPT-3.5) launch in November 2022 and rise steeply thereafter, with aggregate quarterly annualized salary peaking above $3 billion in 2024. Salary figures include Revelio-predicted compensation where reported salary is unavailable. Panel A of Figure 3 plots integrator postings by role. We focus on the cumulative count because each integrator hire continues to contribute to the firm's Generative AI capacity in subsequent quarters. Software engineers and data analysts account for the largest share, with roughly 61,600 and 31,100 cumulative postings by 2025Q2, followed by application engineers and scientists at about 11,100 and 9,500. Panel B shows the Generative AI tools named in job descriptions. GPT and ChatGPT are the most common, followed by Claude, prompt engineering, RAG, and Gemini, which emerge after mid-2023.

[Insert Figures 2 and 3 about here]

of full-text job descriptions for our textual classification pipeline. Revelio also offers another unified product (COSMOS) that aggregates postings from company career pages, job boards, and staffing firms beginning in 2021. We prefer the single-source data to avoid cross-platform variation in description format and completeness. Results are entirely unchanged when restricting the sample to the COSMOS coverage period.

3 Generative AI Adoption and Leverage

3.1 Difference-in-differences estimates

We estimate the effect of Generative AI adoption on corporate leverage using a staggered difference-in-differences design. The baseline specification is

$$Y_{i,t+1} = \beta \cdot \textit{Generative AI Adoption}_{i,t} + \boldsymbol{\gamma}'\mathbf{X}_{i,t} + \alpha_i + \delta_t + \varepsilon_{i,t}, \tag{1}$$

where i indexes firms and t indexes year-quarters. The dependent variable $Y_{i,t+1}$ is a leverage outcome measured one quarter ahead. We lead the outcome by one quarter to allow time for financing decisions to reflect the firm's changed investment composition and to avoid mechanical contemporaneity between the adoption event and balance sheet adjustments recorded in the same quarter. The key variable of interest, $\textit{Generative AI Adoption}_{i,t}$, equals one for all quarters at or after the quarter when firm i first posts a Generative AI integrator vacancy and zero otherwise. For firms that never adopt, the indicator is zero throughout. The coefficient β captures the average change in leverage following adoption, relative to not-yet-adopting and never-adopting firms in the same quarter.

Following the literature on labor market frictions and capital structure (e.g. Matsa, 2010; Agrawal & Matsa, 2013; Simintzi et al., 2015; Kim, 2020), our primary dependent variables are book leverage, total debt divided by total assets, and log interest coverage. Following Faulkender and Petersen (2006) and Agrawal and Matsa (2013), we examine interest coverage as the natural logarithm of one plus the ratio of operating income before depreciation to interest expense, coded as zero for negative operating income where the ratio is otherwise undefined. Book leverage measures the stock of debt on the balance sheet, and interest coverage measures the flow of earnings available to service that debt. We also examine market leverage, net leverage, and long-term leverage as alternative dependent variables. The control vector $\mathbf{X}_{i,t}$ includes standard capital structure determinants: firm size, asset

tangibility, market-to-book ratio, operating return on assets, and R&D. All specifications include firm and year-quarter fixed effects with standard errors clustered at the firm level. To account for time-varying industry conditions that may differentially affect both AI adoption and leverage, we also estimate specifications with industry-by-year-quarter fixed effects.

[Insert Table 2 about here]

Table 2 reports estimates of Equation (1). We find that Generative AI adoption is associated with a significant reduction in leverage. Column (1) shows that adoption reduces book leverage by 1.7 percentage points without additional controls. Relative to the sample mean of 33%, this represents approximately a 5% decline, statistically significant at the 1% level. Column (2) adds standard capital structure controls, and Column (3) replaces year-quarter fixed effects with industry-by-year-quarter fixed effects to absorb time-varying industry trends. The coefficient remains statistically and economically significant in both specifications, confirming that the leverage reduction reflects firm-level financing decisions rather than broader industry dynamics, including aggregate and industry-level macroeconomic shocks that may coincide with the adoption period. The interest coverage estimates further confirm this finding. Column (4) shows that adoption increases log interest coverage by 0.178, which is 11.2% of the sample mean and approximates a 17.8% increase in the interest coverage ratio, statistically significant at the 1% level.[5] The coefficient remains statistically and economically significant in Columns (5) and (6) after adding controls and industry-by-year-quarter fixed effects.

[5]The percentage of the sample mean is $0.178/1.594 = 11.2\%$, where 1.594 is the mean of log interest coverage from Table 1. We interpret the log-points coefficient as the approximate percentage change in interest coverage, $0.178 \approx 17.8\%$.

3.2 Dynamic Effects, Parallel Trends, and Matching

A. Dynamic effects

We estimate dynamic treatment effects around the adoption quarter to test whether adopting firms were on a differential leverage trajectory before adoption. The staggered nature of our setting strengthens this test: at any given quarter, the control group includes both never-adopters and not-yet-adopters who eventually adopt, so the comparison reflects variation in adoption timing among similar firms rather than permanent differences between adopters and non-adopters.

[Insert Figures 4 and 5 about here]

Figures 4 and 5 present these estimates for book leverage and interest coverage. The pre-adoption coefficients are small and statistically insignificant across all specifications, confirming the parallel trends assumption. Post-adoption, the leverage reduction builds gradually, reaching approximately 1.3 percentage points by the fifth quarter after adoption, consistent with the static estimate of 1.3 percentage points in Table 2. The interest coverage improvement follows a similar trajectory, with the effect growing to 0.165 by the fourth quarter. This gradual pattern is consistent with capital structure adjustment being slow: firms do not restructure their balance sheets immediately, and the full effect of changed investment composition on financing decisions takes several quarters to materialize. The monotonically increasing post-treatment effects also argue against reverse causality, as the leverage reduction follows rather than precedes adoption.

We employ a comprehensive set of estimators for staggered difference-in-differences to address concerns that heterogeneous treatment effects across adoption cohorts may bias the baseline estimates. Figures 4 and 5 also report estimates from four robust estimators: Callaway and Sant'Anna (2021), who aggregate group-time average treatment effects using not-yet-treated controls; Borusyak et al. (2024), who impute counterfactual outcomes for treated units; Sun and Abraham (2021), who reweight cohort-specific estimates to avoid

contamination across groups; and De Chaisemartin and d'Haultfoeuille (2024), who estimate treatment effects robust to heterogeneous dynamics. All five estimators produce qualitatively similar patterns: clean pre-trends and gradually increasing post-adoption effects.

B. Parallel trends

Our estimates remain economically and statistically significant under meaningful departures from parallel trends. Roth (2022) argues that insignificant pre-treatment coefficients alone cannot definitively validate the parallel trends assumption, as the standard pretest has limited power to detect violations large enough to explain away treatment effects. To address this concern, we implement the framework of Rambachan and Roth (2023), which constructs robust confidence intervals under progressively relaxed assumptions about post-treatment trend behavior. The key parameter $\bar{M}$ bounds the maximum post-treatment trend deviation as a fraction of the largest pre-treatment violation observed in the data, where $\bar{M} = 0$ corresponds to exact parallel trends and larger values permit increasingly severe departures. We focus on the average post-treatment effect, which aggregates across all post-treatment periods and provides a more stable assessment of overall impact than period-specific estimates. We implement this test using the imputation estimator of Borusyak et al. (2024), which provides efficient estimates with a full variance-covariance matrix suitable for this analysis.[6]

Figure IA4 in the Internet Appendix reports the results. The 95% robust confidence intervals widen as $\bar{M}$ increases from 0 to 0.5, but continue to exclude zero for both book leverage and interest coverage up to $\bar{M} = 0.4$. That is, even if post-treatment trend deviations were 40% as large as the worst pre-treatment violation, the estimated effects remain statistically significant. The consistency of these results under cautious assumptions about trend violations provides strong evidence that the leverage reduction and interest coverage improvement documented above represent genuine treatment effects of Generative AI adoption rather than artifacts of differential pre-existing trends.

[6]Results are qualitatively similar across all five estimators.

C. Matched sample estimates

We construct two matched samples that pair adopting firms with non-adopting firms that are observably similar. The first uses entropy balancing. At 2022Q3, the quarter before ChatGPT's public release, we reweight comparison firms to match adopting firms on size, tangibility, market-to-book, profitability, R&D, and book leverage. As shown in Table IA4 in the Internet Appendix, the adoption and comparison samples have near-exact balance: all standardized differences fall below 1%. The estimates from analysis of the matched sample are reported in Panel A of Table IA3. They are, if anything, larger than the unmatched baseline. Column (1) reports a 1.6 percentage point leverage reduction with mean and variance balancing, matching the baseline estimate from Equation (1) in Column (3) of Table 2. Columns (2), (4), (6), and (8) additionally balance on the four-quarter pre-treatment change in leverage, so that treated and control firms enter the post-period on identical leverage trajectories. Adding the leverage trend increases the estimate to 2.0 percentage points (Column 2). Interest coverage rises by 12.4% without trend balancing (Column 5) and by 13.7% with it (Column 6). That trend-balanced specifications produce similar or larger coefficients indicates pre-existing leverage dynamics do not drive the result.

Panel B reports a stacked-cohort propensity-score-matching design. For each adoption quarter, we match treated firms to never-adopters at the quarter before adoption using nearest-neighbor matching on the same six covariates, and in Columns (2), (4), (6), and (8) restrict matches to the same two-digit SIC industry. Subpanels spanning five quarters before through five quarters after adoption are stacked, with firm-by-cohort and cohort-by-year-quarter fixed effects absorbing cohort-specific shocks. The estimates strengthen as the comparison group narrows. Broad matching with five neighbors yields a 0.9 percentage point leverage reduction in Column (1); within-industry matching with a single nearest neighbor yields 2.7 percentage points in Column (4), at the 1% level. Interest coverage increases range from 11.7% to 17.0% across specifications. Tighter matching produces larger coefficients, confirming that the unmatched baseline is conservative.

3.3 Robustness

A. Intensive margin of adoption

We also estimate the intensive margin of Generative AI adoption, replacing the binary indicator in Equation (1) with two continuous measures that capture the scale of investment. The first is the natural logarithm of one plus cumulative Generative AI postings since adoption; the second is the natural logarithm of one plus cumulative annualized salary allocated to these postings.[7] Table 3 reports the results. Both measures are statistically significant at the 1% level across all specifications. In Panel A, a one standard deviation increase in log cumulative postings among adopter firm-quarters reduces book leverage by 0.9 percentage points, or 3% of the sample mean, and raises the interest coverage ratio by approximately 7.8%.[8] Moving from the 25th to 75th percentile of adoption intensity, corresponding to an increase from 2 to 36 cumulative postings, reduces leverage by 2.0 percentage points and raises interest coverage by approximately 16.5%. Panel B yields results of the same sign and significance when adoption is measured by cumulative salary investment. A one standard deviation increase among adopters reduces leverage by 0.2 percentage points and raises interest coverage by 2.5%; the interquartile range spans approximately $229,000 to $4.0 million in cumulative salary.[9] For comparison, Kim (2020) shows that large manufacturing plant openings, each involving thousands of new workers, increase leverage by 2.5 to 3.9 percentage points. This increase from 2 to 36 integrator hires produces a leverage reduction comparable

[7] As a further robustness check, we replace the cumulative count with a stock that depreciates as workers turn over, $\text{stock}_t = (1-\delta)\,\text{stock}_{t-1} + \text{postings}_t$. In untabulated results using annual depreciation rates of 15%, the organization capital rate of Eisfeldt and Papanikolaou (2013), 20%, and 30%, close to the quit rate for professional and business services in the BLS Job Openings and Labor Turnover Survey, the estimates are essentially unchanged.

[8] Panel A coefficients under industry-by-year-quarter fixed effects: −0.008 (leverage) and 0.066 (interest coverage). Among adopter firm-quarters, the standard deviation of log cumulative postings is 1.18, and the interquartile range is 2.50 (from $\ln(1+2) = 1.10$ to $\ln(1+36) = 3.60$). One standard deviation effect on leverage: $0.008 \times 1.18 = 0.009$. Interquartile-range effect: $0.008 \times 2.50 = 0.020$. Interest coverage: $0.066 \times 1.18 \approx 0.078$ and $0.066 \times 2.50 \approx 0.165$.

[9] Panel B coefficients under industry-by-year-quarter fixed effects: −0.001 (leverage) and 0.013 (interest coverage). Among adopter firm-quarters, the standard deviation of log cumulative salary is 1.89, and the interquartile range is 2.86. The smaller intensive-margin effects reflect that cumulative salary captures both hiring scale and wage-level differences across firms.

in magnitude, consistent with Generative AI's capacity to substitute for a broad range of tasks at low marginal cost.

[Insert Table 3 about here]

B. Alternative leverage measures

Our findings are robust to alternative leverage definitions. Table IA1 reports results for three alternative measures: market leverage, defined as total debt over total debt plus market equity; long-term leverage, defined as long-term debt over total assets; and net debt to assets, defined as total debt minus cash over total assets. All three measures show statistically significant reductions at least at the 5% level across all specifications. For example, Column (2) of Panel B shows that treated firms reduce market leverage by 2.3 percentage points relative to control firms after Generative AI adoption. Long-term leverage shows a similar effect. By isolating long-term debt, this measure confirms that the capital structure adjustment reflects permanent financing decisions rather than transitory changes in short-term borrowing, with a 1.4 percentage point reduction under industry-by-year-quarter fixed effects. For net debt to assets, which accounts for both the debt and cash sides of the balance sheet, Column (6) of Panel B estimates a 2.6 percentage point reduction for treated firms in the post-adoption period. All estimates remain significant across both the baseline and industry-adjusted specifications, indicating that the deleveraging effect of Generative AI adoption does not depend on a particular measurement choice.

C. Measurement and specification alternatives

Our results are robust to the choice of large language model used to classify job postings. We replicate the baseline classification with two alternative LLMs: Meta's Llama-3.1-70B and Google's Gemini-2.5-Flash. We also restrict the baseline GPT-4o-mini classification to postings published from 2022Q4 onward, to align with the economic shock of ChatGPT's public release used in Eisfeldt et al. (2023). Figure IA5 in the Internet Appendix reports coefficient estimates and 95% confidence intervals for book leverage and interest coverage

in Panels A and B, and for three alternative leverage measures in Panels C through E. All three classification approaches produce point estimates that are statistically significant and economically close to the baseline, with overlapping confidence intervals throughout. We also estimate event studies for each alternative classification under all five heterogeneity-robust estimators used in the main analysis. Figure IA6 presents Borusyak et al. (2024) as an example; other event study estimators yield similar patterns. Pre-treatment estimates for all three models cluster near zero for both book leverage and interest coverage. Post-treatment, the dynamic paths diverge from zero in the expected direction and track one another closely across all event-time horizons, confirming that neither the magnitude nor the timing of the estimated effects depends on which LLM classifies the postings.

In Columns (3) and (6) of Table 2, we absorb time-varying industry shocks with SIC two-digit industry-by-year-quarter fixed effects. We replace this classification with three-digit NAICS and Fama-French 48 alternatives and find that results remain economically and statistically similar. Figure IA3 in the Internet Appendix reports coefficients and 95% confidence intervals for book leverage and interest coverage in Panels A and B, and for three alternative leverage measures in Panels C through E.

The estimates are also robust to retaining firms with zero total debt in the sample. Table IA2 in the Internet Appendix reports the baseline regressions on the augmented sample; the leverage and interest coverage coefficients remain statistically and economically similar across all specifications.

D. Synthetic Difference-in-Differences

We implement the synthetic difference-in-differences (SDID) estimator of Arkhangelsky et al. (2021). SDID reweights control firms and pre-treatment periods so that the weighted average tracks each adopter's pre-treatment leverage trajectory. In practice, SDID matches each adopter's pre-treatment leverage trajectory using unit weights on never-adopters, and adjusts for common aggregate time trends using time weights that align pre-treatment and post-treatment outcomes within the control group. Two-way fixed effects remove level dif-

ferences across firms and quarters but cannot account for firm-specific trends. The SDID estimator accounts for these trends, so that only post-treatment deviations enter the estimated effect. The estimator requires a strongly balanced panel, and we restrict the sample to firms with complete quarterly coverage. Table IA5 in the Internet Appendix reports the results. All five leverage measures yield statistically significant coefficients, directionally consistent with the baseline, both with and without controls entered via the projected method. Our findings' robustness to matching pre-adoption leverage trajectories strengthens the causal interpretation of the estimates.

E. Pre-treatment placebo

We further test whether adopting firms were already deleveraging before adoption. We assign each adopting firm a placebo treatment four quarters before its actual first Generative AI posting and restrict the sample to the pre-adoption period. The actual treatment never enters. Never-adopting firms serve as controls throughout. Table IA6 in the Internet Appendix reports the results. All six placebo coefficients are economically small and statistically insignificant across both outcomes and all three specifications.

F. Realized hiring

We verify that the demand recorded in job postings reaches staffing by matching adopters to realized hires. A Generative AI hire is a new Revelio Labs position whose title or description contains one of the paper's keywords. Table IA7 in the Internet Appendix restricts treated firms to adopters who record such a hire in the adoption quarter or the following one, with never-adopters remaining the control group. The leverage and interest coverage effects hold on the restricted sample, and both point estimates are somewhat larger than on the full sample. Adoption identified from postings therefore reflects Generative AI capacity that firms actually build.

3.4 Instrumental Variable Estimates

We examine the causal effect of Generative AI adoption on capital structure using two instrumental variable strategies that exploit distinct sources of exogenous variation. The first is a Bartik-style shift-share instrument based on the evolving capabilities of frontier AI models; the second exploits geographic proximity to AI research universities. Each instrument isolates supply-side variation in firms' access to Generative AI, and the two rely on entirely different sources of cross-sectional and time-series variation, providing independent tests of the causal relationship.

A. LLM capability Bartik

Our first instrument follows the Bartik shift-share design (Bartik, 1991; Goldsmith-Pinkham et al., 2020). Large language models have evolved rapidly since the release of GPT-2 in 2019. Early models were limited to basic text generation; GPT-3.5 and GPT-4 added instruction-following and code generation; subsequent releases introduced web search, mathematical reasoning, and chain-of-thought capabilities; most recently, agentic systems such as Claude Code can execute multi-step software engineering workflows autonomously. Each generation expands the set of occupations whose tasks can be performed or augmented by LLMs. A model capable only of text generation is relevant primarily to writing-intensive occupations; one that also writes code shifts exposure toward software engineers; one that reasons over data and browses the web reaches analysts, researchers, and administrative professionals. The shift component of our instrument captures this time-varying expansion: as new models are released, the alignment between the LLM capability frontier and each occupation's task content changes. The share is the predetermined occupation composition of each firm's industry. The instrument assigns each firm the LLM capability exposure of its NAICS sector:

$$Z_{it}^{\text{Bartik}} = \sum_{s \in \mathcal{S}} \omega_{s,t-1}^{g(i)} \cdot A_{st} \tag{2}$$

where $\omega_{s,t-1}^{g(i)}$ is the lagged employment share of occupation s in firm i's three-digit NAICS industry from the prior year's BLS Occupational Employment and Wage Statistics survey, and A_{st} is the text similarity between O*NET occupation descriptions and the cumulative capability profile of large language models released up to quarter t.[10] Using lagged shares ensures the occupation composition is predetermined relative to the outcome. The instrument is standardized to zero mean and unit variance.

Two features support the exclusion restriction. First, the occupation shares come from BLS at the industry level, not from individual firms, and are lagged one year. A firm's own hiring decisions do not affect its instrument value; only the predetermined aggregate occupation structure of its sector determines the share component. Second, the time-varying alignment reflects the pace of frontier AI research determined by the global research community, not by any individual firm's adoption decision. Firm fixed effects absorb permanent differences across sectors, and year-quarter fixed effects absorb aggregate movements in alignment common to all firms.

Table 4 reports the results under two text-alignment specifications. The first stage is strong in both panels, with Kleibergen-Paap F-statistics of 36.5 and 33.1. The reduced-form and second-stage coefficients are negative for leverage and positive for interest coverage, qualitatively consistent with the OLS estimates.

The second-stage estimates are larger than the OLS estimates, and the difference follows from what each estimator averages over. The OLS coefficient reflects the average response across adopting firms. Two-stage least squares uses only the variation in adoption that the instrument generates, and the instrument moves adoption most where a firm's industry employs a larger share of workers in occupations whose tasks align with advancing LLM capabilities. Firms of this kind carry the heaviest weight in the second stage, which therefore recovers a local average treatment effect for them rather than that average.

[10] We measure A_{st} using TF-IDF cosine similarity between O*NET Detailed Work Activity descriptions and Epoch AI model capability metadata, restricting to Language and Multimodal models. Each panel in Table 4 varies the text fields used. Internet Appendix IA3.1 provides full construction details.

To make the comparison transparent, we report economic magnitudes using a one-standard-deviation increase in adoption for the OLS regressions and a one-standard-deviation increase in instrumented adoption for the two-stage least squares regressions.[11] Leverage falls by 0.34 percentage points under OLS and by 0.63 and 0.73 percentage points under the instrument, and interest coverage rises by 4.3% against 11.1% and 11.2%. The signs agree and the magnitudes are of the same order once expressed in standard-deviation terms.

Table IA8 in the Internet Appendix tests whether the results are driven by firms that develop large language models themselves. Panel A drops eight firms that have released proprietary foundation models or provide the primary computing infrastructure for LLM training: Alphabet, Amazon, Apple, IBM, Meta, Nvidia, Oracle, and Salesforce. Panel B drops all firms in software publishing (NAICS 511), computing infrastructure (NAICS 518), and web services (NAICS 519). The estimates are unchanged in both panels, supporting the exclusion restriction.

[Insert Table 4 about here]

B. Geographic proximity to AI research

Second, we follow Babina et al. (2024) using proximity to university AI research as an instrument for adoption. We specifically use geographic distance as the proximity measure rather than hiring networks, so that the instrument is independent of firms' labor market decisions. Firms headquartered near major AI research institutions face a larger local supply of Generative AI workers after ChatGPT's release, increasing adoption through a labor supply channel. We measure local AI research capacity as the log of total AI and computer science publications at nearby top-ranked US universities, using publication data from the Dimensions database over 2017–2022, and interact this cross-sectional measure with a post-ChatGPT indicator to generate the time-series variation needed for identification.[12]

[11] The adoption indicator has a standard deviation of 0.273; instrumented adoption has a standard deviation of 0.047 and 0.044 under the two alignments.

[12] We rank universities by publication count in leading AI and computer science venues and journals, restrict to education institutions to exclude corporate research laboratories, and retain the top-ranked insti-

A firm's headquarters location relative to research universities is determined by historical factors largely unrelated to its current financing decisions. Firm fixed effects absorb any time-invariant effect of location, and industry-by-year-quarter fixed effects absorb time-varying industry shocks, so identification comes from within-industry geographic variation in proximity to AI research.

Table IA9 in the Internet Appendix reports results for two variants that weight university research output differently. Both yield strong first stages, with Kleibergen-Paap F-statistics above 26, and the second-stage estimates again show that instrumented adoption lowers leverage and raises interest coverage, significant at the 5% level. That two instruments built from entirely different sources of variation produce the same qualitative conclusion strengthens the causal interpretation.

4 Mechanisms

In this section, we examine why Generative AI affects firms' capital structure choices. Like earlier waves of automation, firms replaced human labor with technology, but the nature of the capital embedded in that technology is very different. Whereas automation primarily substitutes machinery for manual labor, Generative AI can perform cognitive and language-intensive tasks previously requiring professional judgment (Trajtenberg, 2018; Eloundou et al., 2024). Its adoption thus reduces firms' reliance on some knowledge-worker tasks while reorganizing production around software, proprietary data, internal workflows, organizational routines, and the employees who develop and operate these systems.

We explore these issues in two stages. We first establish that Generative AI adoption is associated with workforce restructuring and that the financing response is stronger where there is greater scope for labor substitution. The evidence comes from layoffs, ex-ante Generative AI exposure, and several measures of labor intensity. We then ask why replacing

tutions within a fixed radius of each firm's headquarters. Results are robust to alternative ranking cutoffs and distance thresholds. Internet Appendix IA3.2 provides the full construction details, university rankings, and excluded corporate laboratories.

labor reduces rather than increases debt capacity. If the financing effect depends on the nature of the capital that replaces labor, deleveraging should be stronger when that capital is more firm specific, less transferable, and less redeployable. We test this first across adopters whose Generative AI implementation differs in how closely it is tied to their own business, and then across firms that differ in their intangible and organizational-capital intensity. A final comparison with conventional workplace automation shows that the direction of the financing response depends on the type of productive capital that replaces labor rather than on labor substitution itself.

4.1 Labor

Labor market frictions have long been linked to capital structure (Matsa, 2018). Workers who bear greater unemployment risk under high debt may demand compensating wage differentials (Titman, 1984; Berk et al., 2010; Agrawal & Matsa, 2013); financial distress can induce the departure of valuable employees and the loss of firm-specific organizational knowledge (Eisfeldt & Papanikolaou, 2013; Donangelo, 2014; Klasa et al., 2018); and leverage can weaken the job security needed to sustain investment in firm-specific human capital (Jaggia & Thakor, 1994). These mechanisms imply that changes in firms' reliance on labor can affect their optimal financing choices.[13]

Generative AI provides a unique setting because it can substitute for cognitive and language-intensive tasks while reorganizing production around software, proprietary data, internal workflows, organizational routines, and specialized human capital. Recent evidence indicates that Generative AI affects labor demand along particular margins, including junior hiring and highly exposed occupations, although its effect on aggregate firm employment remains mixed (Eloundou et al., 2024; Brynjolfsson et al., 2025; Hampole et al., 2025; Hosseini Maasoum & Lichtinger, 2025). We therefore do not assume that adoption mechanically reduces total employment. Instead, we examine whether adoption is associated with mean-

[13] Serfling (2016) shows that firing costs raise operating leverage and crowd out financial leverage, providing direct evidence that labor adjustment costs can shape capital structure.

ingful workforce restructuring and whether the financing response is stronger where firms rely more heavily on the types of labor that Generative AI can affect.

If Generative AI substitutes for some cognitive tasks while reorganizing productive knowledge into firm-specific workflows, data systems, and organizational routines, the financing consequences of labor substitution need not resemble those of conventional automation. In particular, deleveraging should be more pronounced where adoption represents a more consequential reorganization of labor and productive knowledge. We test this implication using WARN layoffs, ex-ante Generative AI exposure, and several measures of firms' reliance on labor and knowledge workers.

A. WARN layoff notices

We first verify that Generative AI adoption leads to workforce displacement. WARN Act layoff notices provide a legally mandated record of mass layoffs and plant closings at large employers. Column (1) of Panel A in Table 5 shows that adopting firms are 1.5 percentage points more likely to file a layoff notice, and Column (4) estimates a 6.2% increase in the number of affected workers. Both coefficients are statistically significant at the 5% level. The results remain with firm-level controls in Columns (2) and (5) and with industry-by-year-quarter fixed effects in Columns (3) and (6), where the layoff probability coefficient rises to 1.7 percentage points, an 81% increase relative to the sample mean of 2.1%.

Panel B examines whether adoption triggers large-scale workforce reductions and whether it triggers facility shutdowns. Column (1) shows that adoption increases the probability of a mass layoff affecting 500 or more workers by 0.7 percentage points, statistically significant at the 5% level. The estimate is stable in Columns (2) and (3). Plant closings, which involve the shutdown of an employment site, show no response to adoption in Columns (4) through (6). Adopting firms are substituting Generative AI for labor within continuing operations, performing the same work with fewer employees rather than shutting down sites.

[Insert Table 5 about here]

B. Generative AI exposure

We further validate that these layoffs reflect Generative AI substituting for workers rather than other characteristics of adopting firms. We follow Eisfeldt et al. (2023) to construct firm-level exposure from O*NET task descriptions, capturing the share of each firm's workforce, measured as of 2022Q3 before ChatGPT's release made Generative AI widely available, in occupations whose core tasks large language models can perform.[14] They provide evidence that this exposure captures labor substitutability: substitution through firms' core tasks, rather than complementarity through supplemental tasks, accounts for the economic effects of Generative AI. These occupations, as Eloundou et al. (2024) document, are concentrated in cognitive and language-intensive work across professional services, distinct from the routine manual tasks displaced by earlier waves of automation (Autor et al., 2003). Panel A of Table 6 shows that the layoff response concentrates among high-exposure firms. Column (1) of Panel A shows that adoption increases the probability of layoff filing by 1.9 percentage points among high-exposure firms, statistically significant at the 5% level. Column (3) estimates a 7.7% increase in the number of affected workers, also significant at the 5% level. Low-exposure firms show no layoff response in Columns (4) through (6). Generative AI displaces knowledge workers, not production workers.

Generative AI displaces workers in high-exposure occupations. If this substitution also drives the leverage reduction, high-exposure firms should deleverage more. Panel B of Table 6 shows that they do. Column (1) of Panel B shows that high-exposure firms reduce leverage by 3.0 percentage points after Generative AI adoption. The effect holds at 2.2 percentage points with controls in Column (2), 6.6% of the sample mean, with a 16.6% increase in interest coverage in Column (4), both significant at the 1% level. Low-exposure firms show no leverage reduction once controls are included in Column (6). Firms with greater labor substitution potential deleverage more after adopting Generative AI.

[14] We verify the measure using the occupation scores of Eisfeldt et al. (2023), taken directly from https://artificialminushuman.com/ and aggregated with the same position weights and median split. Table IA10 in the Internet Appendix reports both panels; the results are qualitatively similar.

[Insert Table 6 about here]

C. Labor intensity

We next test whether firms that depend more on labor also deleverage more. We measure labor intensity from Revelio Labs employee positions rather than Compustat headcount. Revelio data are drawn from professional profiles, providing granular coverage of the white-collar workforce that Generative AI targets. Compustat reports a single annual headcount that aggregates production and professional workers. With Revelio, employees per PP&E measures reliance on knowledge workers relative to physical capital, sales per employee captures revenue per knowledge worker, and the net hiring rate following Belo et al. (2014) captures recent workforce expansion. All three measures are averaged over firm-quarters through 2022Q3, before ChatGPT's release made Generative AI commercially available.

Table 7 shows that the leverage reduction concentrates among labor-intensive firms. In Panel A, Column (1) shows that firms with high employees per PP&E reduce leverage by 3.0 percentage points. The effect remains economically and statistically significant at 2.1 percentage points with controls in Column (2), with a 22.9% increase in interest coverage in Column (4). Capital-intensive firms show no leverage reduction in Columns (5) and (6). Panel B reverses the ratio: Column (1) shows that firms where each worker generates less revenue reduce leverage by 3.5 percentage points; the effect holds at 2.6 percentage points with controls in Column (2). High-productivity firms show no leverage response in Columns (5) and (6). Panel C uses the net hiring rate. Column (1) shows that firms with high ex-ante hiring reduce leverage by 3.4 percentage points, holding at 2.8 percentage points with controls in Column (2). Low-hiring firms show no leverage effect in Columns (5) and (6). Across all three measures, the leverage reduction is strongest where firms depend most on the knowledge workers that Generative AI can displace.

[Insert Table 7 about here]

4.2 Asset Intangibility

A. Firm-specific Generative AI adoption

A central implication of our argument is that the financing consequences of Generative AI should depend on how closely the resulting productive capital is tied to the adopting firm. Capital structure theory predicts that firms relying on specialized or firm-specific assets support less debt because those assets have lower value outside the organization and impose greater losses in financial distress (Titman, 1984; Titman & Wessels, 1988; Jaggia & Thakor, 1994; Frank & Goyal, 2009). More generally, an investment is firm specific when its value depends on the organization in which it is deployed. Such productive capital has limited alternative use, lower liquidation value relative to its value in place, and weaker collateral value because creditors may be unable to separate it from the firm's internal systems, routines, and know-how.

Generative AI implementation can vary substantially along this dimension. Adoption is more firm specific when firms integrate Generative AI with proprietary data, internal documents, established procedures, and business-specific workflows—for example, through fine-tuning on proprietary records, retrieval over internal information, or agentic systems designed around the firm's own processes. By contrast, adoption that consists primarily of general-purpose tools or broadly transferable capabilities should remain more separable from the adopting firm. If the financing response depends on the nature of the productive capital created through adoption, the leverage reduction should be larger when Generative AI implementation is more closely tied to the firm's own business.

We measure firm specificity from how closely a firm's Generative AI hiring maps to its own business, fixed at adoption. Let $\mathcal{P}_i$ denote firm i's integrator postings in its adoption quarter, $\mathbf{v}_p$ the vector representation of posting p, and $\mathbf{u}_i$ that of the firm's 10-K Item 1 business description filed in 2020, before any adoption in our sample. A firm's score is the

largest posting-filing cosine similarity,

$$S_i = \max_{p \in \mathcal{P}_i} \frac{\mathbf{v}_p \cdot \mathbf{u}_i}{\|\mathbf{v}_p\| \, \|\mathbf{u}_i\|}, \tag{3}$$

which captures whether at least one hire is closely tied to the business.[15] We represent the two documents in two ways, as dense vectors from a pretrained sentence-transformer model and as TF-IDF term vectors over the pooled corpus. Everything after the representation step is identical.

Table 8 splits adopters at the median score, comparing each group with the same never-adopters. Firm-specific adopters reduce leverage by 2.2 percentage points with a 30.7% increase in interest coverage in Columns (2) and (4) of Panel A, while general adopters show no effect on either outcome. Panel B gives the same firm-specific estimates, 2.2 percentage points and 25.9%, though general adopters there also record smaller effects on both outcomes. The leverage response is larger where adoption is more firm-specific, which is where the resulting capital is least pledgeable.

[Insert Table 8 about here]

B. Firm-specific Intangibles

Prior research shows that intangible capital generally supports less debt because it is more difficult to verify, pledge, and liquidate than tangible assets (Rampini & Viswanathan, 2013; Falato et al., 2022). This constraint can matter even for borrowing that is formally based on cash flows rather than specific asset collateral, because creditors ultimately care about the value that can be preserved and recovered in distress (Lian & Ma, 2021). At the same time, not all intangible capital has the same financing consequences. Codified and transferable intangible assets, such as patents, can sometimes be pledged to creditors (Mann, 2018), whereas organizational capital embedded in proprietary workflows, internal systems,

[15] Because S_i rises mechanically with the number of postings and the length of the filing, we take the residual $\hat{\varepsilon}_i$ from $S_i = \alpha + \beta_1 \log|\mathcal{P}_i| + \beta_2 \log W_i + \gamma_{j(i)} + \varepsilon_i$, where W_i is the filing's word count and $\gamma_{j(i)}$ an industry fixed effect, and standardize it across adopters before splitting at the median.

data, and firm-specific know-how is more difficult to separate from the organization and redeploy elsewhere.

This distinction provides a broader test of our mechanism. If Generative AI reorganizes productive knowledge in ways that are more difficult to transfer and redeploy, its financing consequences should be stronger in firms whose production already relies more heavily on intangible and organizational capital. In such firms, Generative AI implementation is more likely to interact with existing proprietary knowledge, internal processes, and organizational routines, deepening the importance of assets whose value depends on remaining within the firm. The leverage reduction should hence be more pronounced among firms with greater ex-ante intangible and organizational-capital intensity.

We test this prediction using three measures of pre-existing intangible and organizational-capital intensity, all measured prior to ChatGPT's release. Panel A of Table 9 splits firms by R&D intensity. Among high-R&D firms, Generative AI adoption reduces leverage by 3.7 percentage points without controls and 2.8 percentage points with controls, while increasing interest coverage by 24.5%; all three estimates are significant at the 1% level. In contrast, low-R&D firms exhibit no significant financing response. Panel B yields a similar pattern using SG&A intensity, which captures investment in organizational capital, brand equity, and employee training. High-SG&A firms reduce leverage by 3.0 percentage points without controls and 2.4 percentage points with controls, and increase interest coverage by 34.0%, whereas low-SG&A firms show no significant response. Panel C provides complementary evidence based on patenting activity. Patenting firms reduce leverage by 2.6 percentage points without controls and 1.7 percentage points with controls, and increase interest coverage by 19.1%, while non-patenting firms exhibit no significant leverage reduction. Taken together, these results show that the financing response to Generative AI adoption is concentrated among firms whose productive activities already rely more heavily on intangible and organizational capital, consistent with the mechanism being strongest where productive knowledge is more difficult to separate from the firm and redeploy elsewhere.

[Insert Table 9 about here]

4.3 Generative AI versus Workplace Automation

The central implication of our framework is that labor-substituting technologies can have opposite effects on debt capacity depending on the nature of the productive capital that replaces labor. Conventional workplace automation provides a useful benchmark because it typically substitutes machinery and other tangible, redeployable assets for workers, whereas Generative AI can reorganize productive knowledge around software, proprietary data, internal workflows, and organizational routines that are more difficult to separate from the firm. We therefore compare the financing effects of Generative AI adoption with those of conventional workplace automation.

To measure exposure to conventional workplace automation, we follow Bates et al. (2025) and measure automation exposure using the substitutability of labor with automated capital (SLAC). SLAC aggregates the occupation-level probabilities of computerization developed by Frey and Osborne (2017) using the firm's workforce composition. A higher SLAC indicates that a larger share of the firm's workforce is employed in occupations that are more susceptible to substitution by automated capital.

We construct SLAC at the firm-quarter level as

$$\text{SLAC}_{i,t} = \frac{\sum_s p_s N_{i,s,t}}{\sum_s N_{i,s,t}} \tag{4}$$

where $N_{i,s,t}$ is the number of active Revelio Labs positions at firm i in occupation s during quarter t, mapped to O*NET-SOC codes, and p_s is the Frey and Osborne (2017) probability that occupation s is susceptible to computerization. Because the occupation-level probabilities are fixed over time, within-firm variation in SLAC arises from changes in the firm's occupational composition.

Table 10 reports the results. Panel A shows that greater exposure to conventional work-

place automation is associated with higher book leverage and interest coverage and lower cash holdings. These patterns are consistent with the liquidity results of Bates et al. (2025) and the leverage results of Cheng et al. (2025). The simultaneous increase in leverage and interest coverage is also consistent with automation expanding debt capacity without weakening firms' ability to service debt. In particular, Cheng et al. (2025) show that labor-replacing automation lowers firms' cost of debt, which can help explain why greater borrowing need not be accompanied by lower interest coverage.

The contrast with Generative AI is economically informative. Conventional automation typically replaces labor with machinery and other tangible, collateralizable assets, thereby increasing the share of productive capital that can support borrowing. Generative AI, by contrast, can substitute for cognitive labor while reorganizing productive knowledge around software, proprietary data, internal workflows, and organizational routines that are more difficult to separate from the firm and redeploy elsewhere. The two technologies can thus substitute for labor while having opposite implications for debt capacity.

Panel B includes Generative AI adoption and SLAC jointly. The Generative AI adoption coefficient remains essentially unchanged, while the two forms of labor-substituting technology continue to exhibit opposite financing effects. The comparison hence reinforces the paper's central interpretation: labor substitution itself does not determine the direction of the financing response. What matters is the nature of the productive capital that replaces labor.

[Insert Table 10 about here]

5 Conclusion

In this study, we examine how the financing consequences of labor-substituting technology depend on the nature of the productive capital that replaces labor. Generative AI provides a distinctive setting because it can substitute for cognitive labor while reorganizing productive

knowledge around firm-specific workflows, proprietary data, internal systems, and organizational know-how. Unlike conventional automation, which often replaces labor with tangible and redeployable capital, Generative AI may shift production toward capital whose value is more deeply embedded in the adopting firm and thus less pledgeable and recoverable in financial distress.

Using more than 350 million LinkedIn job postings from Revelio Labs, we construct a firm-level measure of Generative AI adoption based on hiring for "integrator" roles that build, fine-tune, deploy, or operate large-language-model systems. In a staggered DiD design, adopting firms reduce leverage and improve interest coverage. These effects are economically meaningful, increase with adoption intensity, and are robust across alternative leverage measures, modern staggered-adoption estimators, matching and synthetic DiD approaches, placebo tests, and two instrumental-variable strategies.

The cross-sectional evidence helps explain why Generative AI adoption reduces rather than increases debt capacity. Adoption is associated with meaningful workforce restructuring, and the financing response is strongest where firms rely more heavily on the types of labor that Generative AI can affect. More directly, deleveraging is substantially stronger when Generative AI implementation is closely tied to the firm's own business and is also concentrated among firms with greater ex-ante intangible and organizational-capital intensity. Together, these findings are consistent with the view that adoption has larger financing consequences when productive knowledge becomes more difficult to separate from the firm and redeploy elsewhere.

The contrast with conventional workplace automation sharpens this interpretation. Traditional automation often replaces labor with machinery or other tangible capital that can support borrowing. In our setting, conventional automation is associated with higher leverage, whereas Generative AI adoption is associated with lower leverage. Thus, labor substitution itself does not determine the direction of the financing response; what matters is the nature of the productive capital that replaces the labor.

Our findings extend the capital-structure literature by showing that the financing consequences of intangible capital depend not only on its quantity, but also on its organization and redeployability. Intangible capital that is codified and transferable may retain substantial value outside the firm, whereas organizational capital tied to proprietary workflows, internal systems, and firm-specific know-how is more difficult to separate from the organization and may support less debt. Generative AI adoption provides within-firm evidence on this distinction and shows how technological change can reshape financing policy by altering both firms' reliance on labor and the organization of productive knowledge.

More broadly, our findings suggest that the effect of technological change on capital structure depends on the form that productive capital takes after adoption. Technologies that shift production toward tangible, transferable, and collateralizable assets may expand debt capacity, whereas technologies that reorganize productive knowledge around firm-specific and organizationally embedded capital may reduce it. Understanding the financing consequences of technological change hence requires attention not only to whether technology substitutes for labor, but also to the nature of the productive capital that replaces it.

References

Agrawal, A. K., & Matsa, D. A. (2013). Labor unemployment risk and corporate financing decisions. *Journal of Financial Economics*, *108*(2), 449–470.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., & Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, *111*(12), 4088–4118.

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, *118*(4), 1279–1333.

Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, *151*, 103745.

Bae, K.-H., Kang, J.-K., & Wang, J. (2011). Employee treatment and firm leverage: A test of the stakeholder theory of capital structure. *Journal of Financial Economics*, *100*(1), 130–153.

Bartik, T. J. (1991). Who benefits from state and local economic development policies?

Bates, T. W., Du, F., & Wang, J. J. (2025). Workplace automation and corporate liquidity policy. *Management Science*, *71*(2), 1287–1314.

Belo, F., Lin, X., & Bazdresch, S. (2014). Labor hiring, investment, and stock return predictability in the cross section. *Journal of Political Economy*, *122*(1), 129–177.

Berk, J. B., Stanton, R., & Zechner, J. (2010). Human capital, bankruptcy, and capital structure. *The Journal of Finance*, *65*(3), 891–926.

Borusyak, K., Jaravel, X., & Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *The Review of Economic Studies*, *91*(6), 3253–3285.

Bronars, S. G., & Deere, D. R. (1991). The threat of unionization, the use of debt, and the preservation of shareholder wealth. *The Quarterly Journal of Economics*, *106*(1), 231–254.

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, *140*(2), 889–942.

Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230.

Campello, M., Gao, J., Qiu, J., & Zhang, Y. (2018). Bankruptcy and the cost of organized labor: Evidence from union elections. *The Review of Financial Studies*, *31*(3), 980–1013.

Cheng, X., Lyandres, E., Zhou, K., & Zhou, T. (2025). Labor-replacing automation and finance. *Management Science*, *71*(8), 6997–7028.

De Chaisemartin, C., & d'Haultfoeuille, X. (2024). Difference-in-differences estimators of intertemporal treatment effects. *Review of Economics and Statistics*, 1–45.

Demirci, O., Hannane, J., & Zhu, X. (2025). Who is ai replacing? the impact of generative ai on online freelancing platforms. *Management Science*, *71*(10), 8097–8108.

Donangelo, A. (2014). Labor mobility: Implications for asset pricing. *The Journal of Finance*, *69*(3), 1321–1346.

Eisfeldt, A. L., & Papanikolaou, D. (2013). Organization capital and the cross-section of expected returns. *The Journal of Finance*, *68*(4), 1365–1406.

Eisfeldt, A. L., Schubert, G., & Zhang, M. B. (2023). *Generative ai and firm values* (Working Paper). National Bureau of Economic Research.

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). Gpts are gpts: Labor market impact potential of llms. *Science*, *384*(6702), 1306–1308.

Falato, A., Kadyrzhanova, D., Sim, J., & Steri, R. (2022). Rising intangible capital, shrinking debt capacity, and the us corporate savings glut. *The Journal of Finance*, *77*(5), 2799–2852.

Faulkender, M., & Petersen, M. A. (2006). Does the source of capital affect capital structure? *The Review of Financial Studies*, *19*(1), 45–79.

Favilukis, J., Lin, X., & Zhao, X. (2020). The elephant in the room: The impact of labor obligations on credit markets. *American Economic Review*, *110*(6), 1673–1712.

Frank, M. Z., & Goyal, V. K. (2009). Capital structure decisions: Which factors are reliably important? *Financial Management*, *38*(1), 1–37.

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological forecasting and social change*, *114*, 254–280.

Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, *110*(8), 2586–2624.

Hampole, M., Papanikolaou, D., Schmidt, L. D., & Seegmiller, B. (2025). *Artificial intelligence and the labor market*. National Bureau of Economic Research.

Hershbein, B., & Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, *108*(7), 1737–1772.

Hirsch, B. T., Macpherson, D. A., & Even, W. E. (2026). *Union membership, coverage, and earnings from the CPS*. https://unionstats.com

Hosseini Maasoum, S. M., & Lichtinger, G. (2025). Generative ai as seniority-biased technological change: Evidence from us resume and job posting data. *Available at SSRN*.

Humlum, A., & Vestergaard, E. (2026). *Still waters, rapid currents: Early labor market transformation under generative ai* (NBER Working Paper No. 33777). National Bureau of Economic Research.

Jaggia, P. B., & Thakor, A. V. (1994). Firm-specific human capital and optimal capital structure. *International Economic Review*, 283–308.

Jiang, W., Park, J., Xiao, J., & Zhang, S. (2025). *Ai and the extended workday: Productivity, contracting efficiency, and distribution of rents*. National Bureau of Economic Research.

Kim, H. (2020). How does labor market size affect firm capital structure? evidence from large plant openings. *Journal of Financial Economics*, *138*(1), 277–294.

Klasa, S., Ortiz-Molina, H., Serfling, M., & Srinivasan, S. (2018). Protection of trade secrets and capital structure decisions. *Journal of Financial Economics*, *128*(2), 266–286.

Kogan, L., Papanikolaou, D., Seru, A., & Stoffman, N. (2017). Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics*, *132*(2), 665–712.

Kwan, A., Liu, Y., & Matthies, B. (2026). Institutional investor attention. *The Journal of Finance*.

Li, Y. (2025). Fragile new economy: Intangible capital, corporate savings glut, and financial instability. *American Economic Review*, *115*(4), 1100–1141.

Lian, C., & Ma, Y. (2021). Anatomy of corporate borrowing constraints. *The Quarterly Journal of Economics*, *136*(1), 229–291.

Mann, W. (2018). Creditor rights and innovation: Evidence from patent collateral. *Journal of Financial Economics*, *130*(1), 25–47.

Matsa, D. A. (2010). Capital structure as a strategic variable: Evidence from collective bargaining. *The Journal of Finance*, *65*(3), 1197–1232.

Matsa, D. A. (2018). Capital structure and a firm's workforce. *Annual Review of Financial Economics*, *10*(1), 387–412.

Myers, S. C. (1977). Determinants of corporate borrowing. *Journal of Financial Economics*, *5*(2), 147–175.

Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, *381*(6654), 187–192.

Rajan, R. G., & Zingales, L. (1995). What do we know about capital structure? some evidence from international data. *The Journal of Finance*, *50*(5), 1421–1460.

Rambachan, A., & Roth, J. (2023). A more credible approach to parallel trends. *The Review of Economic Studies*, *90*(5), 2555–2591.

Rampini, A. A., & Viswanathan, S. (2013). Collateral and capital structure. *Journal of Financial Economics*, *109*(2), 466–492.

Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, *4*(3), 305–322.

Serfling, M. (2016). Firing costs and capital structure decisions. *The Journal of Finance*, *71*(5), 2239–2286.

Simintzi, E., Vig, V., & Volpin, P. (2015). Labor protection and leverage. *The Review of Financial Studies*, *28*(2), 561–591.

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, *225*(2), 175–199.

Titman, S. (1984). The effect of capital structure on a firm's liquidation decision. *Journal of Financial Economics*, *13*(1), 137–151.

Titman, S., & Wessels, R. (1988). The determinants of capital structure choice. *The Journal of Finance*, *43*(1), 1–19.

Trajtenberg, M. (2018). *Ai as the next gpt: A political-economy perspective* (NBER Working Paper No. 24245). National Bureau of Economic Research.

Figures

Figure 1: **Example Job Posting: Generative AI**

Figure 1 shows part of a LinkedIn job posting. In the Revelio Labs data, each posting record contains the employer name and identifier, job title, location (country, state, metro area), posting and removal dates, salary range (reported or predicted), O*NET occupation code, job category, remote work type, and the full text description.

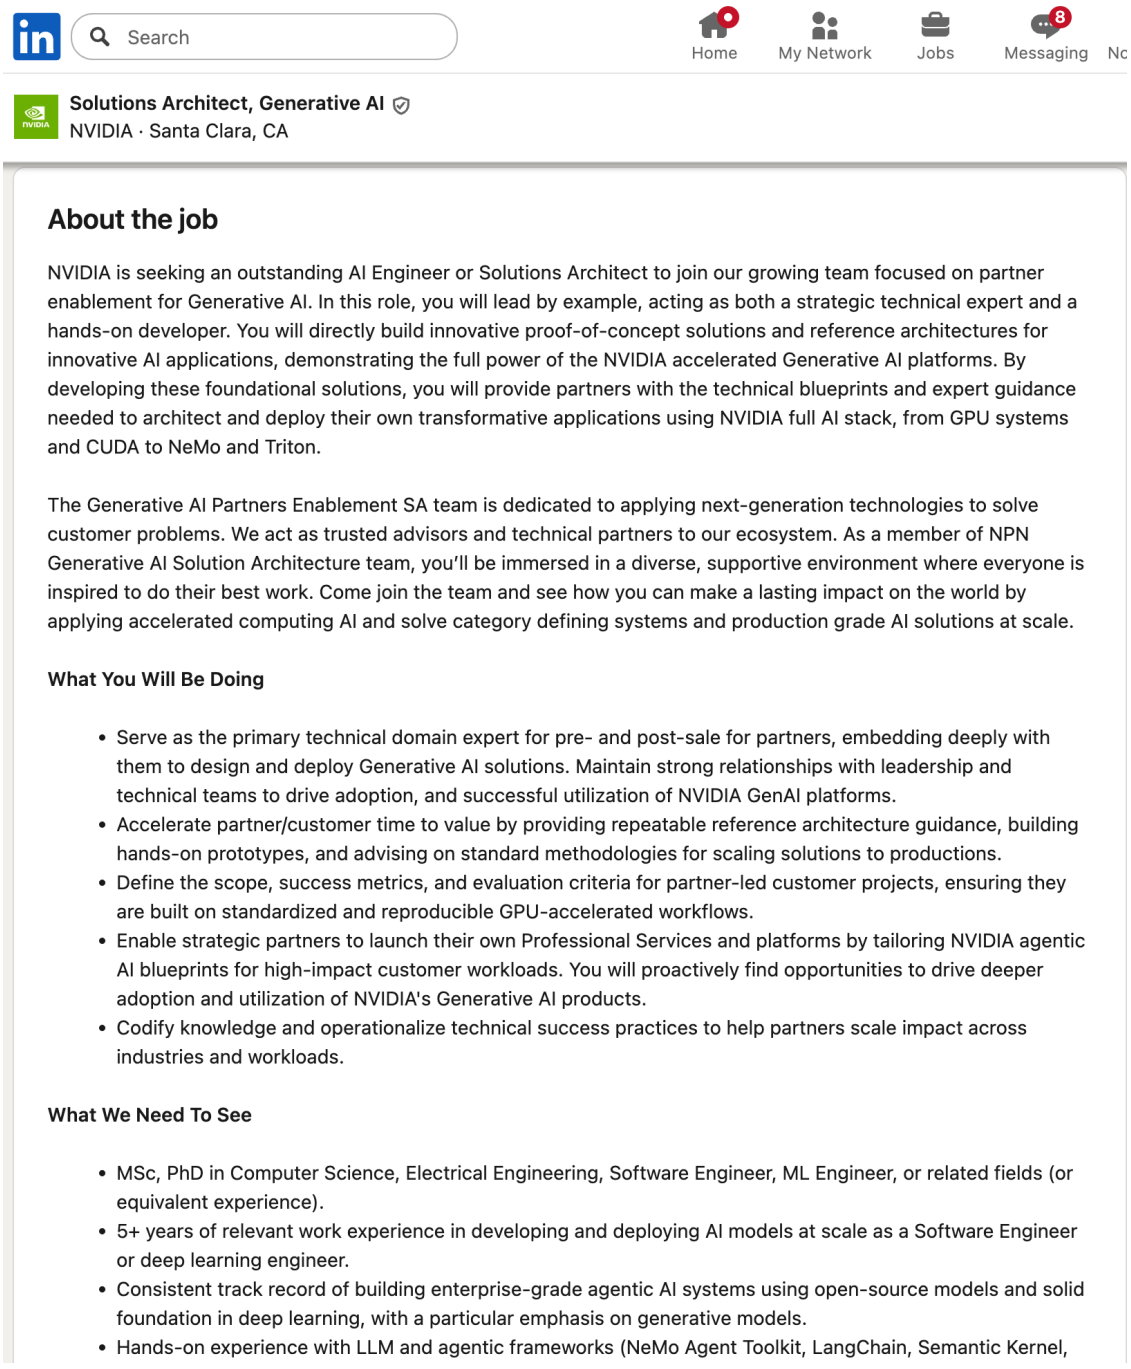


Figure 2: **Human Capital Investment in Generative AI and Job Posting Trends**

Figure 2 plots the aggregate quarterly annualized salary of Generative AI integrator postings (bars, left axis) and the share of Generative AI postings among all technology postings (line, right axis) for matched Compustat firms from 2020Q1 through 2025Q2. The salary aggregate sums posting-level annualized compensation, using Revelio-predicted salaries where reported compensation is unavailable. Generative AI integrator postings cover positions that build, develop, or deploy Generative AI systems; postings that mention Generative AI only as a general skill are classified as user postings and excluded. The vertical dashed line marks the first public release of ChatGPT (GPT-3.5) in November 2022.

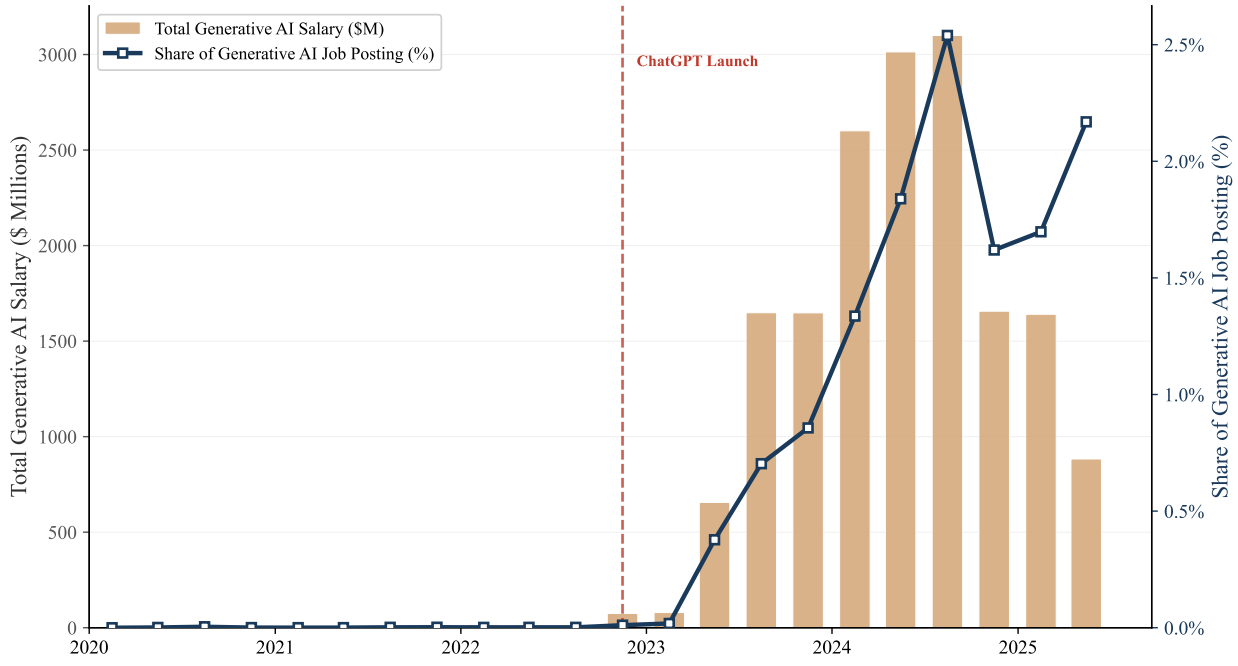
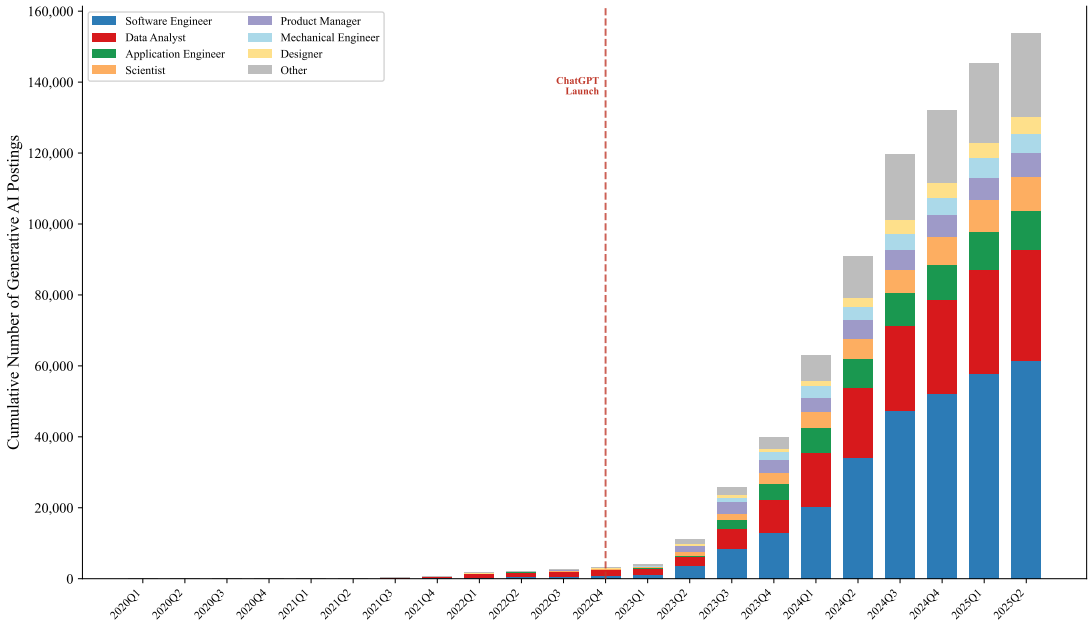
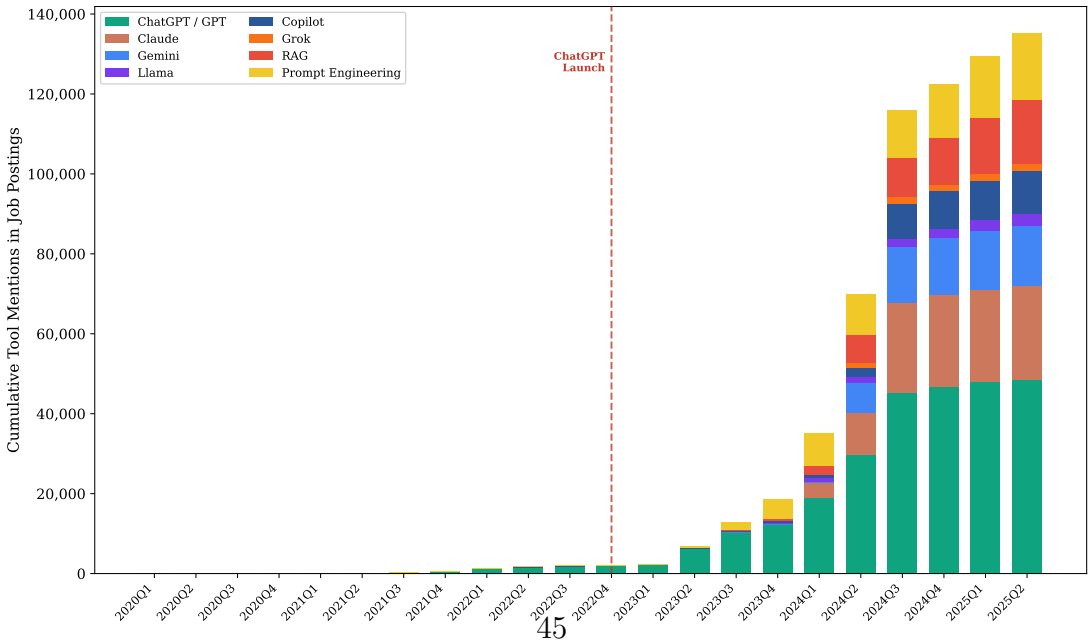


Figure 3: **Generative AI Human Capital Stock by Role and Tool**

Figure 3 plots the cumulative number of Generative AI job postings by quarter from 2020Q1 through 2025Q2 for Compustat firms. We plot cumulative postings because integrator hires accumulate as a stock of Generative AI human capital: a worker hired in one quarter continues to contribute to the firm's AI capacity in subsequent quarters. Panel (a) decomposes postings by the top seven job roles based on the Revelio Labs K-50 taxonomy; postings outside the top seven are grouped as Other. Panel (b) decomposes postings by the Generative AI tools and technologies mentioned in job descriptions or matched keywords. The eight tool categories are: ChatGPT/GPT (including Azure OpenAI and OpenAI API), Claude, Gemini, Llama, Copilot, Grok, Retrieval-Augmented Generation (RAG), and Prompt Engineering. A single posting may mention multiple tools. Panel (b) counts only postings that reference a specific tool or technique; postings that mention Generative AI in general terms without naming a particular tool are not included in the tool counts. The vertical dashed line marks the first public release of ChatGPT (GPT-3.5) in November 2022.



(a) Job Role Composition



(b) Generative AI Tool Adoption

Figure 4: **Event Study: Generative AI Adoption and Leverage**

Figure 4 presents coefficient estimates and 95% confidence intervals for the effect of Generative AI adoption on book leverage. The event is defined as the quarter when a firm first posts a Generative AI integrator job. The figure reports estimates from the baseline difference-in-differences estimator alongside four others that are robust to heterogeneous treatment effects and staggered adoption: Callaway and Sant'Anna (2021), Borusyak et al. (2024), Sun and Abraham (2021), and De Chaisemartin and d'Haultfoeuille (2024). Book leverage is total debt divided by total assets, measured one quarter ahead. All specifications include firm and year-quarter fixed effects. Standard errors are clustered at the firm level.

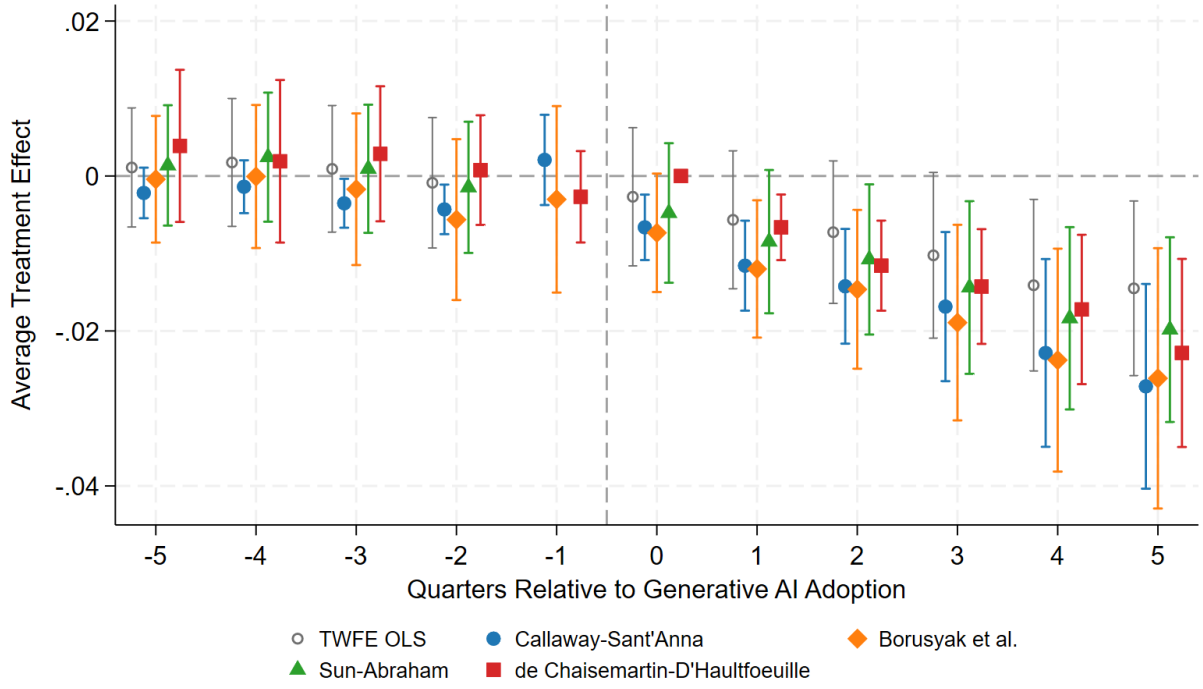
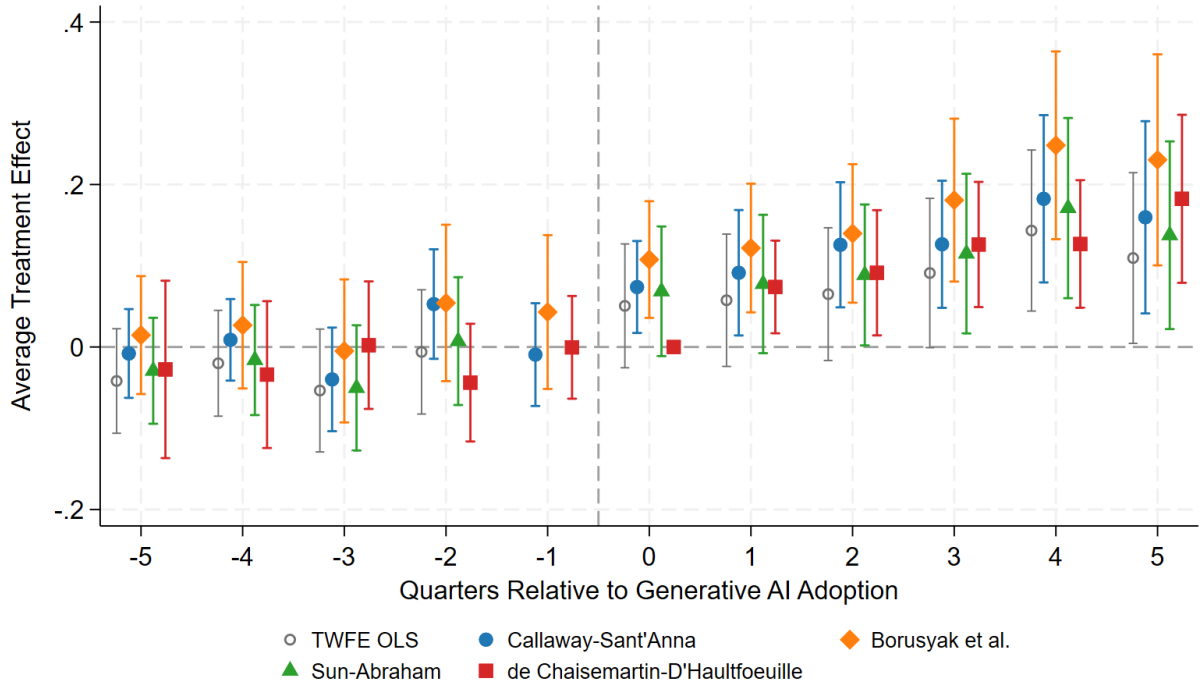


Figure 5: **Event Study: Generative AI Adoption and Interest Coverage**

Figure 5 presents coefficient estimates and 95% confidence intervals for the effect of Generative AI adoption on interest coverage. The event is defined as the quarter when a firm first posts a Generative AI integrator job. The figure reports estimates from the baseline difference-in-differences estimator alongside four others that are robust to heterogeneous treatment effects and staggered adoption: Callaway and Sant'Anna (2021), Borusyak et al. (2024), Sun and Abraham (2021), and De Chaisemartin and d'Haultfoeuille (2024). Log interest coverage is the natural logarithm of one plus EBITDA divided by interest expense, measured one quarter ahead. All specifications include firm and year-quarter fixed effects. Standard errors are clustered at the firm level.



Tables

Table 1: **Summary Statistics**

Table 1 reports summary statistics for the main analysis sample. The sample covers firm-quarters from 2020Q1 to 2025Q2, excluding financial firms (SIC 60–67) and utilities (SIC 49). All continuous variables are winsorized at the 1st and 99th percentiles. Variable definitions are in Table A1.

	N	Mean	S.D.	P25	Median	P75
Generative AI Adoption	41,007	0.081	0.273	0.000	0.000	0.000
Ln(Generative AI Postings)	41,007	0.189	0.718	0.000	0.000	0.000
Ln(Generative AI Salary)	41,007	1.100	3.733	0.000	0.000	0.000
Book Leverage	41,007	0.332	0.251	0.142	0.299	0.456
Log Interest Coverage	41,007	1.594	1.541	0.000	1.570	2.585
Market Leverage	41,007	0.262	0.230	0.075	0.200	0.388
Long-term Leverage	41,007	0.286	0.234	0.096	0.253	0.404
Net Debt/Assets	41,007	0.108	0.386	−0.099	0.162	0.351
Sales Growth	40,528	0.244	1.004	−0.061	0.063	0.228
Cash Ratio	41,007	0.225	0.242	0.046	0.122	0.330
Capex/Assets	40,939	0.008	0.010	0.002	0.005	0.010
Log(Assets)	41,007	7.004	2.091	5.607	7.068	8.424
Tangibility	41,007	0.228	0.213	0.073	0.150	0.320
Market-to-Book	41,007	3.812	9.131	1.072	2.203	4.661
Profitability	41,007	−0.003	0.071	−0.018	0.019	0.035
R&D	41,007	0.018	0.032	0.000	0.001	0.023
Employees/PP&E	40,981	25.861	53.198	3.363	8.520	23.402
Sales/Employees	40,973	1.037	1.867	0.247	0.514	1.010
Net Hiring Rate	37,879	0.006	0.036	−0.008	0.002	0.019
Union Coverage Rate	41,007	0.060	0.068	0.016	0.035	0.082
Has WARN	41,007	0.021	0.143	0.000	0.000	0.000
Mass Layoff (500+)	41,007	0.003	0.052	0.000	0.000	0.000
Plant Closing	41,007	0.007	0.084	0.000	0.000	0.000
Log(1+Workers)	41,007	0.081	0.562	0.000	0.000	0.000
Generative AI Exposure (Firm)	41,007	0.444	0.084	0.393	0.440	0.504
SGA/Assets	36,808	0.296	0.341	0.104	0.210	0.390
Union Membership Rate	41,007	0.053	0.064	0.013	0.028	0.072

Table 2: **Generative AI Adoption and Leverage**

Table 2 reports difference-in-differences estimates of the effect of Generative AI adoption on firm leverage. The dependent variable is the one-quarter-ahead value of the measure indicated in the column header. Leverage is total debt divided by total assets. Log(Interest Coverage) is the natural logarithm of one plus EBITDA divided by interest expense. Generative AI Adoption is an indicator equal to one for firm-quarters at and after the first Generative AI integrator job posting. Columns (1) and (4) include firm and year-quarter fixed effects without controls. Columns (2) and (5) add firm-level controls. Columns (3) and (6) include industry-by-year-quarter fixed effects, to absorb time-varying shocks common to industry. Controls include Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Book Leverage			*Log(Interest Coverage)*		
	(1)	(2)	(3)	(4)	(5)	(6)
Generative AI Adoption	-0.017***	-0.013**	-0.016***	0.178***	0.157***	0.162***
	(0.005)	(0.005)	(0.005)	(0.044)	(0.041)	(0.040)
Log(Assets)		-0.017**	-0.014*		-0.218***	-0.206***
		(0.008)	(0.008)		(0.032)	(0.033)
Tangibility		0.380***	0.364***		-0.419**	-0.168
		(0.036)	(0.037)		(0.184)	(0.182)
Market-to-Book		-0.001***	-0.001***		0.001	0.001
		(0.000)	(0.000)		(0.001)	(0.001)
Profitability		-0.257***	-0.221***		6.009***	5.151***
		(0.058)	(0.060)		(0.361)	(0.361)
R&D		0.263	0.269		6.054***	4.878***
		(0.174)	(0.181)		(0.589)	(0.565)
Observations	41,007	41,007	41,007	41,007	41,007	41,007
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	No	Yes	Yes	No
Ind × YearQuarter FE	No	No	Yes	No	No	Yes
Adjusted R^2	0.847	0.855	0.857	0.721	0.734	0.744

Table 3: **Continuous Measures of Generative AI Adoption and Leverage**

Table 3 reports estimates using continuous measures of Generative AI adoption intensity. Panel A uses Ln(Generative AI Postings), the natural logarithm of one plus cumulative Generative AI job postings since the firm's first Generative AI posting. Panel B uses Ln(Generative AI Salary), the natural logarithm of one plus cumulative Generative AI salary investment since the firm's first Generative AI posting. The dependent variable is the one-quarter-ahead value of the measure indicated in the column header. Columns (1) and (4) include firm and year-quarter fixed effects without controls. Columns (2) and (5) add firm-level controls. Columns (3) and (6) include industry-by-year-quarter fixed effects, to absorb time-varying shocks common to industry. Controls include Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Book Leverage			*Log(Interest Coverage)*		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Count-Based						
Ln(Generative AI Postings)	-0.008***	-0.006***	-0.008***	0.073***	0.066***	0.066***
	(0.002)	(0.002)	(0.002)	(0.020)	(0.019)	(0.019)
Observations	41,007	41,007	41,007	41,007	41,007	41,007
Adjusted R^2	0.847	0.855	0.857	0.721	0.734	0.744
Panel B: Salary-Weighted						
Ln(Generative AI Salary)	-0.001***	-0.001***	-0.001***	0.014***	0.012***	0.013***
	(0.000)	(0.000)	(0.000)	(0.003)	(0.003)	(0.003)
Observations	41,007	41,007	41,007	41,007	41,007	41,007
Adjusted R^2	0.847	0.855	0.857	0.721	0.734	0.744
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	No	Yes	Yes	No
Ind × YearQuarter FE	No	No	Yes	No	No	Yes
Firm-level Controls	No	Yes	Yes	No	Yes	Yes

Table 4: **Bartik IV: LLM Capability Alignment**

Table 4 reports instrumental variable estimates using Bartik shift-share instruments that exploit variation in LLM capabilities across occupations over time. The instrument is $Z_{i,t} = \sum_s \text{Share}(g_i, s, t-1) \times \text{Alignment}(s, t)$, where $\text{Share}(g_i, s, t-1)$ is the employment share of occupation s in firm i's three-digit NAICS industry from the prior year's BLS Occupational Employment and Wage Statistics survey, and $\text{Alignment}(s, t)$ is the TF-IDF cosine similarity between O*NET occupation text and cumulative LLM capability text from the Epoch AI database released up to quarter t. Using lagged shares ensures the occupation composition is predetermined relative to the outcome. Alignment is subject to a monotonicity constraint (values can only increase over time) and the instrument is standardized to zero mean and unit variance. Each panel varies the text used for computing alignment. Panel A aligns O*NET Detailed Work Activities with Epoch AI task-label descriptions of large-scale LLMs. Panel B aligns O*NET Detailed Work Activities with Epoch AI task-and-domain descriptions of large-scale LLMs. Columns (1) and (2) report reduced-form estimates. Column (3) reports the first stage, where the dependent variable is Generative AI Adoption. Columns (4) and (5) report second-stage estimates, where the dependent variable is the one-quarter-ahead value of the indicated measure. All specifications include firm and year-quarter fixed effects and firm-level controls: Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Reduced Form		First Stage	Second Stage	
	Book Leverage	*Log(Interest Coverage)*	*Generative AI Adoption*	*Book Leverage*	*Log(Interest Coverage)*
	(1)	(2)	(3)	(4)	(5)
*Panel A: O*NET Work Activities / LLM Task-Labels*					
LLM Capability Bartik	−0.006**	0.111***	0.047***		
	(0.003)	(0.022)	(0.008)		
$\widehat{\text{Generative AI}}$ Adoption				−0.135**	2.376***
				(0.066)	(0.611)
Kleibergen-Paap F			36.5		
Adjusted R^2	0.855	0.735	0.466	−0.005	−0.251
*Panel B: O*NET Work Activities / LLM (Task + Domain)*					
LLM Capability Bartik	−0.007***	0.112***	0.044***		
	(0.003)	(0.021)	(0.008)		
$\widehat{\text{Generative AI}}$ Adoption				−0.168**	2.559***
				(0.071)	(0.660)
Kleibergen-Paap F			33.1		
Adjusted R^2	0.856	0.735	0.466	−0.043	−0.302
N	40,832	40,832	40,832	40,832	40,832
Firm FE	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes

Table 5: **Generative AI Adoption and Layoffs**

Table 5 reports difference-in-differences estimates of the effect of Generative AI adoption on layoffs using data from the Worker Adjustment and Retraining Notification (WARN) Act. The WARN Act (1988) requires U.S. employers with 100 or more employees to provide 60 days' advance written notice to workers and state governments before plant closings or mass layoffs. WARN notices are aggregated from state government filings via layoffdata.com. Panel A reports estimates for any WARN notice filing (extensive margin) and the log number of affected workers (intensive margin). Panel B decomposes WARN filings into mass layoffs affecting 500 or more workers and plant closings, defined under WARN as permanent or temporary shutdown of a single employment site affecting 50 or more employees. WARN notices are matched to Compustat firms using company name matching with manual verification of major employers. Generative AI Adoption is an indicator equal to one for firm-quarters at and after the first Generative AI integrator job posting. Columns (1) and (4) include firm and year-quarter fixed effects without controls. Columns (2) and (5) add firm-level controls. Columns (3) and (6) replace year-quarter fixed effects with industry-by-year-quarter fixed effects. Controls include Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: WARN Notice Filing and Affected Workers						
	Has WARN			*Log(1+Workers)*		
Generative AI Adoption	0.015**	0.016**	0.017**	0.062**	0.063**	0.068**
	(0.007)	(0.007)	(0.007)	(0.026)	(0.026)	(0.028)
Observations	41,007	41,007	41,007	41,007	41,007	41,007
Adjusted R^2	0.138	0.138	0.149	0.138	0.139	0.150
Panel B: Mass Layoff and Plant Closing						
	Mass Layoff (500+)			*Plant Closing*		
Generative AI Adoption	0.007**	0.007**	0.007**	0.005	0.005	0.005
	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)
Observations	41,007	41,007	41,007	41,007	41,007	41,007
Adjusted R^2	0.060	0.061	0.092	0.100	0.100	0.117
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	No	Yes	Yes	No
Ind × YearQuarter FE	No	No	Yes	No	No	Yes
Firm-level Controls	No	Yes	Yes	No	Yes	Yes

Table 6: **Generative AI Exposure, Layoffs, and Capital Structure**

Table 6 splits firms by ex-ante Generative AI exposure, measured as of 2022Q3, the quarter before ChatGPT's public release. We score each O*NET task for exposure to large language models following Eisfeldt et al. (2023), average task scores to the six-digit SOC occupation level, and aggregate to the firm as a positions-weighted average over its Revelio Labs positions. Firms are split at the median of this cross-sectional index, taken across firms. Panel A reports layoff outcomes: Has WARN is an indicator for filing a WARN Act notice in the quarter, Mass Layoff (500+) is an indicator for layoffs affecting 500 or more workers, and Log(1+Workers) is the natural logarithm of one plus the number of affected workers. Panel B reports capital structure outcomes, measured one quarter ahead: Book Leverage is total debt divided by total assets, and Log(Interest Coverage) is the natural logarithm of one plus EBITDA divided by interest expense. Generative AI Adoption is an indicator equal to one for firm-quarters at and after the firm's first Generative AI integrator job posting. All specifications include firm and year-quarter fixed effects. Panel A includes firm-level controls throughout; in Panel B odd-numbered columns exclude controls and even-numbered columns include them. Controls are Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Layoffs

	High Exposure			*Low Exposure*		
	Has WARN	*Mass (500+)*	*Log(1+Workers)*	*Has WARN*	*Mass (500+)*	*Log(1+Workers)*
	(1)	(2)	(3)	(4)	(5)	(6)
Generative AI Adoption	0.019**	0.005***	0.077**	0.012	0.010	0.044
	(0.008)	(0.002)	(0.031)	(0.012)	(0.006)	(0.046)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-level Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	18,702	18,702	18,702	22,305	22,305	22,305
Adjusted R^2	0.136	0.061	0.135	0.141	0.065	0.143

Panel B: Capital Structure

	High Exposure				*Low Exposure*			
	Book Leverage		*Log(Interest Coverage)*		*Book Leverage*		*Log(Interest Coverage)*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Generative AI Adoption	−0.030***	−0.022***	0.173***	0.166***	−0.013**	−0.009	0.149***	0.109**
	(0.008)	(0.008)	(0.063)	(0.060)	(0.006)	(0.006)	(0.052)	(0.047)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-level Controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	18,702	18,702	18,702	18,702	22,305	22,305	22,305	22,305
Adjusted R^2	0.829	0.839	0.739	0.746	0.863	0.872	0.683	0.706

Table 7: **Labor Intensity, Generative AI Adoption, and Capital Structure**

Table 7 reports subsample TWFE difference-in-differences estimates of the effect of Generative AI adoption on leverage outcomes, split by labor intensity measures averaged over all firm-quarters prior to ChatGPT's release (2022Q4). Panel A splits at the median of employees per PP&E, measured as Revelio Labs distinct employees divided by net property, plant, and equipment. Panel B splits at the median of sales per employee, measured as trailing twelve-month sales divided by Revelio Labs distinct employees. Panel C splits at the median of the net hiring rate following Belo et al. (2014), computed as the symmetric growth rate of Revelio Labs quarterly headcount (distinct employees). The dependent variable is the one-quarter-ahead value of the measure indicated in the column header. Generative AI Adoption is an indicator equal to one for firm-quarters at and after the first Generative AI integrator job posting. All specifications include firm and year-quarter fixed effects. Odd-numbered columns exclude firm-level controls; even-numbered columns include controls: Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Book Leverage		*Log(Interest Coverage)*		*Book Leverage*		*Log(Interest Coverage)*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Labor Intensity								
	High Employees / PP&E				Low Employees / PP&E			
Generative AI Adoption	-0.030***	-0.021***	0.235***	0.229***	-0.005	-0.004	0.104**	0.067
	(0.008)	(0.007)	(0.067)	(0.064)	(0.008)	(0.008)	(0.051)	(0.044)
Observations	19,731	19,731	19,731	19,731	19,707	19,707	19,707	19,707
Adjusted R^2	0.822	0.830	0.709	0.717	0.872	0.881	0.699	0.724
Panel B: Labor Productivity								
	Low Sales / Employees				High Sales / Employees			
Generative AI Adoption	-0.035***	-0.026***	0.174***	0.165**	-0.001	-0.000	0.150***	0.124***
	(0.008)	(0.008)	(0.066)	(0.064)	(0.007)	(0.007)	(0.052)	(0.046)
Observations	19,709	19,709	19,709	19,709	19,743	19,743	19,743	19,743
Adjusted R^2	0.830	0.839	0.712	0.718	0.875	0.881	0.658	0.687
Panel C: Hiring Activity								
	High Net Hiring Rate				Low Net Hiring Rate			
Generative AI Adoption	-0.034***	-0.028***	0.212***	0.181***	0.003	0.006	0.160***	0.136***
	(0.008)	(0.007)	(0.069)	(0.065)	(0.008)	(0.008)	(0.049)	(0.045)
Observations	19,340	19,340	19,340	19,340	19,317	19,317	19,317	19,317
Adjusted R^2	0.834	0.842	0.733	0.746	0.858	0.867	0.689	0.704
Firm-level Controls	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 8: **Firm-Specific Generative AI Adoption and Capital Structure**

Table 8 reports subsample difference-in-differences estimates split by how firm-specific a firm's Generative AI adoption is. Each panel measures firm-specificity differently. Panel A computes the cosine similarity between each adopter's Generative AI integrator postings in its first adoption quarter and its pre-adoption 10-K Item 1 business description, representing both as dense sentence-embedding vectors from a pretrained sentence-transformer model (MPNet). A firm's score is the highest similarity among its postings, which captures whether at least one hire is closely tied to the firm's business. We then net out the log number of postings, the log length of the filing, and industry, because the maximum mechanically rises with both the number of postings and the length of the filing. Panel B replaces the sentence embeddings with TF-IDF term vectors and computes the similarity. Firm-specific adopters are those above the median score; the remaining adopters are the general group, and both are compared with the same never-adopters. The dependent variable is the one-quarter-ahead value of the measure indicated in the column header. Generative AI Adoption is an indicator equal to one for firm-quarters at and after the first Generative AI integrator job posting. All specifications include firm and industry-by-year-quarter fixed effects. Odd-numbered columns exclude firm-level controls; even-numbered columns include controls: Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Firm-Specific Adopters				General Adopters			
	Book Leverage		*Log(Interest Coverage)*		*Book Leverage*		*Log(Interest Coverage)*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: 10-K Similarity, Sentence-Transformer Embedding								
Generative AI Adoption	−0.027***	−0.022***	0.316***	0.307***	−0.015**	−0.011	0.054	0.054
	(0.008)	(0.008)	(0.070)	(0.068)	(0.007)	(0.007)	(0.059)	(0.055)
Observations	35,154	35,154	35,154	35,154	35,150	35,150	35,150	35,150
Adjusted R^2	0.844	0.852	0.734	0.743	0.846	0.853	0.737	0.746
Panel B: 10-K Similarity, TF-IDF Term Vectors								
Generative AI Adoption	−0.026***	−0.022***	0.262***	0.259***	−0.017**	−0.012*	0.121**	0.114**
	(0.008)	(0.008)	(0.076)	(0.073)	(0.007)	(0.007)	(0.053)	(0.050)
Observations	35,146	35,146	35,146	35,146	35,177	35,177	35,177	35,177
Adjusted R^2	0.845	0.853	0.732	0.741	0.845	0.853	0.739	0.748
Firm-level Controls	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ind × YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 9: **Intangible Investment Intensity, Generative AI Adoption, and Capital Structure**

Table 9 reports subsample TWFE difference-in-differences estimates of the effect of Generative AI adoption on leverage outcomes, split by intangible investment and innovation measures determined prior to ChatGPT's release (2022Q4). Panel A splits at the median of R&D intensity (research and development expense divided by total assets), averaged over all firm-quarters prior to 2022Q4. Panel B splits at the median of SGA intensity (selling, general, and administrative expense divided by total assets), averaged over all firm-quarters prior to 2022Q4. Panel C splits on whether the firm holds at least one USPTO utility patent filed in the five years preceding 2022Q3 and subsequently granted. The dependent variable is the one-quarter-ahead value of the measure indicated in the column header. Generative AI Adoption is an indicator equal to one for firm-quarters at and after the first Generative AI integrator job posting. All specifications include firm and year-quarter fixed effects. Odd-numbered columns exclude firm-level controls; even-numbered columns include controls: Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Book Leverage		*Log(Interest Coverage)*		*Book Leverage*		*Log(Interest Coverage)*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: R&D Intensity								
	High R&D				Low R&D			
Generative AI Adoption	−0.037***	−0.028***	0.251***	0.245***	0.001	0.004	0.078	0.030
	(0.009)	(0.008)	(0.065)	(0.062)	(0.006)	(0.006)	(0.054)	(0.050)
Observations	19,732	19,732	19,732	19,732	19,720	19,720	19,720	19,720
Adjusted R^2	0.797	0.808	0.732	0.740	0.893	0.899	0.655	0.682
Panel B: SGA Intensity								
	High SGA				Low SGA			
Generative AI Adoption	−0.030***	−0.024***	0.385***	0.340***	0.005	0.005	0.026	0.016
	(0.009)	(0.008)	(0.077)	(0.072)	(0.006)	(0.006)	(0.050)	(0.045)
Observations	17,942	17,942	17,942	17,942	17,920	17,920	17,920	17,920
Adjusted R^2	0.825	0.833	0.705	0.717	0.893	0.899	0.639	0.680
Panel C: Patent Activity								
	Patenting Firms				Non-Patenting Firms			
Generative AI Adoption	−0.026***	−0.017**	0.196***	0.191***	−0.014**	−0.010	0.125*	0.090
	(0.008)	(0.007)	(0.056)	(0.053)	(0.007)	(0.007)	(0.070)	(0.066)
Observations	17,795	17,795	17,795	17,795	23,212	23,212	23,212	23,212
Adjusted R^2	0.816	0.827	0.753	0.763	0.866	0.874	0.695	0.710
Firm-level Controls	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 10: **Automation Exposure and Capital Structure**

Table 10 reports estimates of the relation between workplace automation exposure and capital structure. Automation Exposure (SLAC) is the substitutability of labor with automated capital, constructed as the positions-weighted average Frey and Osborne (2017) occupational probability of computerization across a firm's workforce in a given quarter. We build the measure at the firm-quarter level from Revelio Labs position records mapped to O*NET occupation codes, and standardize it to mean zero and unit variance, so coefficients are per one standard deviation. Generative AI Adoption is an indicator equal to one for firm-quarters at and after the firm's first Generative AI integrator job posting. In both panels the dependent variable is the one-quarter-ahead value of the measure indicated in the column header. Book Leverage is total debt divided by total assets. Log(Interest Coverage) is the natural logarithm of one plus EBITDA divided by interest expense. Cash Holdings is cash and short-term investments divided by total assets. Panel A includes automation exposure alone, and for each outcome reports results without controls and with the firm-level controls Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Panel B includes automation exposure and Generative AI adoption in the same regression, following the column layout of Table 2: Columns (1) and (4) include firm and year-quarter fixed effects without controls, Columns (2) and (5) add the firm-level controls, and Columns (3) and (6) replace year-quarter with industry-by-year-quarter fixed effects. All specifications include firm fixed effects. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Automation Exposure

	Book Leverage		*Log(Interest Coverage)*		*Cash Holdings*	
	(1)	(2)	(3)	(4)	(5)	(6)
Automation Exposure (SLAC)	0.039***	0.038***	0.184***	0.169***	−0.042***	−0.037***
	(0.012)	(0.012)	(0.047)	(0.044)	(0.011)	(0.011)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-level Controls	No	Yes	No	Yes	No	Yes
Observations	41,007	41,007	41,007	41,007	41,007	41,007
Adjusted R^2	0.847	0.856	0.721	0.734	0.888	0.897

Panel B: Automation Exposure and Generative AI Adoption

	Book Leverage			*Log(Interest Coverage)*		
	(1)	(2)	(3)	(4)	(5)	(6)
Generative AI Adoption	−0.017***	−0.012**	−0.015***	0.181***	0.159***	0.165***
	(0.005)	(0.005)	(0.005)	(0.044)	(0.041)	(0.040)
Automation Exposure (SLAC)	0.038***	0.037***	0.027**	0.190***	0.174***	0.157***
	(0.012)	(0.012)	(0.012)	(0.047)	(0.044)	(0.045)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	No	Yes	Yes	No
Ind × YearQuarter FE	No	No	Yes	No	No	Yes
Firm-level Controls	No	Yes	Yes	No	Yes	Yes
Observations	41,007	41,007	41,007	41,007	41,007	41,007
Adjusted R^2	0.847	0.856	0.857	0.722	0.735	0.745

Appendix

Table A1: **Variable Definitions**

Variable	Definition	Source
Generative AI Adoption	Indicator equal to one for all firm-quarters at or after the quarter when the firm first posts a Generative AI integrator job vacancy.	Revelio Labs
Ln(Generative AI Postings)	Natural logarithm of one plus cumulative Generative AI integrator job postings since the firm's first integrator posting.	Revelio Labs
Ln(Generative AI Salary)	Natural logarithm of one plus cumulative annualized salary allocated to Generative AI integrator postings since adoption.	Revelio Labs
Book Leverage	Total debt (dlttq + dlcq) divided by total assets (atq).	Compustat
Log Interest Coverage	Natural logarithm of one plus EBITDA (oibdpq) divided by interest expense (xintq).	Compustat
Market Leverage	Total debt divided by total debt plus market value of equity (prccq × cshoq).	Compustat
Long-term Leverage	Long-term debt (dlttq) divided by total assets.	Compustat
Net Debt/Assets	Total debt minus cash and short-term investments (cheq) divided by total assets.	Compustat
Has WARN	Indicator equal to one if the firm files at least one WARN Act layoff notice in the quarter.	WARN Act notices
Mass Layoff (500+)	Indicator equal to one if the firm files a WARN notice affecting 500 or more workers.	WARN Act notices
Plant Closing	Indicator equal to one if the firm files a WARN notice classified as a plant closing.	WARN Act notices
Log(1+Workers)	Natural logarithm of one plus the total number of workers named in WARN filings in the quarter.	WARN Act notices
Sales Growth	Quarter-over-quarter percentage change in net sales (saleq).	Compustat
Cash Ratio	Cash and short-term investments (cheq) divided by total assets.	Compustat
Capex/Assets	Capital expenditure (capxy, quarterly change) divided by total assets.	Compustat
Log(Assets)	Natural logarithm of total assets (atq).	Compustat
Tangibility	Net property, plant, and equipment (ppentq) divided by total assets.	Compustat
Market-to-Book	Market value of assets (atq − ceqq + prccq × cshoq) divided by book value of assets (atq).	Compustat
Profitability	Operating income before depreciation (oibdpq) divided by total assets.	Compustat
R&D	Research and development expense (xrdq) divided by total assets. Set to zero when missing.	Compustat
SGA Intensity	Average selling, general, and administrative expense (xsgaq) divided by total assets over all firm-quarters before 2022Q4.	Compustat
Patent Activity	Indicator equal to one if the firm holds at least one USPTO utility patent filed in the five years before 2022Q3.	Kogan et al. (2017)
GenAI Exposure (Firm)	Position-weighted average of six-digit SOC occupation-level Generative AI exposure scores within each firm's workforce, measured as of 2022Q3. Occupation scores are simple averages of O*NET task-level exposure labels classified by GPT-3.5.	Revelio Labs, O*NET
Employees/PP&E	Revelio Labs distinct employees divided by net PP&E (ppentq), averaged over firm-quarters before 2022Q4.	Revelio Labs, Compustat
Sales/Employees	Trailing twelve-month sales divided by Revelio Labs distinct employees, averaged over firm-quarters before 2022Q4.	Compustat, Revelio Labs
Net Hiring Rate	Symmetric growth rate of Revelio Labs quarterly headcount following Belo et al. (2014), averaged over firm-quarters before 2022Q4.	Revelio Labs
Union Coverage Rate	Percentage of workers covered by a collective bargaining agreement at the Census Industry Classification level, fixed at 2022 values.	CPS via Hirsch et al. (2026)
Union Membership Rate	Percentage of workers who are union members at the Census Industry Classification level, fixed at 2022 values.	CPS via Hirsch et al. (2026)

Internet Appendix

This Internet Appendix contains supplementary material referenced in the paper. Section IA1 reports additional tables and figures. Section IA2 describes the firm-level Generative AI adoption measure, including the keyword filtering procedure, the LLM classification prompt and model selection, robustness to alternative classifiers, and illustrative examples. Section IA3 documents the construction of the two instrumental variables. Section IA4 describes the ex-ante Generative AI exposure measure, including the O*NET task classification, occupation-level and firm-level aggregation, and summary statistics by industry.

IA1 Additional Tables and Figures

Figure IA1: **Monthly Generative AI Job Postings by Role**

Figure IA1 plots the monthly number of Generative AI job postings from January 2020 through mid-2025, decomposed by the top seven job roles. Each bar represents the total number of Generative AI integrator postings in a given month, with colored segments indicating the contribution of each role. Roles are ranked by total posting volume across the full sample period; postings outside the top seven roles are grouped as Other. Role classifications are based on the Revelio Labs K-50 taxonomy. The vertical dashed line marks the first public release of ChatGPT (GPT-3.5) in November 2022.

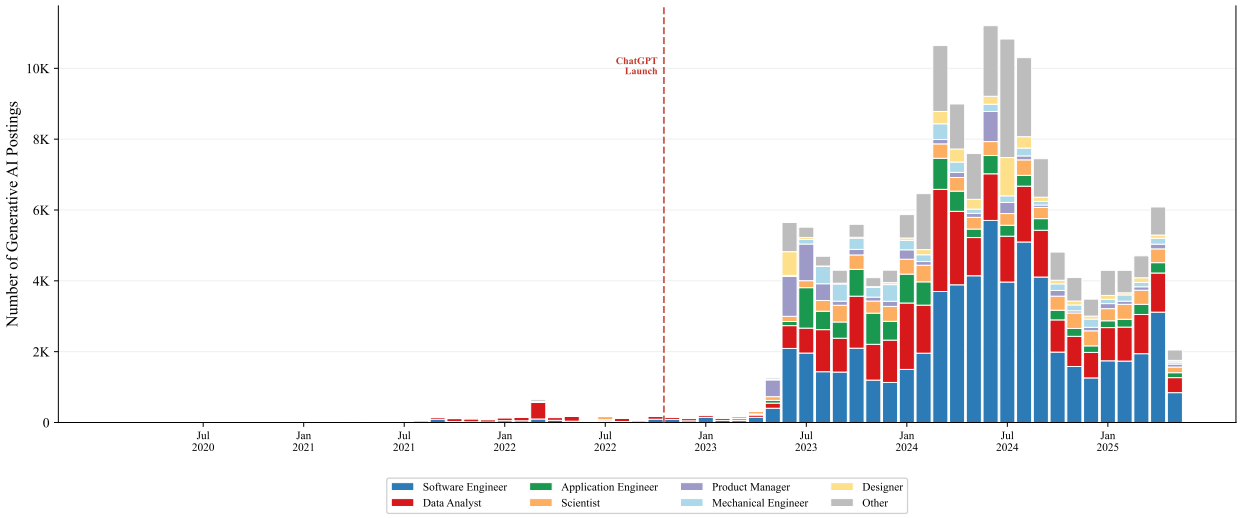


Figure IA2: **Generative AI Adoption Rate by Industry**

Figure IA2 reports the share of firms that post at least one Generative AI integrator role during the sample period, by three-digit NAICS subsector. The sample includes all U.S. public firms matched to Revelio Labs posting data. The top 30 subsectors with at least five firms are shown, sorted by adoption rate.

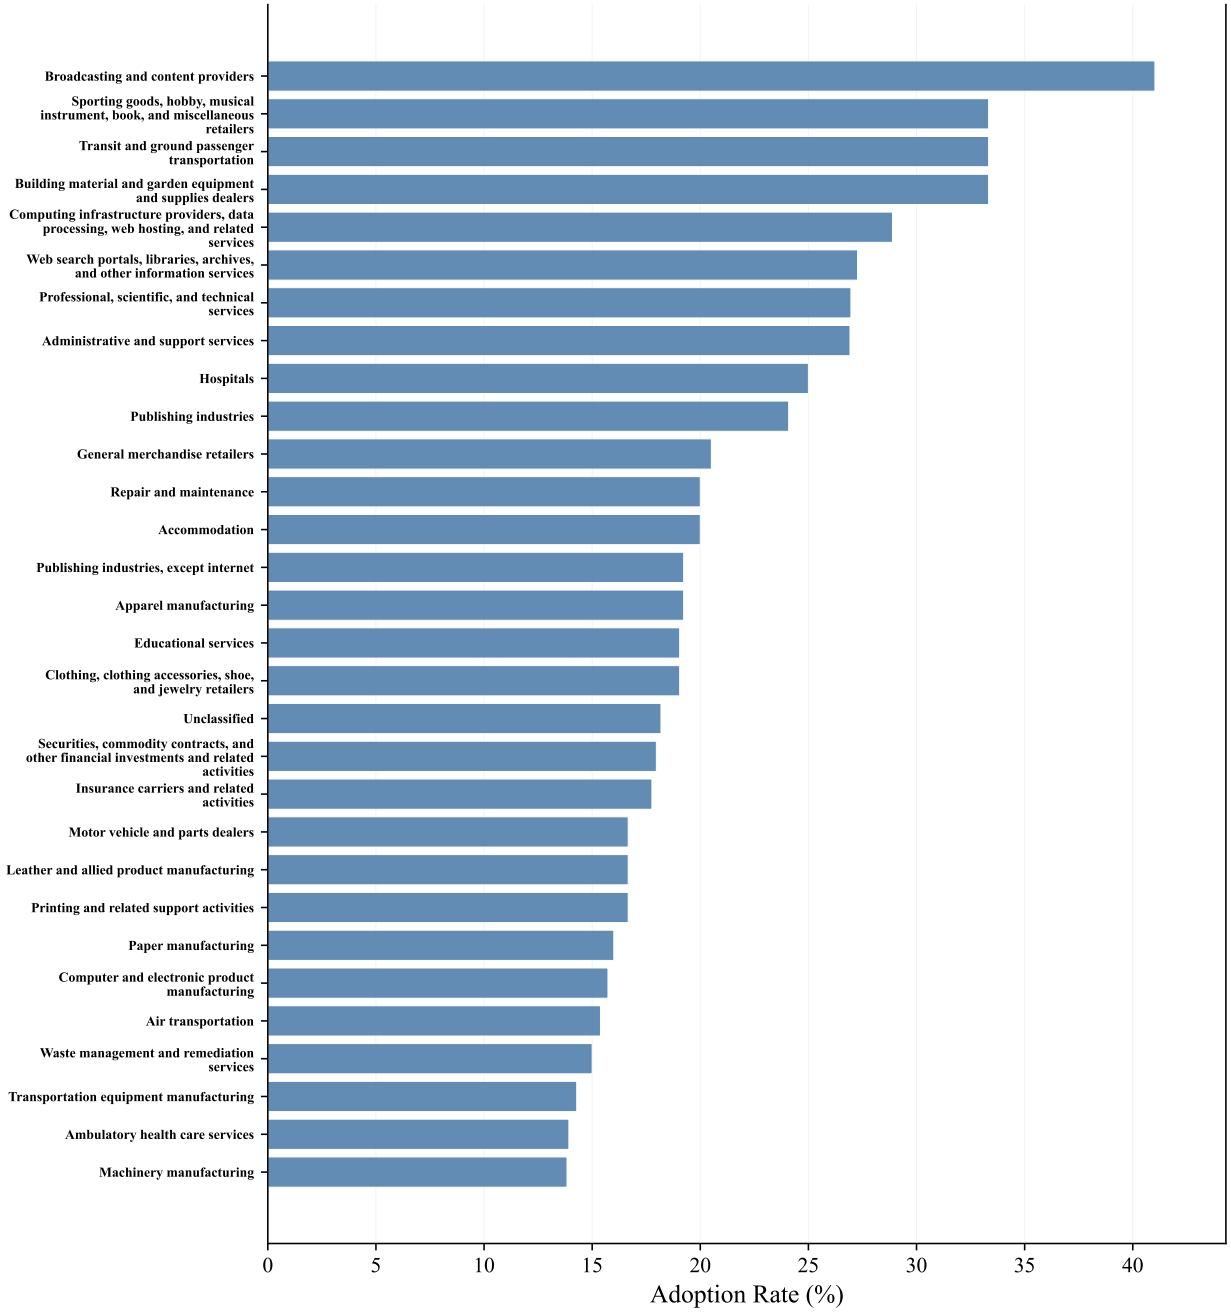


Figure IA3: **Robustness to Alternative Industry Definitions**

Figure IA3 reports coefficient estimates and 95% confidence intervals for the effect of Generative AI adoption on leverage under three alternative industry-by-quarter fixed effect definitions: SIC two-digit, NAICS three-digit, and Fama-French 48. Panel A reports results for book leverage and Panel B for the natural logarithm of one plus interest coverage. Panel C through Panel E extend the analysis to market leverage, long-term leverage, and net debt to assets. All dependent variables are measured one quarter ahead. Each specification includes firm fixed effects, firm-level controls (log assets, tangibility, market-to-book, profitability, and R&D), and standard errors clustered at the firm level. The dashed vertical line marks zero.

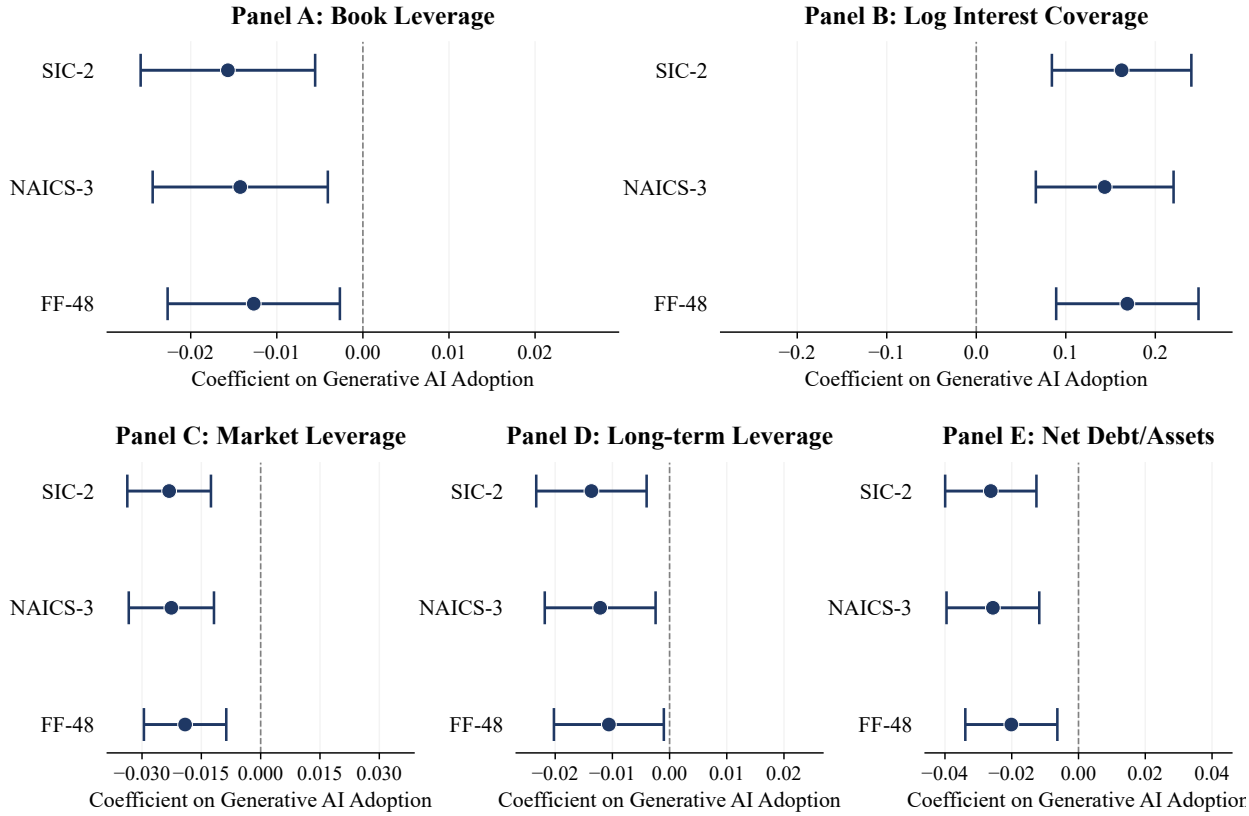
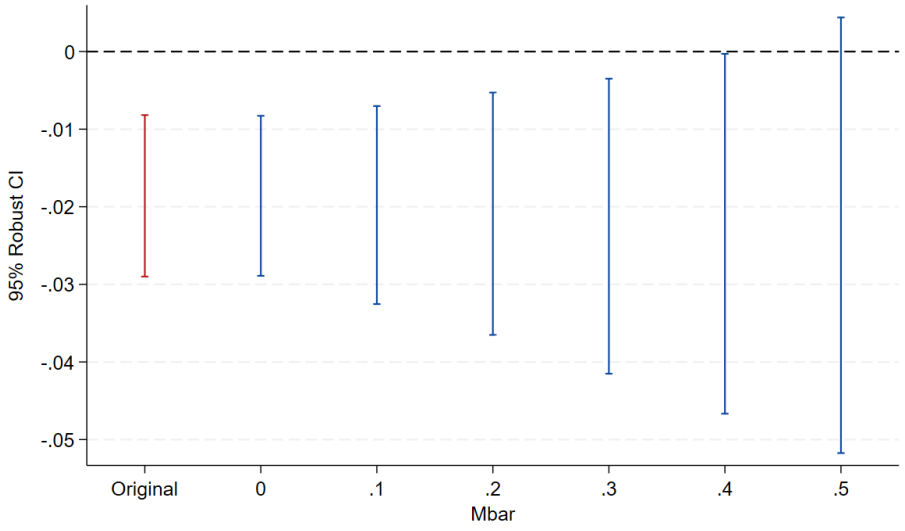
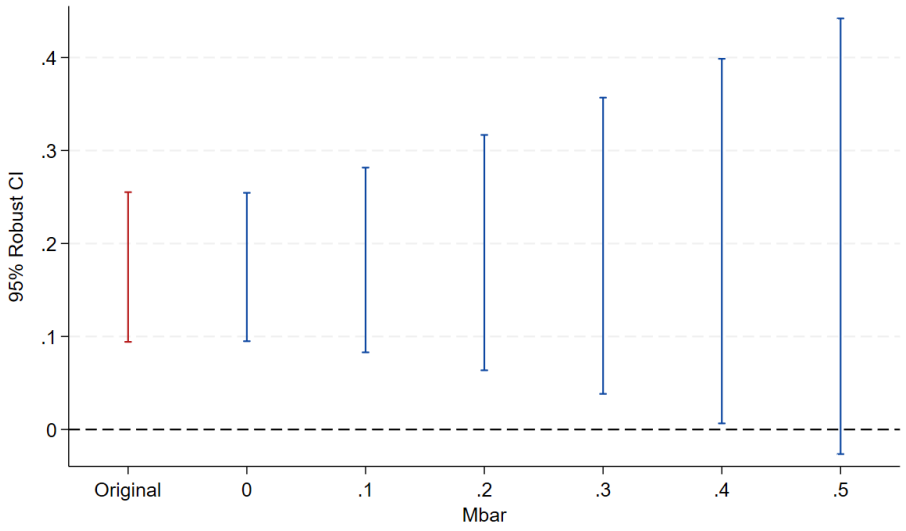


Figure IA4: **Sensitivity to Violations of Parallel Trends**

Figure IA4 reports the sensitivity of average post-treatment effects to violations of the parallel trends assumption, following Rambachan and Roth (2023). The horizontal axis measures $\bar{M}$, the maximum post-treatment violation of parallel trends as a fraction of the largest pre-treatment violation. At $\bar{M} = 0$, exact parallel trends are imposed. As $\bar{M}$ increases, larger deviations are permitted. The shaded region shows the 95% robust confidence interval at each $\bar{M}$. Estimates are based on the imputation estimator of Borusyak et al. (2024). The average effect is constructed as an equally-weighted average of post-treatment period effects (quarters 0 through 5 after adoption).



(a) Book Leverage



(b) Log Interest Coverage

Figure IA5: **Alternative Classification Models: Coefficient Estimates**

Figure IA5 reports coefficient estimates and 95% confidence intervals for the effect of Generative AI adoption on five leverage measures under three alternative classification approaches. GPT (Post) refers to the baseline GPT-4o-mini classification restricted to postings published after the public release of ChatGPT (2022Q4 onward). Llama refers to classification using Llama-3.1-70B, and Gemini refers to Gemini-2.5-Flash. Panel A reports results for book leverage, Panel B for the natural logarithm of one plus interest coverage, Panel C for market leverage, Panel D for long-term leverage, and Panel E for net debt to assets. All specifications include firm fixed effects, SIC two-digit industry-by-year-quarter fixed effects, firm-level controls, and standard errors clustered at the firm level. Dependent variables are measured one quarter ahead. The dashed vertical line marks zero.

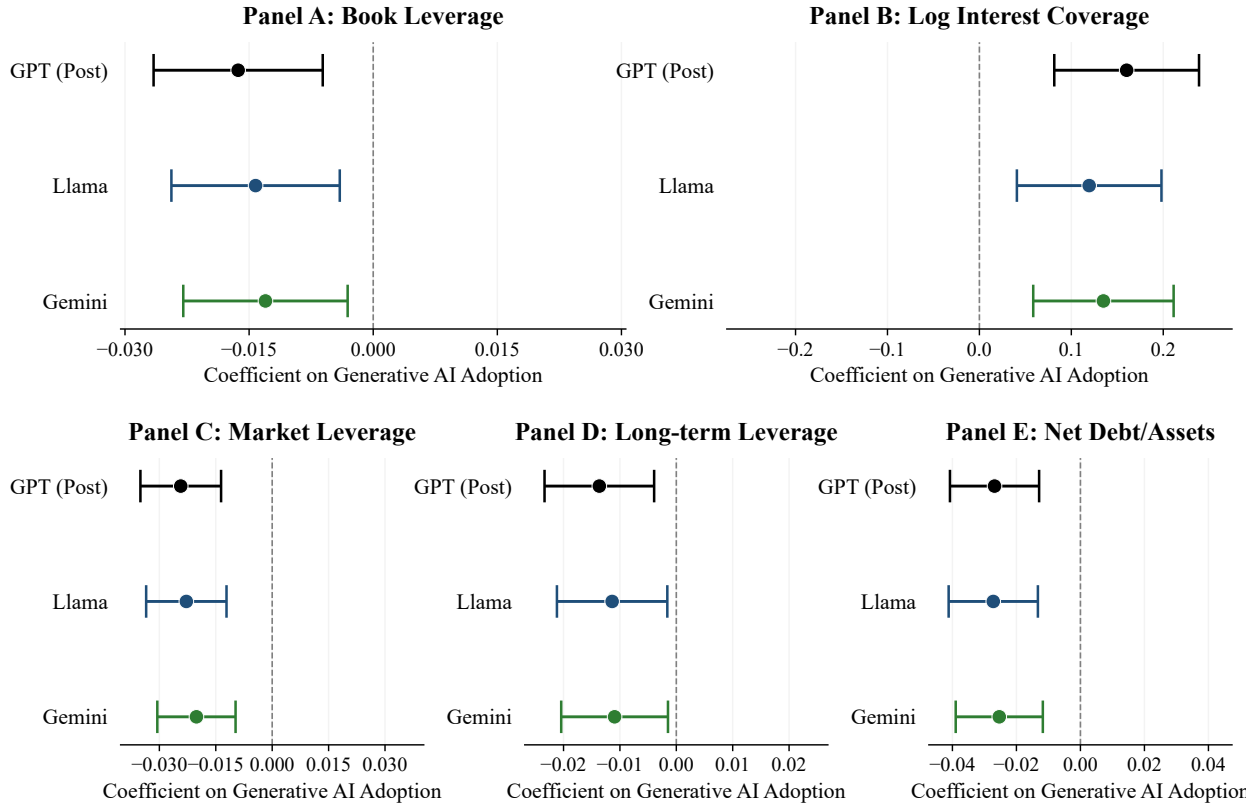
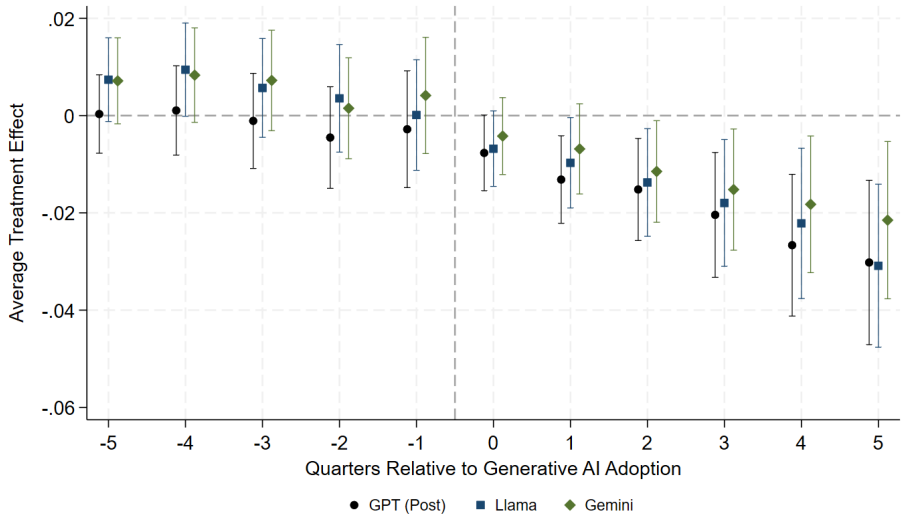
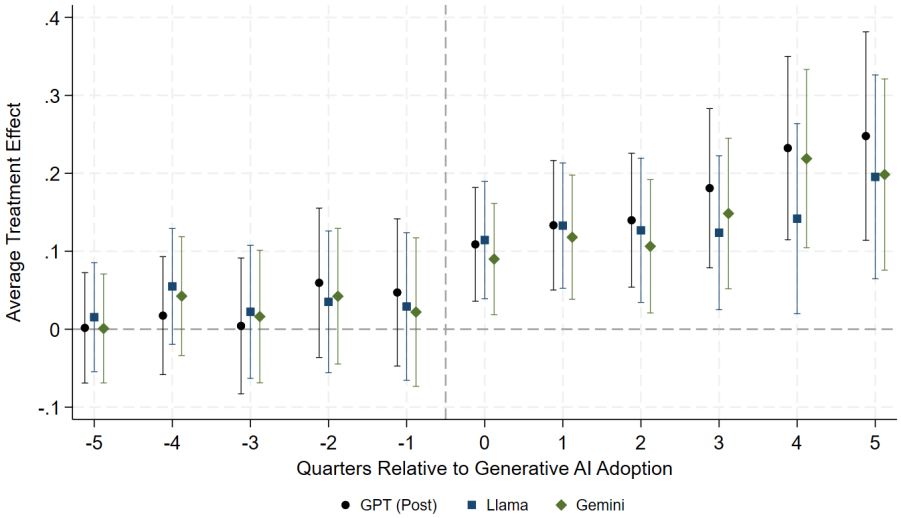


Figure IA6: **Alternative Classification Models: Event Study**

Figure IA6 reports event study estimates under three alternative classification approaches. We estimate event studies under all five heterogeneity-robust estimators reported in the main text; the figure presents Borusyak et al. (2024) as an example; other event study estimators yield similar results. GPT (Post) refers to the baseline GPT-4o-mini classification restricted to postings published after the public release of ChatGPT (2022Q4 onward). Llama refers to classification using Llama-3.1-70B, and Gemini refers to Gemini-2.5-Flash. The horizontal axis measures quarters relative to a firm's first Generative AI job posting. Panel A reports results for book leverage and Panel B for the natural logarithm of one plus interest coverage. All specifications include firm fixed effects, SIC two-digit industry-by-year-quarter fixed effects, firm-level controls, and standard errors clustered at the firm level. Dependent variables are measured one quarter ahead. The dashed vertical line separates the pre- and post-adoption periods.



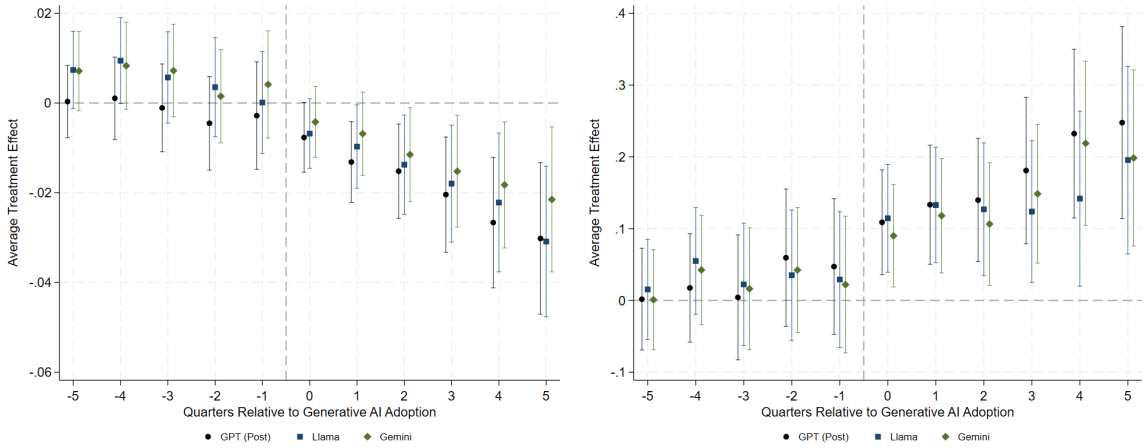
(a) Book Leverage



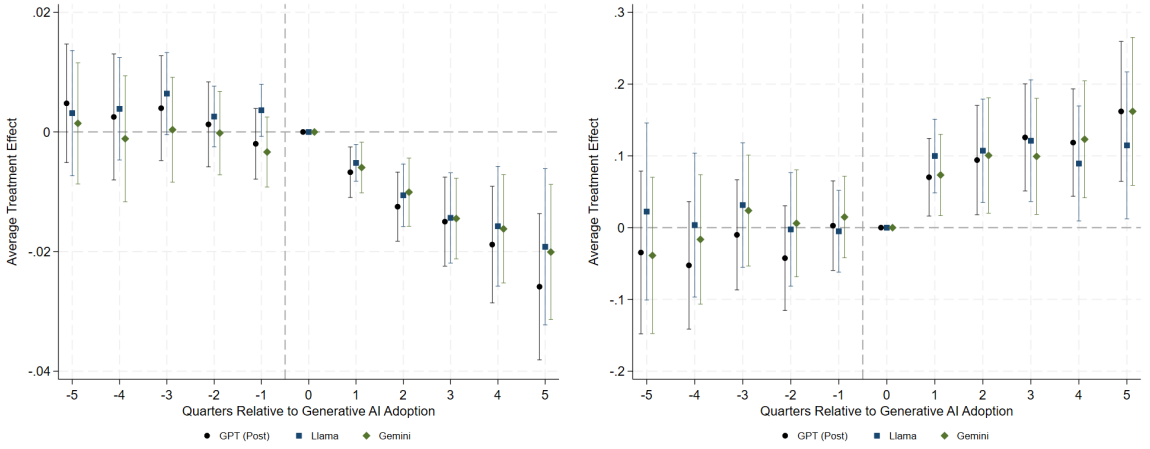
(b) Log Interest Coverage

Figure IA7: **Alternative Classification Models: Event Study**

Figure IA7 reports event study estimates under three alternative classification approaches. We estimate event studies under all five heterogeneity-robust estimators reported in the main text; the figure presents Borusyak et al. (2024) and De Chaisemartin and d'Haultfoeuille (2024) as examples; other estimators yield similar results. GPT (Post) refers to the baseline GPT-4o-mini classification restricted to postings published after the public release of ChatGPT (2022Q4 onward). Llama refers to classification using Llama-3.1-70B, and Gemini refers to Gemini-2.5-Flash. The horizontal axis measures quarters relative to a firm's first Generative AI job posting. Within each panel, the left subplot reports results for book leverage and the right subplot for the natural logarithm of one plus interest coverage. All specifications include firm fixed effects, SIC two-digit industry-by-year-quarter fixed effects, firm-level controls, and standard errors clustered at the firm level. Dependent variables are measured one quarter ahead. The dashed vertical line separates the pre- and post-adoption periods.



(a) Borusyak et al. (2024)



(b) De Chaisemartin and d'Haultfoeuille (2024)

Table IA1: **Generative AI Adoption and Alternative Leverage Measures**

Table IA1 reports difference-in-differences estimates of the effect of Generative AI adoption on alternative leverage measures. In each panel, the dependent variable is the one-quarter-ahead value of the measure indicated in the column header. Market leverage is total debt divided by total debt plus market value of equity. Long-term leverage is long-term debt divided by total assets. Net debt to assets is total debt minus cash divided by total assets. Generative AI Adoption is an indicator equal to one for firm-quarters at and after the first Generative AI integrator job posting. Panel A includes firm and year-quarter fixed effects. Panel B replaces year-quarter fixed effects with industry-by-year-quarter fixed effects to absorb time-varying industry shocks. For each measure, we report results without controls and with firm-level controls. Controls include Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Market Leverage		*LT Leverage*		*Net Debt/Assets*	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Year-Quarter Fixed Effects						
Generative AI Adoption	-0.017***	-0.012**	-0.014***	-0.012**	-0.029***	-0.019***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.007)	(0.007)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-level Controls	No	Yes	No	Yes	No	Yes
Observations	41,007	41,007	41,007	41,007	41,007	41,007
Adjusted R^2	0.808	0.821	0.845	0.849	0.869	0.883
Panel B: Industry By Year-Quarter Fixed Effects						
Generative AI Adoption	-0.026***	-0.023***	-0.015***	-0.014***	-0.032***	-0.026***
	(0.006)	(0.005)	(0.005)	(0.005)	(0.007)	(0.007)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Ind × YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-level Controls	No	Yes	No	Yes	No	Yes
Observations	41,007	41,007	41,007	41,007	41,007	41,007
Adjusted R^2	0.824	0.834	0.847	0.851	0.873	0.885

Table IA2: **Robustness: Including Firms with Zero Total Debt**

Table IA2 reports the baseline regressions on a sample that retains firms with zero total debt, which the main analysis (Table 2) excludes. The dependent variable in Columns (1)–(3) is one-quarter-ahead book leverage and in Columns (4)–(6) is one-quarter-ahead log interest coverage. Generative AI Adoption is an indicator equal to one for firm-quarters at and after the first Generative AI integrator job posting. Columns (1), (2), (4), and (5) include firm and year-quarter fixed effects; Columns (3) and (6) replace year-quarter fixed effects with industry-by-year-quarter fixed effects. Columns (2), (3), (5), and (6) add firm-level controls: Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Log interest coverage is set to zero for firm-quarters with non-positive interest expense or negative operating income, following the convention of Faulkender and Petersen (2006) and Agrawal and Matsa (2013). Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Book Leverage			*Log(Interest Coverage)*		
	(1)	(2)	(3)	(4)	(5)	(6)
Generative AI Adoption	-0.016^{***}	-0.011^{**}	-0.014^{***}	0.177^{***}	0.159^{***}	0.162^{***}
	(0.005)	(0.005)	(0.005)	(0.043)	(0.041)	(0.040)
Observations	46,143	46,143	46,143	43,311	43,311	43,311
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	No	Yes	Yes	No
Ind × YearQuarter FE	No	No	Yes	No	No	Yes
Firm-level Controls	No	Yes	Yes	No	Yes	Yes
Adjusted R^2	0.847	0.856	0.857	0.721	0.732	0.743

Table IA3: **Matching Estimators**

Table IA3 reports difference-in-differences estimates of the effect of Generative AI adoption on leverage using two matching approaches. The dependent variable is the one-quarter-ahead value of the measure indicated in the column header. Book Leverage is total debt divided by total assets. Log(Interest Coverage) is the natural logarithm of one plus EBITDA divided by interest expense. Generative AI Adoption equals one for firm-quarters after the firm's first Generative AI job posting. Panel A uses entropy balancing (EB), which reweights control firms to match the covariate distribution of treated firms at 2022Q3 (pre-ChatGPT release). Balancing covariates are Log(Assets), Tangibility, Market-to-Book, Profitability, R&D/Assets, and Book Leverage; columns (2), (4), (6), and (8) additionally balance on the pre-treatment leverage trend (four-quarter change in leverage). "EB Targets" indicates the number of moments balanced: 2 balances means and variances; 1 balances means only. EB weights are applied to the full firm-quarter panel, and regressions include firm and industry×year-quarter fixed effects. Panel B uses propensity score matching (PSM) under a stacked cohort design: for each adoption quarter, treated firms are matched to never-adopting firms at the quarter immediately before adoption using nearest-neighbor matching on Log(Assets), Tangibility, Market-to-Book, Profitability, R&D/Assets, and Book Leverage. "NN" is the number of nearest neighbors; "Within-Industry" indicates whether matching is restricted to the same two-digit SIC industry. Matched cohort-specific subpanels spanning five quarters before through five quarters after adoption are stacked, and regressions include firm×cohort and cohort×year-quarter fixed effects, which absorb time-invariant firm heterogeneity within each cohort and cohort-specific aggregate time shocks. Controls in both panels are Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D/Assets. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Entropy Balancing

	Book Leverage				*Log(Interest Coverage)*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Generative AI Adoption	−0.016***	−0.020***	−0.014**	−0.016***	0.124***	0.137***	0.155***	0.161***
	(0.005)	(0.006)	(0.006)	(0.006)	(0.040)	(0.043)	(0.039)	(0.041)
Lev. Trend Balanced	No	Yes	No	Yes	No	Yes	No	Yes
EB Targets	2	2	1	1	2	2	1	1
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ind × YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	36,835	32,466	36,835	32,466	36,835	32,466	36,835	32,466
Adjusted R^2	0.890	0.891	0.884	0.891	0.776	0.767	0.766	0.758

Panel B: Stacked Cohort PSM

	Book Leverage				*Log(Interest Coverage)*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Generative AI Adoption	−0.009**	−0.016**	−0.022**	−0.027***	0.117***	0.117**	0.170***	0.164**
	(0.005)	(0.007)	(0.009)	(0.010)	(0.037)	(0.055)	(0.065)	(0.070)
NN	5	5	1	1	5	5	1	1
Within-Industry	No	Yes	No	Yes	No	Yes	No	Yes
Firm × Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cohort × YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	19,056	8,175	4,170	3,818	19,056	8,175	4,170	3,818
Adjusted R^2	0.934	0.928	0.944	0.941	0.828	0.804	0.813	0.808

Table IA4: **Entropy Balancing: Covariate Balance**

Table IA4 compares covariate means between treated and control firms before and after entropy balancing. The sample is a cross-section at 2022Q3, the quarter before the public release of ChatGPT. Treated firms are those that post at least one Generative AI integrator job at any point during the sample period. Control firms never post a Generative AI integrator job. Entropy balancing reweights control observations so that the reweighted control distribution matches the first two moments (means and variances) of the treated distribution on all listed covariates. Log(Assets) is the natural logarithm of total assets. Asset Tangibility is net property, plant, and equipment divided by total assets. Market-to-Book is the market value of assets divided by the book value of assets. Operating ROA is operating income before depreciation divided by total assets. R&D/Assets is research and development expenditure divided by total assets. Book Leverage is total debt divided by total assets. Leverage Trend (Δ_4) is the four-quarter change in book leverage, included to balance pre-treatment leverage dynamics. Standardized differences (Std. Diff.) are computed as the difference in means divided by the pooled standard deviation, expressed as a percentage; values below 1% indicate near-exact balance.

	Treated	Control Mean		Std. Diff. (%)	
	Mean	Raw	Balanced	Raw	Balanced
Log(Assets)	8.771	6.565	8.768	123.5	0.2
Asset Tangibility	0.169	0.243	0.169	-39.0	0.0
Market-to-Book	4.796	2.359	4.804	24.2	-0.1
Operating ROA	0.021	-0.004	0.021	41.1	0.2
R&D/Assets	0.015	0.018	0.015	-9.3	-0.2
Book Leverage	0.367	0.328	0.367	16.2	-0.0
Leverage Trend (Δ_4)	0.019	0.026	0.019	-6.6	-0.2
Observations	458	1,211			

Table IA5: **Synthetic Difference-in-Differences**

Table IA5 reports average treatment effect on the treated (ATT) estimates from the synthetic difference-in-differences (SDID) estimator of Arkhangelsky et al. (2021). SDID assigns data-driven unit and time weights to construct a synthetic counterfactual that matches the pre-treatment outcome trajectory of treated firms, subsuming conventional unit and time fixed effects. The estimator requires a strongly balanced panel; we restrict the sample to firms with complete quarterly coverage throughout the sample period. Under staggered adoption, SDID computes cohort-specific treatment effects for each adoption quarter and aggregates them into a weighted ATT. Panel A reports results for the two primary leverage measures. Panel B reports results for three alternative leverage measures. The dependent variable is the one-quarter-ahead value of the measure indicated in the column header. Book leverage is total debt divided by total assets. Log(Interest Coverage) is the natural logarithm of one plus EBITDA divided by interest expense. Market leverage is total debt divided by total debt plus market value of equity. Long-term leverage is long-term debt divided by total assets. Net debt to assets is total debt minus cash divided by total assets. Generative AI Adoption is an indicator equal to one for firm-quarters at and after the first Generative AI integrator job posting. Odd-numbered columns report estimates without controls. Even-numbered columns include firm-level controls: Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D, entered via the projected method, which residualizes the outcome using parameters estimated from untreated observations only. Bootstrap standard errors (100 replications), shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Primary Leverage Measures

	Book Leverage		*Log(Interest Coverage)*	
	(1)	(2)	(3)	(4)
Generative AI Adoption	−0.015***	−0.014***	0.217***	0.195***
	(0.005)	(0.004)	(0.045)	(0.045)
Firm-level Controls	No	Yes	No	Yes
Observations	21,420	21,420	21,420	21,420

Panel B: Alternative Leverage Measures

	Market Leverage		*LT Leverage*		*Net Debt/Assets*	
	(1)	(2)	(3)	(4)	(5)	(6)
Generative AI Adoption	−0.017***	−0.017***	−0.014***	−0.014***	−0.016**	−0.015***
	(0.005)	(0.005)	(0.005)	(0.004)	(0.006)	(0.006)
Firm-level Controls	No	Yes	No	Yes	No	Yes
Observations	21,420	21,420	21,420	21,420	21,420	21,420

Table IA6: **Pre-Treatment Placebo Test**

Table IA6 tests whether firms that later adopt Generative AI were already on a differential leverage trajectory before adoption. For each adopting firm, we assign a placebo adoption date four quarters (one year) before the firm's actual first Generative AI integrator posting. The sample is then restricted to the pre-adoption period: for adopting firms, we retain only firm-quarter observations before the actual adoption quarter; for never-adopting firms, we retain all observations. The placebo treatment indicator equals one for adopter firm-quarters at or after the placebo date and zero otherwise. Because the sample excludes the post-adoption period, the actual Generative AI treatment never turns on. The dependent variable is the one-quarter-ahead value of the measure indicated in the column header. Book Leverage is total debt divided by total assets. Log(Interest Coverage) is the natural logarithm of one plus EBITDA divided by interest expense. Columns (1) and (4) include firm and year-quarter fixed effects without controls. Columns (2) and (5) add firm-level controls: Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Columns (3) and (6) replace year-quarter fixed effects with industry-by-year-quarter fixed effects. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Book Leverage			*Log(Interest Coverage)*		
	(1)	(2)	(3)	(4)	(5)	(6)
Placebo Generative AI Adoption $(t-4)$	−0.001	−0.000	−0.002	0.024	0.020	0.055
	(0.004)	(0.004)	(0.004)	(0.036)	(0.034)	(0.035)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	No	Yes	Yes	No
Ind × YearQuarter FE	No	No	Yes	No	No	Yes
Firm-level Controls	No	Yes	Yes	No	Yes	Yes
Observations	34,521	34,521	34,513	34,521	34,521	34,513
Adjusted R^2	0.856	0.864	0.866	0.733	0.745	0.757

Table IA7: **Realized-Hiring Validation**

Table IA7 re-estimates the baseline difference-in-differences on adopters who hire. A Generative AI hire is a new Revelio Labs position whose title or description contains a Generative AI keyword from the paper's keyword set, with a start date between 2020Q1 and 2025Q2, attributed to a firm directly or through its ultimate parent. Treated firms are adopters with such a hire in the adoption quarter or the following quarter (t or t+1), leaving never-adopters as controls. The dependent variable is the one-quarter-ahead value of the measure indicated in the column header. Book Leverage is total debt divided by total assets. Log(Interest Coverage) is the natural logarithm of one plus EBITDA divided by interest expense. Generative AI Adoption is an indicator equal to one for firm-quarters at and after the first Generative AI integrator job posting. Columns (1) and (4) include firm and year-quarter fixed effects without controls. Columns (2) and (5) add firm-level controls: Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Columns (3) and (6) replace year-quarter with industry-by-year-quarter fixed effects. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Adopters with a Realized Hire (t or t+1)					
	Book Leverage			*Log(Interest Coverage)*		
	(1)	(2)	(3)	(4)	(5)	(6)
Generative AI Adoption	−0.022***	−0.016***	−0.020***	0.200***	0.177***	0.176***
	(0.006)	(0.006)	(0.006)	(0.056)	(0.053)	(0.052)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	No	Yes	Yes	No
Ind × YearQuarter FE	No	No	Yes	No	No	Yes
Firm-level Controls	No	Yes	Yes	No	Yes	Yes
Observations	36,045	36,045	36,041	36,045	36,045	36,041
Adjusted R^2	0.842	0.852	0.853	0.722	0.735	0.746

Table IA8: **Bartik IV Robustness: Excluding Technology Firms**

Table IA8 examines whether the Bartik IV estimates from Table 4 are robust to excluding firms that may directly contribute to LLM development, which would threaten the exclusion restriction. The instrument is $Z_{i,t} = \sum_s \text{Share}(g_i, s, t-1) \times \text{Alignment}(s, t)$, where $\text{Share}(g_i, s, t-1)$ is the employment share of occupation s in firm i's three-digit NAICS industry from the prior year's BLS Occupational Employment and Wage Statistics survey, and $\text{Alignment}(s, t)$ is the TF-IDF cosine similarity between O*NET occupation text and cumulative large-scale LLM capability text from the Epoch AI database. Panel A excludes eight firms that have publicly released proprietary foundation models or provide the primary computing infrastructure for LLM training: Alphabet, Amazon, Apple, IBM, Meta, Nvidia, Oracle, and Salesforce. Panel B excludes all firms in NAICS three-digit industries 511 (Software Publishers), 518 (Computing Infrastructure and Data Processing), and 519 (Web Search Portals and Other Information Services), which encompass the primary sectors where LLM developers are classified. Columns (1)–(5) use the O*NET Work Activities / LLM Task-Labels alignment; columns (6)–(10) use the O*NET Work Activities / LLM (Task + Domain) alignment. Within each alignment, the first two columns report reduced-form estimates, the third reports the first stage, where the dependent variable is Generative AI Adoption, and the last two report second-stage estimates, where the dependent variable is the one-quarter-ahead value of the indicated measure. All specifications include firm and year-quarter fixed effects and firm-level controls: Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Standard errors clustered at the firm level in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

	Work Activities / Task-Labels					*Work Activities / (Task + Domain)*				
	Reduced Form		First Stage	Second Stage		Reduced Form		First Stage	Second Stage	
	Book Leverage	*Log(Interest Coverage)*	*Generative AI Adoption*	*Book Leverage*	*Log(Interest Coverage)*	*Book Leverage*	*Log(Interest Coverage)*	*Generative AI Adoption*	*Book Leverage*	*Log(Interest Coverage)*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: Excluding LLM Developers (8 firms)										
LLM Capability Bartik	−0.006**	0.111***	0.045***			−0.007**	0.112***	0.042***		
	(0.003)	(0.022)	(0.008)			(0.003)	(0.022)	(0.008)		
Generative $\widehat{\text{AI}}$ Adoption				−0.139**	2.465***				−0.173**	2.654***
				(0.069)	(0.643)				(0.074)	(0.695)
Kleibergen-Paap F			34.4					31.4		
Adjusted R^2	0.855	0.733	0.460	−0.008	−0.271	0.855	0.733	0.460	−0.048	−0.325
N	40,677	40,677	40,677	40,677	40,677	40,677	40,677	40,677	40,677	40,677
Panel B: Excluding Software, Cloud, and Internet Industries (NAICS 511, 518, 519)										
LLM Capability Bartik	−0.009***	0.068***	0.029***			−0.009***	0.072***	0.027***		
	(0.003)	(0.021)	(0.007)			(0.003)	(0.021)	(0.007)		
Generative $\widehat{\text{AI}}$ Adoption				−0.289**	2.322**				−0.345**	2.672**
				(0.118)	(0.946)				(0.135)	(1.078)
Kleibergen-Paap F			16.1					14.3		
Adjusted R^2	0.856	0.735	0.438	−0.231	−0.226	0.856	0.735	0.438	−0.363	−0.319
N	37,191	37,191	37,191	37,191	37,191	37,191	37,191	37,191	37,191	37,191
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table IA9: **IV Estimates: AI Research Proximity**

Table IA9 reports instrumental variable estimates using geographic proximity to AI research universities. The instrument is $Z_{it} = \log\left(\sum_{u \in \mathcal{U}} \mathbb{1}[d_{iu} \leq \bar{d}] \cdot R_u\right) \times \text{Post}_t$, where d_{iu} is the distance between firm i's headquarters and university u, R_u is university u's research output, $\mathcal{U}$ is the set of top-ranked US research universities, and $\text{Post}_t = \mathbb{1}[t \geq 2023\text{Q1}]$. We set $\bar{d} = 250$ miles and rank universities by publication count in leading AI and computer science venues and journals from the Dimensions database, 2017–2022 (see Table IA18). Panel A counts only first-author papers at each institution, which proxies primary research contribution; Panel B counts all papers. Column (1) reports first-stage estimates where the dependent variable is Generative AI Adoption. Columns (2) and (3) report second-stage 2SLS estimates where the dependent variable is the one-quarter-ahead value of the indicated measure. All specifications include firm and industry-by-year-quarter fixed effects and firm-level controls: Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Standard errors clustered at the firm level in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

	First Stage	2SLS	
	Generative AI Adoption	*Book Leverage*	*Log(Interest Coverage)*
	(1)	(2)	(3)
Panel A: First-Author			
AI Research Proximity	0.020***		
	(0.004)		
$\widehat{\text{Generative AI}}$ Adoption		−0.197**	1.705**
		(0.099)	(0.695)
Kleibergen-Paap F	26.6		
Adjusted R^2	0.502	−0.117	−0.142
Panel B: All Papers			
AI Research Proximity	0.017***		
	(0.003)		
$\widehat{\text{Generative AI}}$ Adoption		−0.189**	1.702**
		(0.095)	(0.728)
Kleibergen-Paap F	26.7		
Adjusted R^2	0.502	−0.105	−0.141
N	40,326	40,326	40,326
Firm FE	Yes	Yes	Yes
Ind × YearQuarter FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes

Table IA10: **Generative AI Exposure, Layoffs, and Capital Structure**

Table IA10 splits firms by ex-ante Generative AI exposure, measured as of 2022Q3, the quarter before ChatGPT's public release. We aggregate the occupation-level exposure scores of Eisfeldt et al. (2023) to the firm as a positions-weighted average over its Revelio Labs positions, and split firms at the median. Panel A reports layoff outcomes: Has WARN is an indicator for filing a WARN Act notice in the quarter, and Log(Workers) is the natural logarithm of one plus the number of affected workers. Panel B reports capital structure outcomes: Book Leverage is total debt divided by total assets, and Log(Interest Coverage) is the natural logarithm of one plus EBITDA divided by interest expense, each measured one quarter ahead. Both panels use the same firms and the same split. All columns include firm and year-quarter fixed effects; even-numbered columns add the firm-level controls Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Layoffs

	High Exposure			*Low Exposure*		
	Has WARN	*Mass (500+)*	*Log(1+Workers)*	*Has WARN*	*Mass (500+)*	*Log(1+Workers)*
	(1)	(2)	(3)	(4)	(5)	(6)
Generative AI Adoption	0.019**	0.004**	0.076**	0.010	0.011	0.040
	(0.008)	(0.002)	(0.032)	(0.012)	(0.007)	(0.045)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-level Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	18,654	18,654	18,654	22,353	22,353	22,353
Adjusted R^2	0.137	0.062	0.135	0.142	0.064	0.144

Panel B: Capital Structure

	High Exposure				*Low Exposure*			
	Book Leverage		*Log(Interest Coverage)*		*Book Leverage*		*Log(Interest Coverage)*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Generative AI Adoption	−0.025***	−0.017**	0.172***	0.171***	−0.015**	−0.011*	0.133**	0.083
	(0.008)	(0.007)	(0.057)	(0.055)	(0.007)	(0.007)	(0.063)	(0.057)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-level Controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	18,654	18,654	18,654	18,654	22,353	22,353	22,353	22,353
Adjusted R^2	0.835	0.844	0.745	0.751	0.853	0.863	0.683	0.705

Table IA11: **Knowledge-Worker Share and Capital Structure**

Table IA11 reports subsample difference-in-differences estimates split by Knowledge-Worker Share, measured as Revelio Labs observed headcount divided by the firm's Compustat-reported employment. Online professional profiles capture white-collar workers far better than manual workers, so a higher ratio identifies a more knowledge-worker-intensive workforce, which is the labor Generative AI substitutes for. The split variable is each firm's average over firm-quarters before 2022Q4, and firms are divided at the median. The dependent variable is the one-quarter-ahead value of the measure indicated in the column header. Generative AI Adoption is an indicator equal to one for firm-quarters at and after the first Generative AI integrator job posting. All columns include firm and year-quarter fixed effects. Controls are Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Book Leverage		*Log(Interest Coverage)*		*Book Leverage*		*Log(Interest Coverage)*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	High Knowledge-Worker Share				Low Knowledge-Worker Share			
Generative AI Adoption	-0.030***	-0.020**	0.168**	0.155**	-0.008	-0.006	0.154***	0.120***
	(0.009)	(0.008)	(0.066)	(0.064)	(0.006)	(0.006)	(0.050)	(0.044)
Observations	19,739	19,739	19,739	19,739	19,713	19,713	19,713	19,713
Adjusted R^2	0.820	0.831	0.701	0.711	0.886	0.892	0.680	0.708
Firm-level Controls	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table IA12: **Leverage Reduction: Debt Levels and Interest Expense**

Table IA12 reports estimates of the effect of Generative AI adoption on debt levels and interest expense at horizons one through four quarters. Both dependent variables are scaled by lagged total assets at t. Panel A uses total debt at $t+h$ as the numerator. Panel B uses quarterly interest expense at $t+h$, measured in percentage points. Generative AI Adoption is an indicator equal to one for firm-quarters at and after the first Generative AI integrator job posting. Odd-numbered columns include firm and year-quarter fixed effects without controls. Even-numbered columns add firm-level controls: Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	$h=1$		$h=2$		$h=3$		$h=4$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Total Debt at t+h / Total Assets at t								
Generative AI Adoption	−0.016***	−0.011**	−0.019***	−0.014***	−0.021***	−0.016***	−0.023***	−0.018***
	(0.005)	(0.005)	(0.006)	(0.005)	(0.007)	(0.006)	(0.008)	(0.007)
Observations	37,819	37,819	34,764	34,764	32,195	32,195	29,775	29,775
Adjusted R^2	0.824	0.838	0.811	0.827	0.799	0.821	0.788	0.818
Panel B: Interest Expense at t+h / Total Assets at t								
Generative AI Adoption	−0.056***	−0.049***	−0.059***	−0.052***	−0.055***	−0.048***	−0.053***	−0.046***
	(0.015)	(0.014)	(0.015)	(0.014)	(0.016)	(0.015)	(0.017)	(0.016)
Observations	36,251	36,251	33,421	33,421	30,969	30,969	28,667	28,667
Adjusted R^2	0.746	0.756	0.746	0.758	0.748	0.763	0.751	0.770
Firm-level Controls	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table IA13: **Profitability Mediation Test**

Table IA13 tests whether the leverage and interest coverage effects of Generative AI adoption are mediated through changes in profitability. If the mechanism were purely that adoption raises profits, which accumulate as retained earnings and mechanically reduce leverage, then controlling for profitability changes should absorb the adoption coefficient. Columns (1) and (3) add the one-quarter change in profitability to the baseline regression of Table 2; Columns (2) and (4) add the four-quarter change in profitability. Leverage is total debt divided by total assets. Log(Interest Coverage) is the natural logarithm of one plus EBITDA divided by interest expense. Δ Profitability (1Q) is the change in operating ROA from the prior quarter; Δ Profitability (4Q) is the change from four quarters prior. All specifications include firm and year-quarter fixed effects and firm-level controls: Log(Assets), Tangibility, Market-to-Book, and R&D. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Book Leverage		*Log(Interest Coverage)*	
	(1)	(2)	(3)	(4)
Generative AI Adoption	−0.014***	−0.013***	0.180***	0.163***
	(0.005)	(0.005)	(0.044)	(0.042)
Δ Profitability (1Q)	0.080***		2.274***	
	(0.026)		(0.235)	
Δ Profitability (4Q)		0.037		3.048***
		(0.039)		(0.245)
Firm FE	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	Yes	Yes
Firm-level Controls	Yes	Yes	Yes	Yes
Observations	37,819	29,775	37,819	29,775
Adjusted R^2	0.866	0.893	0.731	0.760

Table IA14: **Balance Sheet Composition and Firm Performance**

Table IA14 examines the balance sheet and performance channels through which Generative AI adoption affects leverage. Panel A tests whether adoption changes asset size, equity composition, and retained earnings flows. Log(Assets) is the natural logarithm of total assets. Book Equity/Assets is common equity divided by total assets. Δ RE/Assets is the quarterly change in retained earnings divided by beginning-of-quarter total assets. Panel B tests whether adoption changes operating performance and payout policy. ROA is operating income divided by total assets. Cash Ratio is cash and short-term investments divided by total assets. Payout/Assets is dividends plus share repurchases divided by total assets. The dependent variable is the one-quarter-ahead value of the measure indicated in the column header. Generative AI Adoption is an indicator equal to one for firm-quarters at and after the first Generative AI integrator job posting. Odd-numbered columns include firm and year-quarter fixed effects without controls. Even-numbered columns add firm-level controls: Log(Assets), Tangibility, Market-to-Book, Profitability (excluded when ROA or Cash Ratio is the dependent variable), and R&D. Standard errors, shown in parentheses, are clustered at the firm level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Log(Assets)		*Book Equity/Assets*		*Δ RE/Assets*	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Balance Sheet Composition						
Generative AI Adoption	0.017***	0.013***	0.024***	0.018***	0.005***	0.004***
	(0.005)	(0.004)	(0.007)	(0.007)	(0.001)	(0.001)
Observations	37,819	37,819	37,819	37,819	37,621	37,621
Adjusted R^2	0.996	0.996	0.836	0.841	0.645	0.649
	ROA		*Cash Ratio*		*Payout/Assets*	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel B: Operating Performance and Payout						
Generative AI Adoption	0.007***	0.007***	0.011***	0.007*	0.001	0.001*
	(0.002)	(0.001)	(0.004)	(0.004)	(0.000)	(0.000)
Observations	41,007	41,007	41,007	41,007	37,819	37,819
Adjusted R^2	0.629	0.640	0.887	0.897	0.396	0.402
Firm-level Controls	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
YearQuarter FE	Yes	Yes	Yes	Yes	Yes	Yes

IA2 Construction and Validation of the Generative AI Adoption Measure

We measure firm-level Generative AI adoption by identifying job postings that recruit workers to build or operate LLM-powered systems. The procedure has two steps: flag postings containing Generative AI-related keywords, then classify each flagged posting with an LLM to separate genuine integrator roles from incidental mentions. Similar to Hosseini Maasoum and Lichtinger (2025), who use this integrator/user framework to study Generative AI's labor market effects, we develop our own classification prompt with a conservative positive safeguards, and validate the results across three LLMs from different providers.

The underlying data come from Revelio Labs, which compiles job postings from LinkedIn and other sources. Each posting includes a firm identifier, posting date, and raw job description text. Our sample covers 2020Q1 through 2025Q2.

IA2.1 Identifying Generative AI Adoption

A. Step 1: Keyword Filtering

We compile a list of 27 Generative AI-related keywords spanning model names, technical infrastructure, and implementation terminology. A posting is flagged if its description contains at least one keyword, matched via case-insensitive regular expressions with word-boundary constraints to avoid partial matches. For multi-word phrases, we allow flexible whitespace and hyphens between tokens (e.g., "retrieval-augmented generation" matches "retrieval augmented generation").

```
KEYWORDS = [
    # AI Models/Systems
    "ChatGPT", "GPT", "Claude", "Anthropic Claude API", "Gemini", "Copilot",
    "large language model", "LLM", "Generative AI", "Gen AI",
    # Technical Infrastructure
    "LangChain", "LlamaIndex", "RAG", "retrieval-augmented generation",
    "vector embeddings", "vector database", "Pinecone", "Weaviate", "Milvus",
    "transformer-based model",
    # Implementation Terms
    "prompt engineering", "prompt design", "OpenAI API", "Azure OpenAI",
    "Google Vertex AI Generative", "HuggingFace Transformers", "RetrievalQA"
]
```

Temporal Filtering. We track the release date of each product and technology in our keyword list and require that a posting appear after the referenced technology became pub-

licly available. A posting that mentions "Gemini" in 2022 cannot refer to Google's model, which launched in December 2023, and a "Claude" posting from 2022 predates Anthropic's March 2023 release. As shown in the Table IA15, we manually verify the product release date of the following products: ChatGPT from 2022Q4, Claude from 2023Q1, Gemini from 2023Q4, GitHub Copilot from 2022Q2, Azure OpenAI from 2023Q1. Generic terms ("large language model," "LLM," "Generative AI") enter in 2022Q4. The OpenAI API is an exception: developers could access the GPT-3 API from June 2020, so we allow these postings from 2020Q2. For "GPT" without a version number we default to 2022Q4; when the description specifies GPT-2, GPT-3, or GPT-4, we apply the corresponding release quarter. A posting is dropped only when all of its matched keywords predate their release; if at least one keyword is valid, the posting is retained.

B. Step 2: LLM Classification

We classify each flagged posting into one of four categories:

- **Integrator:** Roles that build, deploy, or maintain LLM-powered systems, for example, developing RAG pipelines, fine-tuning models, or integrating LLM APIs into products.
- **User:** Roles that use Generative AI tools (ChatGPT, Copilot, etc.) for drafting, summarizing, or coding assistance, without building the underlying systems.
- **Both:** Roles combining integration and usage responsibilities.
- **Neither:** Postings where the keyword match is incidental or the role has no substantive Generative AI component.

Only integrator postings indicate that a firm is building Generative AI into its operations. User postings signal awareness but not active implementation.

Classification Prompt. The prompt is conservative: when evidence for integration is ambiguous, the classifier defaults to user or neither. It includes disambiguation rules for keywords with non-AI meanings in specific industries[16] and enforces a strict boundary between skill requirements and system-building. Listing "experience with GPT" or "familiarity with Transformers" does not qualify a posting as integrator; only action verbs directed at LLM systems qualify: build, develop, deploy, integrate. The full prompt is reproduced below.

```
# System message
```

[16]For example, "Gemini" (Capgemini), "RAG" (Red-Amber-Green project status), "Copilot" (aviation), "Claude" (first name), "GPT" (Topcon surveying equipment).

System: "You are an expert in Generative AI workforce analysis. Classify whether jobs involve BUILDING LLM systems (integrator) or just USING LLM tools (user). Output JSON only."
Classification prompt
Classify this job posting. Be CONSERVATIVE - only classify as INTEGRATOR if there is clear evidence of BUILDING LLM-powered systems.
Categories
INTEGRATOR: Job involves BUILDING, DEVELOPING, or DEPLOYING systems that use Large Language Models as a core component.
Strong integrator signals (must describe BUILDING, not just knowledge):
- "Build/develop/create chatbot using GPT/Claude/LLM"
- "Integrate OpenAI API / Anthropic API / Azure OpenAI into product"
- "Build RAG (retrieval-augmented generation) pipeline"
- "Develop LLM-powered application/feature"
- "Deploy LLM inference infrastructure"
- "Fine-tune LLM for production use"
- "Build with LangChain / LlamaIndex / Semantic Kernel"
- "Implement prompt engineering for product"
- "Create vector search/embeddings pipeline for LLM app"
USER: Job involves USING LLM knowledge or tools WITHOUT building LLM-powered systems. This includes:
1) End-user usage:
- "Use ChatGPT/Claude to draft content"
- "Leverage AI tools for productivity"
- "Use Copilot for coding assistance"
2) Skill/knowledge requirements (NOT building):
- "Experience with BERT, GPT, Transformers"
- "Knowledge of large language models"
- "Familiarity with HuggingFace / PyTorch / TensorFlow"
- "Background in NLP, NLU, NLG"
- ML Engineer / Data Scientist / NLP Engineer roles listing LLM as skills
BOTH: Job involves BOTH building LLM systems AND using LLM tools.
NEITHER: No clear LLM/GenAI involvement, or the matched keyword is a false positive.
FALSE POSITIVE WARNINGS - Check carefully:
- "Gemini" $\rightarrow$often Cap Gemini (consulting company), NOT Google Gemini AI
- "RAG" in project management $\rightarrow$Red-Amber-Green status reports, NOT retrieval-augmented generation
- "Copilot" in aviation/non-tech $\rightarrow$co-pilot or other software, NOT GitHub/AI Copilot
- "Claude" as person's name $\rightarrow$NOT Anthropic Claude AI
- "GPT" in surveying/construction $\rightarrow$Topcon GPT equipment, NOT language model
Critical Rule: SKILL $\neq$INTEGRATOR
"Experience with GPT/BERT/Transformers" = USER (skill requirement)
"Build chatbot using GPT API" = INTEGRATOR (actually building)

```
The difference is ACTION VERBS: build, develop, create, deploy, integrate, implement
## Classification Logic
1. Does job mention BUILDING/DEPLOYING LLM systems? →INTEGRATOR
2. Does job list LLM as skill/experience requirement? →USER
3. Does job mention using ChatGPT/Copilot as tools? →USER
4. Does job have both building AND using? →BOTH
5. Is keyword a false positive (Gemini=consulting, RAG=status)? →NEITHER
6. When uncertain →USER or NEITHER (be conservative)
## Output Format (JSON only, no prose):
{"is_integrator": boolean, "role_type": "integrator"|"user"|"both"|"neither", "confidence
": 0.0-1.0}
## Job Posting:
{description}
```

Model Selection. We use OpenAI's `gpt-4o-mini` with temperature set to zero and a fixed random seed of 66. The GPT-5 family (`gpt-5-mini`, `gpt-5-nano`) does not accept temperature or seed parameters, so output cannot be made reproducible.[17] With temperature zero the model always selects the highest-probability next token; the fixed seed controls the remaining randomness in the generation process. OpenAI describes this as "best effort" determinism: the system tries to return the same tokens given the same inputs, though server-side updates can occasionally introduce variation.[18] Since our output is a single JSON label and a confidence score, there is little room for such variation to matter. We submit requests through OpenAI's Batch API at a 50% discount relative to synchronous calls.

Knowledge Cutoff. One might worry that the classifier's training data cutoff biases results: a model trained before a given framework was released could miss postings that reference it. But our task is to classify whether a posting describes building LLM systems, not to verify whether a particular tool exists. The posting text itself says "build RAG pipeline" or "integrate OpenAI API into our product." The classifier reads these descriptions; it does not need to know what LangChain is from its own training data to recognize that a posting describes system-building.

C. Firm Adoption Definition

A firm is a Generative AI adopter if it has posted at least one integrator vacancy. The adoption date is the posting date of the firm's first integrator role.

The following examples illustrate the distinction between integrator and user postings.

[17] See https://community.openai.com/t/temperature-in-gpt-5-models/1337133. The same applies to OpenAI's o-series models (o1, o3, o4-mini).

[18] See https://cookbook.openai.com/examples/reproducible_outputs_with_the_seed_parameter.

IA2.2 Robustness

Robustness of Classification. We re-run classification using the original Hosseini Maasoum and Lichtinger (2025) prompt, which differs from ours in three ways: it favors integrator classification when signals are ambiguous ("Prefer TRUE for integrator when listed signals appear"), it lacks false-positive warnings for keywords with non-AI meanings, and it does not distinguish between skill requirements and active system-building. We also run our prompt on two additional models: `llama-3.1-8b-instant` (Meta, served via Groq) and `gemini-2.5-flash` (Google). Results are similar across all specifications, suggesting that classification reflects textual content in the posting rather than idiosyncrasies of any single model or prompt.

Table IA15: **Generative AI Keyword Temporal Filter Rules**

This table reports the earliest valid quarter for each Generative AI keyword with a product-specific release date. A job posting matching a given keyword is retained only if the posting date falls at or after the keyword's earliest valid quarter. Release dates are verified from official product announcements and documentation. All remaining keywords not listed below (e.g., large language model, LLM, Generative AI, RAG, vector database, prompt engineering) use 2022Q4, the quarter of ChatGPT's public release, as the default earliest valid quarter.

Keyword	Earliest Valid Quarter	Rationale
ChatGPT	2022Q4	Released Nov 30, 2022
GPT-2	2019Q1	Released Feb 2019
GPT-3	2020Q2	API released Jun 2020
GPT-4	2023Q1	Released Mar 14, 2023
Claude	2023Q1	Anthropic release Mar 2023
Gemini	2023Q4	Google release Dec 2023
Copilot	2022Q2	GitHub Copilot GA Jun 2022
LangChain	2022Q4	Released Oct 2022
LlamaIndex	2022Q4	Released Nov 2022
Pinecone	2021Q1	Released Jan 2021
Weaviate	2019Q1	Released 2019
Milvus	2019Q4	Released Oct 2019
OpenAI API	2020Q2	GPT-3 API access Jun 2020
Azure OpenAI	2023Q1	GA Jan 16, 2023
Google Vertex AI Generative	2023Q1	Released Mar 2023

IA2.3 Examples

Table IA16: **Example Job Postings: Generative AI Integrators and Users**

This table presents examples of job postings classified as Generative AI integrator roles (Panel A) and Generative AI user roles (Panel B). Integrator postings recruit for building, deploying, or maintaining LLM-powered systems. User postings reference Generative AI only among desired skills or as a general-purpose tool for routine tasks, without responsibilities for building such systems.

Posting Details	Job Description (Excerpt)
Panel A: Integrator Postings	
JPMorgan Chase May 4, 2025 *LLM/SLM Specialist*	(Job ID: 4212806281) **LLM**/SLM Specialist to drive the development, deployment, and optimization of **large language models** (LLMs) and small language models (SLMs). You will be responsible for training, fine-tuning, evaluating, and integrating language models into various applications using Infinite AI Training/serving architecture to accelerate our ability to deliver **Generative AI** in the Technology Operate and Production Management Team. Key Responsibilities: Research, develop, and optimize large and small language models for various applications. Fine-tune pre-trained models using domain-specific datasets. Implement, evaluate, and improve NLP pipelines to enhance model performance. Design efficient training workflows and optimize inference performance. Required: Experience with LLMs (e.g., **GPT**, LLaMA, Falcon) and SLMs. Knowledge of vector databases, **RAG** (**Retrieval-Augmented Generation**), and **prompt engineering**. Strong coding skills in Python and familiarity with relevant libraries (Transformers, Hugging Face, **LangChain**).
Qualcomm Apr 25, 2025 *Senior Machine Learning Engineer, Generative AI*	(Job ID: 4213957095) We are seeking an experienced Machine Learning Engineer specializing in **Generative AI** to join our core AI team. The ideal candidate will be responsible for designing, developing, and deploying cutting-edge **Generative AI** solutions, with a focus on **Large Language Models** (LLMs), **Retrieval-Augmented Generation** (**RAG**), and Intelligent agent systems. Key Responsibilities: Design and implement **RAG**-based solutions to enhance **LLM** capabilities with external knowledge sources. Develop and optimize **LLM** fine-tuning strategies for specific use cases and domain adaptation. Create robust evaluation frameworks for measuring and improving model performance. Build and maintain agentic workflows for autonomous AI systems. Preferred: Experience with popular LLM frameworks (**LangChain**, **LlamaIndex**, Transformers). Knowledge of **prompt engineering** and chain-of-thought techniques.
The Cigna Group Aug 26, 2023 *Senior AI Engineer*	(Job ID: 3716737940) You will be joining a team transforming healthcare and improving the lives of the millions of members we serve. We leverage cutting edge AI and ML algorithms to develop solutions for **automated document processing** and **customer service chat bots**. We are looking for Sr AI Engineers with strong engineering, full stack expertise to build the best fit solutions leveraging **Large Language Models** (LLMs) and **Generative AI** solutions. The work you do will impact millions of customers, members, and employers that rely on Cigna every day. Responsibilities: Build scalable software solutions using **LLM**s and other ML models to solve challenges in healthcare. Build enterprise grade AI solutions with focus on privacy, security, fairness. Build extensible API Integrations, low code UI/UX solutions, to extract information from sources, integrate with **GPT-4**, receive insights and make them available in intuitive, high performing dashboards.
Panel B: User Postings	
Dentsu International Sep 22, 2024 *Content Writer*	(Job ID: 4032072829) Generating SEO articles from assigned topics without plagiarism. Edit and proofread content to ensure accuracy, consistency, trustworthiness and adherence to brand guidelines and client requirements. Produce quality writing that connects with target audiences, client objectives, and brand. Complete writing assignments within project specifications and deadlines. Follow the content outline as assigned by SEO specialists or project managers. Requirements: Bachelor's degree in Marketing, Journalism, Mass Communications, Finance, Economy, or a related field. 1–2 years of experience in web content writing. Excellent writing, editing, and proofreading skills. Familiarity with **Generative AI such as ChatGPT or Claude** to assist in content writing **is a plus**.

Continued on next page

Table IA16: **Example Job Postings: Generative AI Integrators and Users** (continued)

Posting Details	Job Description (Excerpt)
TE Connectivity Dec 6, 2024 *Lead, Digital Finance Center of Excellence*	(Job ID: 4092101285) To lead the Digital Finance Center of Excellence (COE) by driving excellence in financial reporting, data visualization, and process optimization. This role aims to ensure accurate, timely, and insightful financial analysis to support strategic decision-making. Responsibilities: Take ownership of projects related to Tableau dashboards, ensuring improvements and innovations are consistently introduced. Utilize data visualization skills to derive actionable insights and provide strategic recommendations for financial decision-making. Desired Candidate Profile: MBA degree in Finance. Proficiency in Tableau (must have) and Power BI (preferred). Working knowledge of SQL. 7–10 years of professional experience in finance technology-related roles. **Have used prompt engineering in the past to automate repeated tasks. (Use of ChatGPT, Co-pilot).**

Table IA17: **Pre-ChatGPT Generative AI Integrator Job Postings**

This table presents six examples of pre-ChatGPT job postings classified as Generative AI integrator roles. All postings were published before the public release of ChatGPT (November 30, 2022).

Posting Details	Job Description (Excerpt)
Klaviyo May 8, 2022 *Senior Software Engineer, Data Science, Creative Guidance*	(Job ID: 3106244632) The team recently launched Showcase, an algorithmically selected library of high-performing email and SMS campaigns sent with Klaviyo that users can search through for creative inspiration for their next campaign. The team is actively conducting a number of experiments to help streamline and simplify the creative process. We are also working on a subject line generator powered by **GPT-3** that enables users to quickly generate subject line ideas for email campaigns. As a Senior Software Engineer on the Creative Guidance Team, you will primarily be responsible for prototyping and testing new features that utilize NLP models like **GPT-3** to generate copy for users. You will also be responsible for helping us improve and expand our Showcase product, in addition to designing and productionalizing new features with the intent of adding value to users along their creative journey. Technologies we use: Python, React, Typescript, **GPT-3**, Amazon Web Services.
CarMax Aug 9, 2022 *Senior Manager, App Sys*	(Job ID: 319026809600000000002) Online search at CarMax is used by people looking to buy their next car. Users can be anywhere in their shopping journey, from initial consideration to closer to taking the decision. Search is one of the most important steps in user satisfaction and a sale. As the largest retailer of used cars, CarMax has a lot of data about cars and user needs. We are re-imagining online search to support users in their complex car shopping journey. We are doing that through some of the most advanced technologies like cloud, **Open AI**, machine learning, natural language processing and personalization. Our goal is to help users spend more time finding than searching for the right car, through highly personalized search experiences. Read more about how CarMax is using **Open AI** to build car research tools. The position partners with a number of business areas including SEO, Online Platform, DevOps, Marketing, and plays a pivotal role in building online buying, as CarMax disrupts the industry for the second time.
Udemy Oct 11, 2022 *Senior Software Engineer, Deep Learning*	(Job ID: 3305492524) Udemy is looking for a software engineer to join our new Deep Learning team. The Deep Learning team is focused on exploring novel applications of deep learning and **foundational models** (e.g., **GPT-3**, OPT-175B, BLOOM) to improving the learning experience. In this role, you will design, build, and integrate scalable production systems, platforms, and tools that leverage deep learning. You will work in a wide stack including big data and streaming technologies as part of a cross-functional agile team that includes software engineers, machine learning engineers, data scientists, designers, and product managers.
Adobe Sep 15, 2021 *Sr Applied ML Scientist, Creative Co-Pilot*	(Job ID: 2716598081) The Intelligent Services group's latest mission is to build the Creative **Copilot**, an intelligent assistant that will help Adobe customers create original, high quality content. Creative **Copilot** is a system using transformers to assist creators with their workflows by providing a variety of next step recommendations. Building on earlier work on the creative assistant and latest advances in transformer technology (e.g. the GitHub **Copilot**), the group is currently training machine learning models on large-scale data sets from Adobe applications and assets. The goal is to ultimately have the **Copilot** present high-quality visual recommendations unobtrusively within Adobe applications. Familiarity with the latest developments in transformers (**GPT**, HuggingFace transformer models) and applications.
ServiceNow Jul 27, 2022 *Senior Research Scientist, Natural Language Processing*	(Job ID: 3191015085) The Human-Machine Interaction Through Language (HMITL) team at ServiceNow Research is exploring how natural language processing (NLP) and **large language models** (**LLM**) can revolutionize the way humans interact with software and knowledge systems. The topics include but are not limited to: task-oriented dialog and semantic parsing, interactive **code generation**, compute- and memory-efficient **LLMs**, **retrieval-augmented LLMs**, lightweight **LLM** adaptation for downstream tasks, using **LLM** for data synthesis, large context transformers, transformers with memory, controllable AI, **chain-of-thought** methods, integration with symbolic AI, trustworthy **natural language generation**; combatting hallucinations.
NVIDIA Nov 24, 2022 *Deep Learning Engineer, Games and Graphics*	(Job ID: 3375228181) We are looking for a Deep Learning Software Engineer to join the NVIDIA Developer Technology group (Devtech) and work with a team of experienced engineers on innovative uses of deep learning for games and graphics. The Devtech team works with NVIDIA researchers and leading game developers to bring cutting edge AI research from across NVIDIA and the industry to gamers and 3D professionals in high performance packages such as Deep Learning Super Sampling (DLSS). Use deep learning to enhance the quality and improve the performance of 3D rendering, and use **Generative AI** to create contents for games. Apply deep learning techniques to real-time 3D rendering, computer games and film production pipelines. Develop tools and products for real-time deep learning model development, PC games, VR and educational applications.

IA3 Instrumental Variable Estimates

We report two instrumental variable strategies that exploit distinct sources of exogenous variation in firms' Generative AI adoption. Section IA3.1 constructs a Bartik-style shift-share instrument from the evolution of frontier AI model capabilities. Section IA3.2 instruments adoption using geographic proximity to AI research universities.

IA3.1 LLM Capability Bartik

Our first instrument is a Bartik shift-share IV (Goldsmith-Pinkham et al., 2020) that interacts sector-level occupation structures with the evolving capabilities of large language models. As new LLMs are released, their capabilities align more closely with some occupation tasks than others, so sectors concentrating employment in occupations whose tasks overlap with LLM capabilities face a larger supply-side shock to the feasibility of Generative AI integration. The instrument assigns each firm the AI capability exposure of its NAICS sector, combining predetermined occupation employment shares from the Bureau of Labor Statistics with time-varying text similarity between occupation descriptions and LLM capability descriptions.

A. Text Alignment

We measure how closely each occupation's work content aligns with the cumulative capabilities of large language models in each quarter.

On the occupation side, we use the O*NET 30.0 database, which provides structured text descriptions for each 6-digit SOC occupation. Our main specification uses Detailed Work Activities (DWA), which describe workplace behaviors at finer granularity than task statements; O*NET maps each occupation's tasks to a standardized set of DWA descriptions.

On the model side, we use the Epoch AI large-scale AI models database, which catalogs AI systems trained with substantial compute, with metadata for publication date, demonstrated capabilities, and application domain. We restrict to models in the Language or Multimodal domains published after 2010; the cumulative set grows each quarter as new systems are released. Each record contains a Task field listing demonstrated capabilities (e.g., "question answering," "summarization," "code generation") and a Domain field for the application area.

To compute alignment, we represent both occupation and model texts as TF-IDF vectors over a shared vocabulary fitted on all texts, with 10,000 features, unigrams and bigrams, English stop words removed, and sublinear term-frequency scaling. For each quarter t, the cumulative AI capability frontier is the average $\bar{\mathbf{m}}_t$ of the TF-IDF vectors of all models

released up to t. The raw alignment between occupation s and the frontier is the cosine similarity:

$$\tilde{A}_{st} = \cos(\mathbf{d}_s, \bar{\mathbf{m}}_t)$$

where $\mathbf{d}_s$ is the TF-IDF vector of occupation s's O*NET text. We impose a monotonicity constraint:

$$A_{st} = \max_{t' \leq t} \tilde{A}_{st'}$$

so alignment can only increase over time, reflecting the cumulative nature of AI capabilities: once demonstrated, a capability persists in subsequent quarters. The constraint also smooths transient fluctuations in the similarity measure from changes in the mix of newly released models.

B. Instrument Construction

The instrument for firm i in quarter t is:

$$Z_{it} = \sum_{s \in \mathcal{S}} \omega_{s,t-1}^{g(i)} \cdot A_{st} \tag{5}$$

where $\mathcal{S}$ is the set of 6-digit SOC occupations, $\omega_{s,t-1}^{g(i)}$ is the lagged employment share of occupation s in firm i's three-digit NAICS industry $g(i)$ from the prior-year BLS Occupational Employment and Wage Statistics survey, and A_{st} is the text alignment defined above. The instrument is standardized to zero mean and unit variance before estimation.

We use sector-level employment shares from BLS rather than each firm's own workforce composition. All firms in the same NAICS sector receive identical occupation shares, so individual hiring decisions do not affect the instrument value. Cross-sectional variation comes from differences in sector occupation structures: a technology firm and a manufacturing firm receive different shares because their sectors employ different occupation mixes. Time-series variation comes from A_{st} as new models shift the alignment frontier differentially across occupations.

Each panel in Table 4 varies the alignment text inputs. Panel A aligns O*NET Detailed Work Activities against Epoch AI task-label descriptions. Panel B aligns the same O*NET text against Epoch AI task-and-domain descriptions, which add the application domain to the task categories. The two panels test robustness to the alignment text representation.

C. Exclusion Restriction

The exclusion restriction requires that the Bartik instrument affects capital structure only through Generative AI adoption. Two features support this. First, the occupation shares come from BLS at the three-digit NAICS level, not from individual firms; a firm's

decision to hire more software engineers does not change its instrument value, since only the aggregate occupation structure of its industry determines the share component. Using lagged BLS shares (each year's quarters use the prior year's survey) further ensures the shares are predetermined, severing the link between firm-level workforce choices and the instrument. Second, the time-varying component, LLM capability alignment, reflects the pace of frontier AI research set by the global research community, not by any individual firm's adoption decision.

Firm fixed effects absorb permanent across-sector differences in instrument levels, and year-quarter fixed effects absorb aggregate movements in alignment common to all firms. Identification comes from differences across sectors in how quickly AI capabilities align with their occupations as new models are released. One remaining concern is that sectors whose occupations align closely with LLM capabilities may face correlated leverage trends from other technology-driven forces. The two text-representation panels in Table 4 provide a partial check: if results hold across alignment measures, they are less likely to reflect a spurious correlation specific to one text representation.

IA3.2 Geographic Proximity to AI Research

Our second instrument uses geographic proximity to AI research universities. Firms headquartered near major AI research institutions face a larger local supply of Generative AI talent after ChatGPT's release, increasing adoption through a labor supply channel. The instrument is cross-sectional: headquarters location relative to research universities reflects historical factors unrelated to current financing decisions, and firm fixed effects absorb any time-invariant effect of location.

A. Publication Data

We collect AI and computer science publications from the Dimensions database, which indexes scientific publications primarily through Crossref DOIs and publisher partnerships.[19] We count publications in leading AI and computer science venues and journals over 2017–2022, the window beginning with the transformer architecture and ending before ChatGPT's release.

[19]Dimensions provides comprehensive coverage of AI venues published through established channels such as ACM (KDD, WWW, SIGIR), IEEE (CVPR, ICCV), and the ACL Anthology (ACL, EMNLP, NAACL), as well as major journals (JMLR, TACL, PAMI, and others). Conference proceedings published through independent platforms, such as NeurIPS, ICML, and ICLR, may be partially covered since these venues do not always deposit metadata with Crossref. Any undercoverage is uniform across institutions and does not bias the cross-sectional instrument. As a robustness check, we reconstruct the instrument using OpenAlex, the successor to the Microsoft Academic Graph, which provides broader coverage of independently published conference proceedings. The results are qualitatively similar.

Each publication–author–institution combination constitutes one observation. For each US institution, the GRID identifier (from the Global Research Identifier Database) classifies institutions by type. We restrict the instrument to education institutions, excluding corporate research laboratories (Google, Microsoft, Adobe, IBM Research, Nvidia, Meta, and others), government facilities (Army Research Laboratory, NIST), and nonprofits (Allen Institute for AI, SRI International). Several corporate laboratories rank among the most productive AI research organizations, but we exclude them because corporate publication activity may directly reflect firms' own AI investments, violating the exclusion restriction. Table IA18 reports the top-75 education institutions by paper count, alongside a selection of excluded corporate laboratories for comparison.

B. Instrument Construction

We rank universities by research output R_u and retain the top 75. The instrument for firm i is:

$$Z_{it} = \begin{cases} \log\left(\sum_{u \in \mathcal{U}} \mathbb{1}\left[d_{iu} \leq \bar{d}\right] \cdot R_u\right) \times \text{Post}_t & \text{if at least one } u \in \mathcal{U} \text{ has } d_{iu} \leq \bar{d}, \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

where $\mathcal{U}$ is the set of top-ranked universities, d_{iu} is the great-circle distance between firm i's headquarters and university u, R_u is research output at university u, and $\text{Post}_t = \mathbb{1}[t \geq 2023\text{Q1}]$. We set $\bar{d} = 250$ miles; the instrument equals zero for firms with no top-ranked university within this radius. Headquarters locations come from SEC filings.

We construct two variants of R_u. The first counts only first-author papers, proxying primary research contribution and reducing double-counting from multi-institution collaborations. The second counts all unique papers. Both variants count at the paper level: a paper with multiple authors at the same institution contributes one unit. Firm fixed effects absorb time-invariant location effects, and industry-by-year-quarter fixed effects absorb time-varying industry shocks, so identification comes from within-industry geographic variation in proximity to AI research.

C. Exclusion Restriction

The exclusion restriction requires that geographic proximity to AI research universities affect capital structure only through Generative AI adoption. Two concerns merit discussion. First, proximity could proxy for local economic conditions or agglomeration effects. Firm fixed effects absorb permanent location characteristics, including local industry composition, venture capital access, and historical university ties; the remaining variation is the differential effect of proximity after ChatGPT's release, difficult to attribute to pre-existing agglomeration. Second, firms near AI universities might differ in unobservable ways

that independently affect capital structure. Industry-by-year-quarter fixed effects absorb any industry-level trends correlated with geography. Identification thus comes from within-industry differences in distance to AI research, after versus before ChatGPT's release.

Table IA18: **Top AI Research Institutions**

Table IA18 lists the top-75 US education institutions by AI/CS publication count used to construct the geographic proximity instrument, alongside selected corporate laboratories excluded from the instrument. Publications are from the Dimensions database, covering leading AI and computer science venues and journals, 2017–2022. Papers = total publication count (used in the paper-weighted variant). Institutions are classified using the Global Research Identifier Database (GRID); only education institutions enter the instrument. GRID identifiers are available in the replication package.

#	Institution	Papers	#	Institution	Papers	#	Institution	Papers
1	Carnegie Mellon	969	26	Texas A&M	181	51	Brown	93
2	Stanford	552	27	Duke	178	52	Ohio State	93
3	UIUC	500	28	NYU	175	53	Wash. U. St. Louis	92
4	UC Berkeley	492	29	U. Virginia	172	54	Rochester Inst. Tech.	91
5	MIT	409	30	U. Central Florida	171	55	UT Dallas	90
6	U. Maryland	345	31	U. Buffalo, SUNY	160	56	George Mason	85
7	USC	343	32	U. Rochester	155	57	Indiana U.	84
8	UT Austin	331	33	Stony Brook	153	58	UT Arlington	82
9	Georgia Tech	322	34	Virginia Tech	151	59	Caltech	82
10	Rutgers	306	35	Northeastern	148	60	UC Davis	81
11	UCLA	293	36	UC Merced	143	61	Emory	76
12	Cornell	291	37	U. Minnesota	136	62	Temple	74
13	UC San Diego	286	38	Boston U.	135	63	UC Santa Cruz	73
14	U. Washington	280	39	U. Pennsylvania	133	64	Stevens Inst. Tech.	70
15	U. Michigan	277	40	Princeton	132	65	U. Kentucky	68
16	Johns Hopkins	271	41	U. Wisconsin	127	66	Oregon State	68
17	Harvard	260	42	UNC Chapel Hill	125	67	NC State	67
18	Arizona State	253	43	Notre Dame	122	68	U. Albany, SUNY	61
19	Michigan State	225	44	Northwestern	106	69	Worcester Polytech.	60
20	UMass Amherst	221	45	UC Riverside	105	70	UT San Antonio	57
21	Penn State	212	46	U. Florida	104	71	Syracuse	55
22	UIC	210	47	UC Irvine	102	72	U. Arizona	54
23	Columbia	193	48	Rensselaer	101	73	U. Georgia	54
24	Purdue	186	49	Rice	99	74	Washington State	54
25	U. Pittsburgh	183	50	UC Santa Barbara	95	75	Missouri S&T	54
Selected corporate laboratories (excluded from instrument)								
	Google	846		Microsoft	655		Adobe Systems	415
	IBM Research	345		Nvidia	213		Meta Platforms	154
	Intel	109		Allen Inst. for AI	89			

IA4 Ex-ante Generative AI Exposure

IA4.1 Ex-ante Exposure Measurement

Following Eisfeldt et al. (2023), we build a firm-level measure of Generative AI exposure: we classify whether each occupational task in O*NET can be performed more efficiently with large language models (LLMs), aggregate task-level scores to occupations, and then to firms by weighting occupations by each firm's workforce composition using Revelio Labs position data. Since we study market reactions to the ChatGPT release on November 30, 2022, our exposure measures must reflect information available to investors *before* that date; we therefore draw task descriptions from O*NET Database Version 27.1, released in November 2022 shortly before ChatGPT became public, and classify tasks using `gpt-3.5-turbo`, a model whose capabilities roughly match those of the original ChatGPT.

IA4.2 Occupation-level Exposure

Each of the 19,265 tasks in O*NET 27.1 is assigned to one of four exposure categories based on how much an LLM could reduce the time needed to complete it: E1 (direct, score = 1.0), E2 (indirect via LLM-powered software, = 0.5), E3 (image-dependent, = 0.0), or E0 (no exposure, = 0.0). We query the OpenAI API using `gpt-3.5-turbo` with temperature set to zero, random seed fixed at 66, maximum output tokens capped at 100, and JSON response format enabled. We adopt the classification prompt, rubric, and few-shot examples from Eisfeldt et al. (2023), with one modification: we add a third few-shot example illustrating the E2 category because in pilot runs the model rarely assigned the E2 label without a worked demonstration, tending to classify ambiguous cases as either E1 or E0.[20] We also randomize the order of few-shot examples across tasks to mitigate ordering effects. Across two independent runs with identical random seeds, label distributions are nearly identical (Run 1 / Run 2: E0 62.3%/62.2%, E1 22.4%/22.3%, E2 14.4%/14.6%, E3 0.9%/0.9%). As a robustness check, classifying O*NET 30.0 tasks using GPT-5.1 yields E0 58.4%, E1 26.3%, E2 13.6%, E3 1.6%; the higher direct-exposure share reflects the broader capabilities of frontier models.

[20] The first two few-shot examples are from Eisfeldt et al. (2023); the third draws on a task they report as E2 in their Table I, but the reasoning text is ours. The complete prompts are available in the replication code.

IA4.3 Firm-level Aggregation

We aggregate task scores to occupations by simple averaging, $X_O = \frac{1}{|T_O|} \sum_{t \in T_O} X_t$, where T_O is the set of tasks belonging to occupation O and $X_t \in \{0, 0.5, 1\}$ is the exposure score for task t. We then map 8-digit O*NET-SOC codes to 6-digit SOC codes using equal weights across sub-occupations. Firm-level exposure is the employment-weighted average of occupational exposures, $X_f = \sum_{o \in O_f} w_{fo} X_o$, where w_{fo} is occupation o's share of firm f's workforce; employment shares come from Revelio Labs position records as of 2022Q3, the quarter immediately preceding the ChatGPT release.

Table IA19 reports summary statistics for our occupation-level Generative AI exposure scores.

Table IA19: **Occupation-Level Generative AI Exposure**

Table IA19 reports summary statistics for occupation-level Generative AI exposure scores at the 6-digit SOC level. Each occupation's score is the simple average of its task-level exposure scores from O*NET 27.1, classified using `gpt-3.5-turbo`. For comparison, Eisfeldt et al. (2023) report a mean of 0.234 and standard deviation of 0.209.

Variable	Obs.	Mean	S.D.	p10	Median	p90
Occupation-Level Exposure	798	0.276	0.263	0.000	0.182	0.667